

*Robert Krimmer, Melanie Volkamer,
Bernhard Beckert, Véronique Cortier, Ardita Driza Maurer, David Duenas-Cid,
Jörg Helbach, Reto Koenig, Iuliia Krivonosova, Ralf Küsters, Peter Rønne,
Uwe Serdült, Oliver Spycher (Eds.)*

Fourth International Joint Conference on Electronic Voting

E-Vote-ID 2019

1-4 October 2019, Lochau/Bregenz, Austria

Co-organized by:

*Tallinn University of Technology
Ragnar Nurkse Department of Innovation and Governance
Karlsruhe Institute of Technology
Institute of Applied Informatics and Formal Description Methods
E-Voting.CC GmbH
Competence Center for Electronic Voting and Participation
Gesellschaft für Informatik
German Informatics Society, SIG SEC/ECOM
Kastel
Competence Center for Applied Security Technology*

PROCEEDINGS

Robert Krimmer, Melanie Volkamer,
Bernhard Beckert, Véronique Cortier, Ardita Driza Maurer,
David Duenas-Cid, Jörg Helbach, Reto Koenig, Iuliia Krivosova,
Ralf Küsters, Peter Rønne, Uwe Serdült, Oliver Spycher (Eds.)

4th Joint International Conference on Electronic Voting

E-Vote-ID 2019

1-4 October 2019, Lochau/Bregenz, Austria

**Co-organized by the Tallinn University of Technology,
Karlsruhe Institute of Technology, E-Voting.CC, Gesellschaft für
Informatik and Kastel**



Proceedings E-Vote-ID 2019
TalTech Press

ISBN 978-9949-83-473-0

Volume Editors

Prof. Dr. Robert Krimmer
Tallinn University of Technology
Ragnar Nurkse Department of Innovation and Governance
Akadeemia tee 3
12618 Tallinn
Estonia
robert.krimmer@taltech.ee

Prof. Dr. Melanie Volkamer
Karlsruhe Institute of Technology
Institute of Applied Informatics and Formal Description Methods
Kaiserstr. 89
76131 Karlsruhe,
Germany
melanie.volkamer@secuso.org

Bernhard Beckert
Karlsruhe Institute of Technology
E-mail: beckert@kit.edu

Reto Koenig
Bern University of Applied Sciences
E-mail: reto.koenig@bfh.ch

Véronique Cortier
Laboratoire Lorrain de Recherche en
Informatique et ses Applications
E-mail: veronique.cortier@loria.fr

Iuliia Krivonosova
Tallinn University of Technology
E-mail: Iuliia.krivonosova@taltech.ee

Ardita Driza-Maurer
Zentrum für Demokratie Aarau/Zurich
University
E-mail: ardita.driza@sefanet.ch

Ralf Küsters
University of Stuttgart
E-mail: ralf.kuesters@sec.uni-stuttgart.de

David Duenas-Cid
Tallinn University of Technology /
Kozminski University
E-mail: david.duenas@taltech.ee

Peter Rønne
University of Luxembourg
E-mail: peter.roenne@gmail.com

Jörg Helbach
Rheinische Fachhochschule Köln
E-mail: joerg@helbach.info

Uwe Serdült
Ritsumeikan University / Centre for
Democracy Studies
E-mail: uwe.serdult@zda.uzh.ch

Oliver Spycher
Swiss Federal Chancellery
E-mail: spycher.oliver@gmail.com

This conference is co-organized by:



Tallinn University of Technology - Ragnar Nurkse
Department of Innovation and Governance



Karlsruhe Institute of Technology - Institute of Applied
Informatics and Formal Description Methods



E-Voting.CC GmbH - Competence Center for
Electronic Voting and Participation

Gesellschaft
für Informatik



Gesellschaft für Informatik, German Informatics
Society, SIG SEC/ECOM

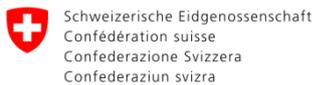


Kastel, Competence Center for Applied Security
Technology

Supported by:



Regional Government of Vorarlberg



Swiss Federal Chancellery

General Chairs

Krimmer, Robert (Tallinn University of Technology - Ragnar Nurkse Department of Innovation and Governance, Estonia)

Volkamer, Melanie (Karlsruhe Institute of Technology - Institute of Applied Informatics and Formal Description Methods)

Track on Security, Usability and Technical Issues

Cortier, Véronique

Laboratoire Lorrain de Recherche en Informatique et ses Applications, France

Beckert, Bernhard

Karlsruhe Institute of Technology, Germany

Küsters, Ralf

University of Stuttgart, Germany

PhD Colloquium

Koenig, Reto

Bern University of Applied Sciences, Switzerland

Driza-Maurer, Ardita

Zentrum für Demokratie Aarau/Zurich University, Switzerland

Track on Administrative, Legal, Political and Social Issues

Serdült, Uwe

Ritsumeikan University, Japan / Centre for Democracy Studies, Switzerland

Duenas-Cid, David

Tallinn University of Technology, Estonia / Kozminski University, Poland

Organizational Committee

Traxler, Gisela

E-Voting.CC, Austria (Main Contact)

Track on Election and Practical Experiences

Helbach, Jörg

Rheinische Fachhochschule Köln, Germany

Spycher, Oliver

Swiss Federal Chancellery, Switzerland

Outreach Chairs

Rønne, Peter

University of Luxembourg, Luxembourg

Krivosova, Iuliia

Tallinn University of Technology, Estonia

Preface

This volume contains papers presented at E-Vote-ID 2019, the Fourth International Joint Conference on Electronic Voting, held during October 1-4, 2019, in Bregenz, Austria. It resulted from the merging of EVOTE and Vote-ID and counting up to 15 years since the first E-Vote conference in Austria. Since the first conference in 2004, over 1000 experts have attended the venue, including scholars, practitioners, authorities, electoral managers, vendors and PhD Students. The conference collected the most relevant debates on the development of Electronic Voting, from aspects relating to security and usability through to practical experiences and applications of voting systems, also including legal, social or political aspects, amongst others; turning out to be an important global referent in relation to this issue.

Also, this year, the conference consisted of:

- Security, Usability and Technical Issues Track
- Administrative, Legal, Political and Social Issues Track
- Election and Practical Experiences Track
- PhD Colloquium, Poster and Demo Session on the day before the conference

E-VOTE-ID 2019 received 45 submissions, being, each of them, reviewed by 3 to 5 program committee members, using a double blind-review process. As a result, 23 papers were accepted for this volume, representing 51% of the submitted proposals. The selected papers cover a wide range of topics connected with electronic voting, including experiences and revisions of the real uses of E-voting systems and corresponding processes in elections.

We would also like to thank the German Informatics Society (Gesellschaft für Informatik) with its ECOM working group and KASTEL for their partnership over many years. Further we would like to thank the Swiss Federal Chancellery for their kind support. Special thanks go to the members of the international program committee for their hard work in reviewing, discussing, and shepherding papers. They ensured the high quality of these proceedings with their knowledge and experience.

October 2019
Bregenz, Austria

Robert Krimmer
Melanie Volkamer
Bernhard Beckert
Véronique Cortier
Ardita Driza Maurer
David Duenas-Cid
Jörg Helbach
Reto Koenig
Iuliia Krivososova
Ralf Küsters
Peter Rønne
Uwe Serdült
Oliver Spycher

Table of Contents

Introductory paper

- E-Voting – an overview of the development in the past 15 years and current discussions 1
Robert Krimmer, Melanie Volkamer and David Duenas-Cid

Invited Speaker

- The Academic Debate on Electronic Voting in a Socio-Political Context . . 17
Anne-Marie Oostveen and Peter van den Besselaar

Risk Limiting Audit and its applications

- Auditing Indian Elections 37
Vishal Mohanty, Chris Culnane, Philip Stark and Vanessa Teague
- Risk-Limiting Tallies 53
Peter Y. A. Ryan, Wojciech Jamroga, Peter Roenne and Philip Stark
- VAULT: Verifiable Audits Using Limited Transparency 69
Josh Benaloh, Philip Stark and Vanessa Teague

The Swiss Voting Experience

- The Swiss Postal Voting Process and its System and Security Analysis . . . 90
Christian Killer and Burkhard Stiller
- How Do the Swiss Perceive Electronic Voting? Social Insights from a Qualitative Field Survey 106
Emmanuel Fragniere, Sandra Grèzes and Randolph Ramseyer
- The Swiss Post/Scytl transparency exercise and its possible impact on internet voting regulation 122
Ardita Driza Maurer

Models and schemes

- Security models for everlasting privacy 140
Panagiotis Grontas, Aris Pagourtzis and Alexandros Zacharakis
- Improvements in Everlasting Privacy: Efficient and Secure Zero Knowledge Proofs 155
Thomas Haines and Clementine Gritti
- Modeling Requirements Conflicts in Secret Ballot Elections 171
Aaron Wilson

User Experience Design for E-Voting: How mental models align with security mechanisms	187
<i>Marie-Laure Zollinger, Verena Distler, Peter Roenne, P. Y. A. Ryan, Carine Lallemand and Vincent Koenig</i>	
Internet Voting Governance: Canada and Estonia	
The curse of knowledge? Does having more technology skills lead to less trust towards ivoting?	204
<i>Mihkel Solvak and Robert Krimmer</i>	
Online Voting in a First Nation in Canada: Implications for Participation and Governance	208
<i>Brian Budd, Chelsea Gabel and Nicole Goodman</i>	
How increasing use of Internet voting impacts the Estonian election management	226
<i>Iuliia Krivonosova, Robert Krimmer, David Duenas-Cid and Radu Antonio Serrano Iova</i>	
Analysis of Deployed Systems	
The Danish Party Endorsement System	229
<i>Carsten Schürmann and Alessandro Bruni</i>	
Online Voting in Ontario Municipal Elections: A Conflict of Legal Principles and Technology?	245
<i>Anthony Cardillo, Nicholas Akinyokun and Aleksander Essex</i>	
Election Integrity and Electronic Voting Machines in 2018 Georgia (USA)	261
<i>Kellie Ottoboni and Philip Stark</i>	
E-Voting, practical approaches	
Cybersecurity and Electoral Integrity: The Case of Ukraine, 2014-present	278
<i>Beata Martin-Rozumilowicz and Thomas Chanussot</i>	
GI Elections with POLYAS: a Road to End-to-End Verifiable Elections	293
<i>Ralf Kuesters, Tomasz Truderung, Melanie Volkamer, Bernhard Beckert, Achim Brelle, Rüdiger Grimm, Nicolas Huber, Michael Kirsten, Jörn Müller-Quade, Maximilian Noppel, Kai Reinhard, Jonas Schwab, Rebecca Schwerdt and Cornelia Winter</i>	
Pakistan’s Internet Voting Experiment	295
<i>Hina Binte Haq, Ronan McDermott and Syed Taha Ali</i>	

Implementing a public security scrutiny of an online voting system: the Swiss experience	311
<i>Jordi Puiggali</i>	

Attacks and Security Requirements in Practice

UnclearBallot: Automated Ballot Image Manipulation	327
<i>Matthew Bernhard, Jeremy Wink, Kartikeya Kandula and J. Alex Halderman</i>	

Election Manipulation with Partial Information	346
<i>Michelle Blom, Peter Stuckey and Vanessa Teague</i>	

On practical aspects of coercion-resistant voting systems	362
<i>Jan Willemson and Kristjan Krips</i>	

PhD Colloquium

How can Internet Voting be implemented in Portuguese elections? A comparison with Estonia	379
<i>Marlon Freire</i>	

How elections with Internet voting are administered? The case of the 2019 Parliamentary elections in Estonia	381
<i>Iuliia Krivonosova</i>	

Modelling Strategic Capabilities in Tamarin: Pros and Cons	383
<i>Damian Kurpiewski</i>	

Voters' Understanding of the Coercion Mitigation Mechanism in Selene ..	385
<i>Marie-Laure Zollinger</i>	

DEMO Session

Features of Polys Blockchain Voting: A Survey	388
<i>Roman Alyoskin</i>	

Unsupervised electronic voting machines and methods	390
<i>Promitheas Christophides</i>	

Verifiability and security of Scytl's online voting system	392
<i>Jordi Puiggali</i>	

Program Committee

Marta Aranyossy	Corvinus University of Budapest
Myrto Arapinis	The University of Edinburgh
Jordi Barrat i Esteve	eVoting Legal Lab
Bernhard Beckert	Karlsruhe Institute of Technology
Josh Benaloh	Microsoft
David Bismark	Votato
Nadja Braun Binder	University of Zurich
Christian Bull	The Norwegian Ministry of Local Government and Regional Development
Susanne Caarls	Election Consultant
Gianpiero Catozzi	UNDP
Veronique Cortier	CNRS, Loria
Ardita Driza Maurer	Zentrum für Demokratie Aarau/Zurich University
David Duenas-Cid	Tallinn University of Technology / Kozminski Uni- versity
Noella Edelmann	Danube University Krems
Ulle Endriss	University of Amsterdam
Aleksander Essex	University of Western Ontario
Joshua Franklin	National Institute of Standards and Technology
David Galindo	University of Birmingham
Micha Germann	University of Bath
J Paul Gibson	Mines Telecom
Kristian Gjøsteen	Norwegian University of Science and Technology
Nicole Goodman	University of Toronto
Rajeev Gore	The Australian National University
Ruediger Grimm	University of Koblenz
Rolf Haenni	Bern University of Applied Sciences
Thomas Haines	Queensland University of Technology
Thad Hall	MPR
Jörg Helbach	Rheinische Fachhochschule Köln
Toby James	University of East Anglia
Tarmo Kalvet	Tallinn University of Technology
Reto Koenig	Berne University of Applied Sciences
Iuliia Krivosova	Tallinn University of Technology
Ralf Kuesters	University of Stuttgart
Oksana Kulyk	ITU Copenhagen
Leontine Loeber	University of East Anglia
Beata Martin-Rozumilowicz	IFES
Anu Masso	Tallinn University of Technology
Ronan Mcdermott	Medis
Juan Mecinas	Universidad de las Américas Puebla
Magdalena Musiał-Karg	Adam Mickiewicz University
András Nemaslaki	BME

Hannu Nurmi	University of Turku
Jon Pammett	Carleton University
Olivier Pereira	UCLouvain
Goran Petrov	OSCE
Stéphanie Plante	University of Ottawa
Josep M ^a Reniu	University of Barcelona
Peter Roenne	SnT, University of Luxembourg
Mark Ryan	University of Birmingham
P. Y. A. Ryan	University of Luxembourg
Peter Sasvari	National University of Public Service
Steve Schneider	University of Surrey
Berry Schoenmakers	Eindhoven University of Technology
Carsten Schuermann	IT University of Copenhagen
Uwe Serdült	Centre for Research on Direct Democracy
Oliver Spycher	Swiss Federal Chancellery
Philip Stark	University of California, Berkeley
Vanessa Teague	The University of Melbourne
Tomasz Trudering	University of Trier
Priit Vinkel	State Electoral Office of Estonia
Kåre Vollan	Quality AS
Bogdan Warinschi	University of Bristol
Roland Wen	The University of New South Wales
Gregor Wenda	BMI
Jan Willemson	Cybernetica
Peter Wolf	International IDEA
Michael Yard	IFES

Introductory paper

E-Voting – an overview of the development in the past 15 years and current discussions

Robert Krimmer¹ [0000-0002-0873-539X], Melanie Volkamer² [0000-0003-2674-4043] and David Duenas-Cid^{1,3} [0000-0002-0451-4514]

¹ Tallinn University of Technology, Akadeemia tee 3, 12618 Tallinn, Estonia
robert.krimmer,david.duenas@taltech.ee

² Karlsruhe Institute of Technology, Kaiserstr. 89, 76131 Karlsruhe, Germany
melanie.volkamer@kit.edu

³ Kozminski University, Jagiellonska 57/59, 03-301 Warsaw, Poland

Abstract. This opening article introduces the Fourth International Joint Conference on Electronic Voting and, on the occasion of the 15 years since the first E-Vote conference in Austria, presents an analysis of the network of co-authorships based on the books published by the Electronic Voting Conference Series. The goal of the analysis is to provide an overview of the development of the network of authors involved in the conference and to give some insights on the internal dynamics of collaboration within the field. Its comprehension sheds light on the creation of influence, internal norms and performance of the publications, enlarging the knowledge on the field and highlighting the contribution of the conferences on its development.

Keywords: E-Vote-ID Conference Series, Network of co-authorships, Social Network Analysis

1 Looking back in time: a network analysis

On the occasion of celebrating the first 15 years since starting the Electronic Voting Conference Series, we introduce an opening chapter to conduct a retrospective analysis of evolution of the Conference with regards to its impact on creating a community of scholars interested in this field. After these years elapsing, the network of scholars working on it has increased, but are no more certain of its extent than the personal perceptions derived from interactions between colleagues. In our humble opinion, one of the main successes of these 15 years of conferences on electronic voting has been the contribution to the consolidation of a field of research and to the creation of a regular meeting point for the researchers interested in the topic. We assume that these meetings helped to create new connections that consolidated in common projects and publications. In order to shed some light on the impact of the conference, we conducted an analysis of the network of scholars created amongst the participants in the conference, based on their collaborative work. As Cugmas et al. [6] cite, there are different ways to collaborate scientifically [25], but most are invisible [26]: collaboration involving a division of labor, service collaboration, providing access to research equipment, the transmission of know-how, mutual stimulation, and trusted assessment. Amongst the “visible” and formal ones, the same authors [5] reference six indicators of collaborative relationships between scientists [12]: co-authorship; shared editing of publications; joint supervision of PhD projects; writing research proposals together; participation in formal research programs; and jointly organizing scientific conferences.

In this analysis we will focus on one particular type of collaboration, the network resulting from the collaboration in common published papers. Networks of co-authorships have been analyzed in diverse fields such as digital libraries [29], organizational studies [2] or healthcare [7], and they normally share a common background sustained by the idea that collaboration is a basic element of academic life as it increases the productivity of the researcher [27] and impact of his/her research [10], helping to increase the citation rate for publications [31] and the overall quality of research [15]. This process helps to weave a wider network of interconnected researchers together, contributes to sharing resources that are relevant in the field, such as information, common understanding and knowledge [28], and serves as a way of introduce new members (PhD students, for example) to the field of research [1].

On the other hand, analyzing social networks helps us describe and understand the underlying processes in social relations. As Fitzhugh and Butts [9] describe, the structural influence of countless social processes rarely occur in isolation, but in relation to the social network in which they are inserted. Networks allow us to understand, for example, how interaction processes are influenced by the homophilic tendency to relate to people whom we are similar to [30], how networks expand and create relational clusters [35] or how weak connections can bring new opportunities for existing relationships [11]. In this case, formalization of informal relations (as a co-authored publication) may be considered a representation of a common set of interests and goals, crystalizing in a common joint project (a research or a publication), fostering the exchange of information and the improving the resulting work. The accumulative process of creating and developing networks of collaboration helps to create internal dynamics and implicit

rules in a given environment. Every organizational environment (field) is defined by a non-written set of standards and norms, behavioral patterns and dynamics and certain forms of capital that help us navigate within it, and to understand how the environment works when you are submerged in it [8]. Following this approach, a field must be conceptualized as a configuration of relationships between nodes represented in a network and the positions that those nodes happen to occupy [8]. The use and analysis of networks helps reveal the internal dynamics of a given field, and to understand the different symbolic positions and strategies followed by their components. Assuming that the Electronic Voting Conference Series is a representation of the electronic voting community, and understanding the network of co-authorships as a representation of internal dynamics in this field of research, analyzing the resulting network gives us clues for detecting how this particular field of knowledge works and to celebrate the contribution of the conference series to the creation, development and consolidation of this field.

2 Methodology

We analyze the network of co-authorships on books resulting from the E-Vote, Vote-ID and E-Vote-ID Conference series in this paper. The data regarding publications, co-authorships and citations has been extracted from the profile in Google Scholar which compiles all the information on the Electronic Voting Conference Series¹. The data was extracted using the Harzing Publish or Perish (version 6)² software, a program for looking up scholarly citations and calculating impact metrics. The data was analyzed using Gephi (Version 0.9.2)³, a program for manipulating and visualizing network graphs.

The data presented here consists of an exploratory analysis of evolution of the network of co-authorships⁴ as an indicator for creating a structure of scholars around (and thanks to) the electronic voting conference series and its insights. The relations between nodes (edges) are considered as undirected (two nodes are related, being irrelevant the direction of the relation), and the analysis (Number of nodes connected, Betweenness, PageRank, Detection of hubs) is done taking the cumulative network (2004-2018) as an analytical frame.

3 Analysis

A total of 14 books has been published resulting from the presentations held in the conference [3, 13, 14, 16–24, 32, 33], including a total of 228 articles and 628 collaborators, between authors and editors, since 2004. These publications have been cited up

¹ scholar.google.com/citations?user=KsBkbjkAAAAJ&hl=en (Last Access, 7 May 2019)

² Available at: harzing.com/resources/publish-or-perish?source=pop_6.46.6370.7005 (Last Access, 7 May 2019)

³ Available at: gephi.org/ (Last Access, 7 May 2019)

⁴ Co-Editions are included as Co-Authorships.

to 2.764 times, an average of 184,3 times per year and 12,1times per article⁵. Co-authoring articles seems to be the most common practice in the conference series, since just 20% of the articles published are by a single author.

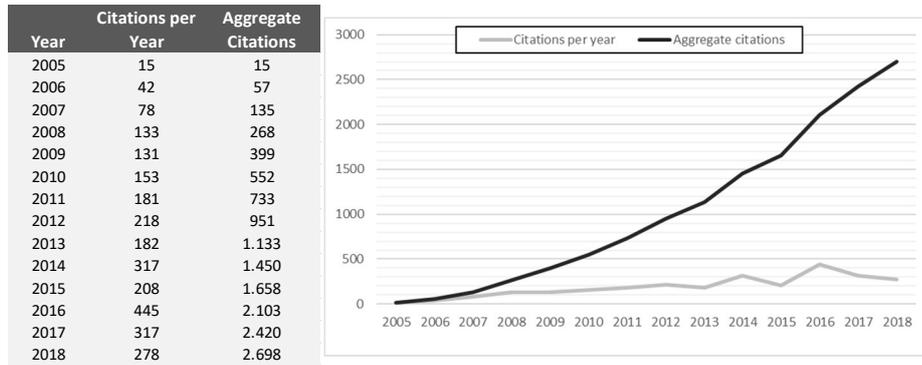


Fig. 1. Evolution in the number of citations

To create the co-authorship network, we identified every author and publication as a node in the network and every collaboration as an edge connecting the different nodes involved: authors and publications. As a result, we have a network composed of 616 nodes and 1.518 edges⁶. The process creating the network (see **Fig. 2**) was gradual, up until the current configuration of the network. The nodes are of different sizes depending on their *degree* (the number of nodes they are connected to), highlighting those nodes that are more “popular” in terms of shared publications. The ten most popular nodes are those which had been more active in publishing, co-authoring and involved in co-editing some of the conference proceedings (see **Table 1**).

Author	Nodes
M Volkamer	87
R Krimmer	58
V Teague	44
PYA Ryan	41
P Vora	33
C Schürmann	27
S Schneider	25
R Goré	22
A Essex	21
S Popoveniuc	21

Table 1. Nodes with higher degree in the network

⁵ Last Update of citations, April 2019, Publish or Perish

⁶ Note that the total number of nodes does not correspond to the sum of authors and publications, due to the fact many researchers published more than one text, and, therefore, they are counted just once.

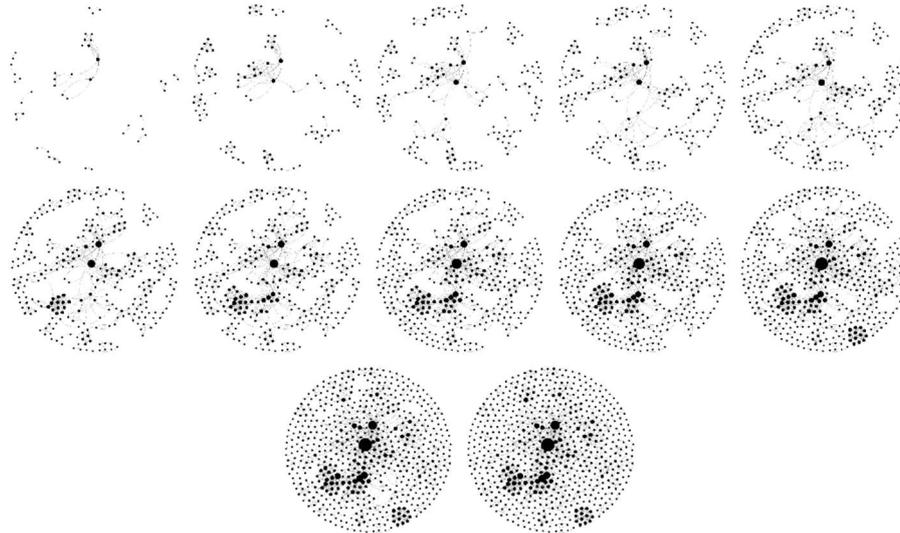


Fig. 2. Cumulative evolution of the network of co-authorships 2004-2018⁷

The degree is considered as a measure of centrality, as it allows to identify relevant nodes in a network according to their “popularity”. Even so, centrality in a network is not only about being connected to many nodes, but also being connected to relevant nodes. The relevance of connections is analyzed using the *betweenness centrality* measure, which collects information on how often a node appears on the shortest path between different networks or, in other words, which are the nodes that are better placed in a network. The position occupied is directly connected to the amount of information that one node can accumulate in the network and is directly related with holding symbolic power within the field. There is a tendency to correlate between having more nodes and being placed in the center of the network, but there is also a dependency on the nodes to which one is connected. For example, a strategic connection with well-connected nodes gives you a more central position than a large number of poorly connected nodes. In the case we are analyzing, shared publications with relevant nodes can influence the position in the network (see **Fig. 3**).

⁷ Every network corresponds to one Edition of the Electronic Voting Conference Series. The nodes are weighted according to the number of nodes they are connected to, the larger the number, the greater the number of contacts.

Author	Betweenness
M Volkamer	251.578
R Krimmer	18.892
C Schürmann	9.710
V Teague	8.330
PYA Ryan	8.240
P Vora	6.278
J Benaloh	4.978
R Goré	4.485
J Barrat	4.340
J Helbach	3.516

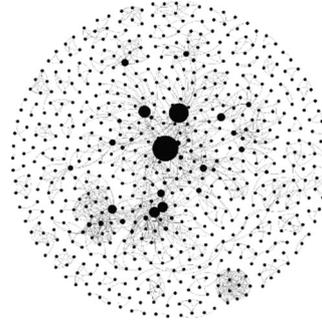


Fig. 3. Betweenness Centrality

Betweenness allows us to determine which is the symbolic center of the network, and it relates to the capacity to influence other nodes in the network. Another estimate of the capacity for influencing is given by the *PageRank* measure, which estimates the relevance of one node in relation with the nodes with whom it is connected, giving a wider approach of the capacity to influence a node and being used, for example, for detecting priorities in networks [34]. Having high values of PageRank in a network represents a combined indicator of the position of a certain node in the network in relation to the capacity of the node to create a hub around it.

Node	PageRank
M Volkamer	0,013
R Krimmer	0,010
C Schürmann	0,007
PYA Ryan	0,005
V Teague	0,005
P Vora	0,004
J Willemson	0,004
RM Alvarez	0,004
L Loeber	0,004
J Puiggalí	0,004

Table 2. PageRank Distribution

The measures presented allowed us to detect the most relevant individual nodes within the network. In the case analyzed, the most relevant nodes are researchers who participated in the Conference Series in many times and managed to gather people around them with whom they shared publications. But networks extend beyond the individual position of nodes by including internal dynamics of group behavior. Combining related nodes with dense relationships with themselves create hubs of interrelated nodes, and detecting these hubs is particularly relevant for understanding the internal standards and norms for a network. In the case analyzed here, using the *Modularity* algorithm provided by Gephi, we detected a set of nodes that create hubs or clusters of interrelated nodes due to their greater interconnectedness with themselves than with the rest of the network [4], meaning, in this case, a greater collaboration in co-authoring papers (see

Fig. 4). Not all the hubs have been colored to avoid interference in visualizing the network, choosing just those that represented larger numbers of scholars. The singular analysis of each hub provides interesting results but, in order to make the analysis easier, only the most paradigmatic typologies will be presented.

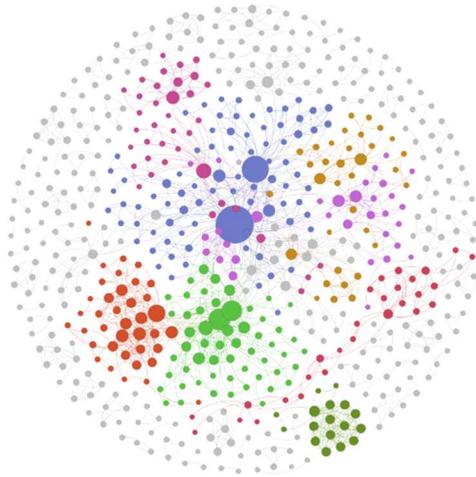


Fig. 4. Division of the network into hubs

Depending on the internal data of each hub detected, we could find some patterns of similarity in the features and outcomes of different hubs. The main element of similarity and difference relates to the typology of members composing the hub. Two different types of hub are detected in this regard, those which are created based on a diverse group of researchers interacting due to common interests and a second group of hubs that are created based on similarity of the origin of their members. Amongst the first group of hubs, the differences amongst them arise from the position they occupy in the network and the internal patterns of functioning. Amongst the second one, the difference is based on the country/environment of origin of their members. Therefore, four main geographically based groups were detected: An Estonian hub, a Swiss one, an Australian one and one based on their relation with the private vendor Scytl. The first three (Estonian, Swiss and Australian) are composed of scholars who are researching the countries where internet voting systems are deployed, showing the research attractiveness of pioneering examples. The Scytl hub is created around the publications done by researchers link to a private vendor that offers its services worldwide. Their publications, therefore, are generally connected to the places where they deliver their services.

	Central hub I	Central hub II	Horizontal hub	Central hub III	Peripheral hub
Number of members	47	34	28	18	16
Average Degree	8,43	11,09	11,70	6,61	5,63
Median Degree	4	8	12	6	5
Density	0,07	0,14	0,24	0,13	0,16
Average Betweenness	1.114	840	354	862	38
Average Closeness	0,31	0,30	0,25	0,30	0,42
Average Year of first publication	2.010	2.013	2.011	2.012	2.009
Number of Publications	28	19	12	13	12
Average cites per year	1,44	2,32	0,98	0,73	1,76
Average number of authors	3,18	3,89	4,22	2,38	2,50
	Estonian hub	Australian hub	Swiss hub	Scytl hub	Total
Number of members	14	10	9	9	392
Average Degree	6,14	7,30	9,11	6,56	6,24
Median Degree	4	5	7	5	4
Density	0,19	0,34	0,26	0,29	0,01
Average Betweenness	264	449	465	477	342
Average Closeness	0,26	0,23	0,27	0,26	0,35
Average Year of first publication	2.014	2.016	2.012	2.012	2.011
Number of Publications	10	7	12	8	222
Average cites per year	3,86	1,18	1,86	1,99	1,51
Average number of authors	2,80	3,14	2,92	2,88	2,81

Table 3. Main clusters detected in the network

A group of three hubs has been defined as central hubs (I, II and III). The reason is that they all share a high degree of betweenness as a common salient feature. They gather a relevant number of researchers who occupy the centrality of the network, but they differ in the internal structure of the hub. While central hubs I and III display a clear priority (Melanie Volkamer and Robert Krimmer for Central hub I -see **Fig. 5**- and Carsten Schürmann and Jordi Barrat for Central hub III -**Fig. 7**), Central hub II (**Fig. 6**) presents a softer and more horizontal distribution, with a higher average interconnectedness (represented by a higher average and median degree). The second central hub also joined the Conference Series, on average, at a later stage, when some internal dynamics had already been created and this might also help create a different relational structure based on commonality of interests⁸.

⁸ For reading the forthcoming figures: in order to ease the comprehension and comparison of data, average values for all the network have been calculated (Index 1). The data for each hub is presented in comparison with the average value.

Name	Degree	Betweenness	PageRank
M Volkamer	87	25.158	0,013
R Krimmer	58	18.892	0,010
R Grimm	20	908	0,004
S Neumann	19	805	0,003
J Helbach	12	3.516	0,002
K Reinhard	11	1.314	0,003
N Meissner	10	987	0,002
C Feier	9	200	0,001
M Traxl	9	145	0,001
A Prosser	8	4	0,001

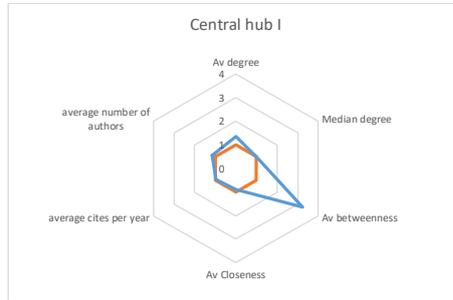


Fig. 5. Values for Central Hub I⁹

Name	Degree	Betweenness	PageRank
V Teague	44	8.330	0,005
PYA Ryan	41	8.240	0,005
S Schneider	25	1.951	0,002
R Wen	19	694	0,002
Z Xia	19	200	0,002
J Heather	17	78	0,002
J Benaloh	16	4.978	0,002
C Culnane	15	257	0,002
JA Halderman	14	1.026	0,002
D Demirel	13	304	0,002

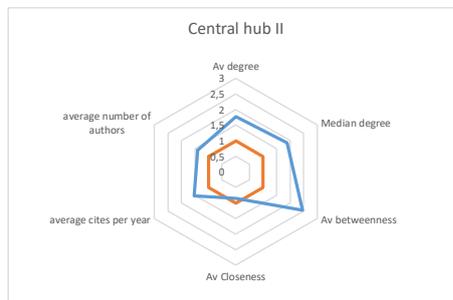


Fig. 6. Values for Central Hub II

Name	Degree	Betweenness	PageRank
C Schürmann	27	9.710	0,007
J Barrat	17	4.340	0,004
O Pereira	11	1.458	0,003
B Goldsmith	6	0	0,001
D Jandura	6	0	0,001
J Turner	6	0	0,001
M Chevallier	6	0	0,001
N Kersting	6	0	0,001
NB Binder	6	0	0,001
R Sharma	6	0	0,001

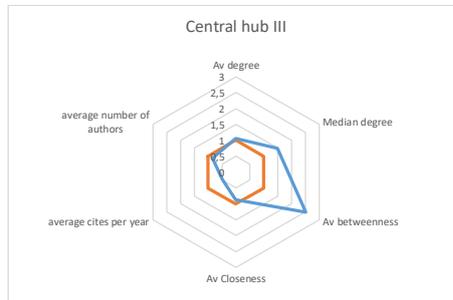


Fig. 7. Values for Central Hub III

Regarding the other two non-geographically based hubs, named as peripheral and horizontal hubs, these portray a different perception. The peripheral hub (Fig. 8) was more active in the past and now only some of their members still attend the conference and, probably, they will jump across to a different hub in the future if their publishing activity continues with active members of the community. This “lethargy” moved the position of this hub to a non-central space (low betweenness) but with adequate proximity

⁹ Note that the scale for Fig. 5 is different

and impact on the overall network. The horizontal hub (**Fig. 9**) follows the pattern presented for the Central hub II, representing a group of scholars with intense internal dynamics and mutual interconnectedness (higher degree and average number of authors) which might take on more central positions in the future. We would like to remind the reader that networks are live and change, and that this description is just a fixed snapshot in time of a temporary reality that will evolve in years to come.



Fig. 8. Values for Peripheral Hub



Fig. 9. Values for Horizontal Hub

Finally, the geographically based hubs (Estonia **-Error! Reference source not found.-**, Switzerland **-Fig. 11-**, Australia **-Fig. 12-** and Scytl **-Fig. 13-**) tend to have smaller numbers of researchers and are structured around the link to a certain environment relating to use or research on Electronic Voting. In the Estonian, Swiss and Australian cases, the connection is based on implementation of Electronic Voting in their electoral systems, with many publications connected to the observation of use and development, impact, improvements or conditions, amongst others. In the Estonian case, the performance in terms of citations per publication works clearly better than average, showing the general interest of the community by the application of Internet Voting in a real context.

For the other case, the hub relates more to creating a cluster of expertise around their professional practice and the experiences and research carried out using the cases where Scytl offers its services.

Name	Degree	Betweenness	PageRank
J Willemson	19	2.301	0,004
S Heiberg	17	541	0,003
P Vinkel	9	84	0,002
T Martens	6	656	0,002
A Koitmae	5	37	0,001
D Duenas-Cid	5	37	0,001
I Krivososova	5	37	0,001
A Parsovs	3	0	0,001
I Kubjas	3	0	0,001
K Krips	3	0	0,001

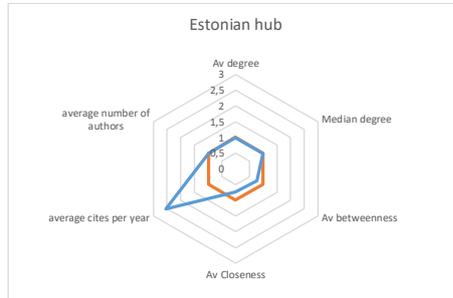


Fig. 10. Values for the Estonian Hub

Name	Degree	Betweenness	PageRank
R Haenni	19	718	0,003
RE Koenig	19	2.481	0,003
O Spycher	12	274	0,003
E Dubuis	9	12	0,002
P Locher	7	43	0,001
A Driza-Maurer	6	657	0,002
A Weber	4	0	0,001
G Taglioni	4	0	0,001
M Schlapfer	2	0	0,001

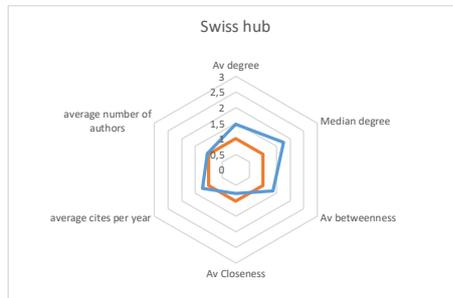


Fig. 11. Values for the Swiss Hub

Name	Degree	Betweenness	PageRank
R Goré	22	4.485	0,003
D Pattinson	12	4	0,001
M Tiwari	9	2	0,001
MK Ghale	7	1	0,001
LB Moses	5	0	0,001
R Levy	5	0	0,001
T Meumann	5	1	0,001
B Beckert	3	0	0,001
JE Dawson	3	0	0,001
E Lebedeva	2	0	0,001

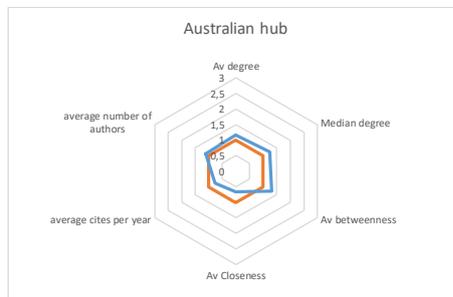


Fig. 12. Values for the Australian Hub

Name	Degree	Betweenness	PageRank
J Puiggalí	17	2.750	0,004
S Guasch	11	799	0,002
J Cucurull	7	631	0,002
A Fornós	5	1	0,001
J Lladós	5	1	0,001
Jl Toledo	5	1	0,001
D Galindo	3	0	0,001
M Soriano	3	109	0,001
V Morales-Rocha	3	1	0,001

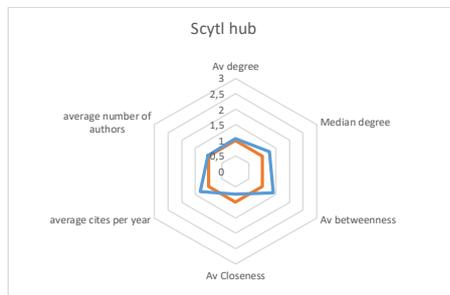


Fig. 13. Values for the Scytll Hub

4 Conclusions

This analysis allowed a better understanding of internal patterns of functioning for the field created regarding the Electronic Voting Conference Series. Every field or research creates a set of internal and or written standards and norms, by standardizing behaviors and activities which in turn become accepted in the community. In this case, starting with the assumption that the network of co-authorships represents a formal representation of the real network of relationships in the field, through analyzing it, we were able to detect which are the more central nodes (researchers) in the field and the priorities being created at the same time. A common publication can be understood to be an indicator of social capital and symbolic power, as it means a common interest in publishing together (social capital) and a form of academic respect of the researcher's potentials (symbolic power). The maximization and dispersion of co-authored publications involves a greater centrality in the network and, in practical terms, increasing the authority on the topic in the research field.

Detecting hubs helps visualize the existence of different relational environments and secondary centralities within the same network. The presence of smaller but closely interrelated networks within the network shows the existence of partial priorities and differentiated publishing strategies or patterns of collaboration. Based on the results, the hubs detected present different results in relation with to position, interconnection and performance, increasing the knowledge available on the field and its dynamics.

These results, finally, open some new possibilities for future research on meta-analysis of the Electronic Voting field, by questioning which is the best possible strategy for making an impact in the community. Centrality seems to work well for gaining authority in the field, but the singularity of certain cases and experiences performs very well in terms of academic impact or citations. This research could be widened in soon by expanding analysis to the citation network, deepening the impact of certain pieces for creating and consolidating of the field and, by extension, the knowledge of Electronic Voting.

Acknowledgements

The work of Krimmer and Duenas received support from ETAG personal research grant 1361.

References

1. Abbasi, A. et al.: Identifying the effects of co-authorship networks on the performance of scholars: A correlation and regression analysis of performance measures and social network analysis measures. *J. Informetr.* 5, 594–607 (2011). <https://doi.org/10.1016/j.joi.2011.05.007>.
2. Acedo, F.J. et al.: Co-authorship in management and organizational studies: An empirical and network analysis. *J. Manag. Stud.* 43, 5, 957–983 (2006). <https://doi.org/10.1111/j.1467-6486.2006.00625.x>.
3. Alkassar, A., Volkamer, M.: E-Voting and Identity: First International Conference, VOTE-ID 2007, Bochum, Germany, October 4-5, 2007, Revised Selected Papers. Springer-Verlag Berlin Heidelberg, Bochum (2007). <https://doi.org/10.1007/978-3-540-77493-8>.
4. Blondel, V.D. et al.: Fast unfolding of communities in large networks. *J. Stat. Mech. Theory Exp.* (2008). <https://doi.org/10.1088/1742-5468/2008/10/P10008>.
5. Cugmas, M. et al.: Scientific Co-Authorship Networks. In: Doreian, P. et al. (eds.) *Advances in Network Clustering and Blockmodeling*. p. 460 Wiley-Blackwell, New Jersey (2019).
6. Cugmas, M. et al.: The stability of co-authorship structures. *Scientometrics*. (2016). <https://doi.org/10.1007/s11192-015-1790-4>.
7. E Fonseca, B. de P.F. et al.: Co-authorship network analysis in health research: Method and potential use. *Heal. Res. Policy Syst.* 14, 34, 1–10 (2016). <https://doi.org/10.1186/s12961-016-0104-5>.
8. Emirbayer, M., Johnson, V.: Bourdieu and organizational analysis. *Theory Soc.* (2008). <https://doi.org/10.1007/s11186-007-9052-y>.
9. Fitzhugh, S.M., Butts, C.T.: Patterns of co-membership: Techniques for identifying subgraph composition. *Soc. Networks.* 55, 1–10 (2018). <https://doi.org/10.1016/j.socnet.2018.03.006>.
10. Gazni, A., Didegah, F.: Investigating different types of research collaboration and citation impact: A case study of Harvard University’s publications. *Scientometrics.* 87, 251–265 (2011). <https://doi.org/10.1007/s11192-011-0343-8>.
11. Granovetter, M.: The Strength of Weak Ties: A Network Theory Revisited. *Sociol. Theory.* (2006). <https://doi.org/10.2307/202051>.
12. De Haan, J.: Authorship patterns in Dutch sociology. *Scientometrics.* 39, 2, 197–208 (1997). <https://doi.org/10.1007/BF02457448>.

13. Haenni, R. et al.: E-Voting and Identity 5th International Conference, VoteID 2015, Bern, Switzerland, September 2-4, 2015, Proceedings. Springer-Verlag Berlin Heidelberg, Bern (2015). <https://doi.org/10.1007/978-3-319-22270-7>.
14. Heather, J. et al.: E-Voting and Identity 4th International Conference, Vote-ID 2013, Guildford, UK, July 17-19, 2013, Proceedings. Springer-Verlag Berlin Heidelberg, Guildford (2013). <https://doi.org/10.1007/978-3-642-39185-9>.
15. Katz, J.S., Martin, B.R.: What is research collaboration? *Res. Policy.* 26, 1, 1–18 (1997). [https://doi.org/10.1016/S0048-7333\(96\)00917-1](https://doi.org/10.1016/S0048-7333(96)00917-1).
16. Kiayias, A., Lipmaa, H.: E-Voting and Identity Third International Conference, VoteID 2011, Tallinn, Estonia, September 28-20, 2011, Revised Selected Papers. Springer-Verlag Berlin Heidelberg, Tallinn (2011). <https://doi.org/10.1007/978-3-642-32747-6>.
17. Krimmer, R. et al.: Electronic Voting - Third International Joint Conference, E-Vote-ID 2018 Bregenz, Austria, October 2–5, 2018 Proceedings. Springer, Cham (2018). <https://doi.org/10.1007/978-3-319-68687-5>.
18. Krimmer, R.: Electronic Voting 2006 2nd International Workshop. *GI Lecture Notes in Informatics*, Bregenz (2006).
19. Krimmer, R. et al.: Electronic Voting First International Joint Conference, E-Vote-ID 2016, Bregenz, Austria, October 18-21, 2016, Proceedings. Springer Cham, Bregenz (2016).
20. Krimmer, R. et al.: Electronic Voting Second International Joint Conference, E-Vote-ID 2017, Bregenz, Austria, October 24-27, 2017, Proceedings. Springer Cham, Bregenz (2017). <https://doi.org/10.1007/978-3-319-68687-5>.
21. Krimmer, R., Grimm, R.: Electronic Voting 2008 (EVOTE08) 3 rd International Conference. *GI Lecture Notes in Informatics*, Bregenz (2008).
22. Krimmer, R., Grimm, R.: Electronic Voting 2010 (EVOTE2010) 4 th International Conference. *GI Lecture Notes in Informatics*, Bregenz (2010).
23. Krimmer, R., Volkamer, M.: 6th International Conference on Electronic Voting EVOTE2014 28–31 October 2014, Lochau/Bregenz, Austria. TUT Press, Bregenz (2014).
24. Kripp, M. et al. eds: 5th International Conference on Electronic Voting 2012. Köllen Druck + Verlag GmbH, Bonn (2012).
25. Laudel, G.: *Interdisziplinäre Forschungskoperation: Erfolgsbedingungen der Institution Sonderforschungsbereich*. Sigma Ed, Berlin (1999).
26. Laudel, G.: What do we measure by co-authorships? *Res. Eval.* 11, 1, 3–15 (2002).
27. Lee, S., Bozeman, B.: The impact of research collaboration on scientific productivity. *Soc. Stud. Sci.* 35, 5, 673–702 (2005). <https://doi.org/10.1177/0306312705052359>.
28. Li, E.Y. et al.: Co-authorship networks and research impact: A social capital perspective. *Res. Policy.* 42, 9, 1515–1530 (2013). <https://doi.org/10.1016/j.respol.2013.06.012>.
29. Liu, X. et al.: Co-authorship networks in the digital library research community. *Inf. Process. Manag.* 41, 1462–1480 (2005). <https://doi.org/10.1016/j.ipm.2005.03.012>.

30. McPherson, M. et al.: Birds of a Feather: Homophily in Social Networks. *Annu. Rev. Sociol.* 27, 1, 415–444 (2001). <https://doi.org/10.1146/annurev.soc.27.1.415>.
31. Narin, F. et al.: Scientific co-operation in Europe and the citation of multinationally authored papers. *Scientometrics.* 21, 3, 313–323 (1991). <https://doi.org/10.1007/BF02093973>.
32. Prosser, A., Krimmer, R.: *Electronic Voting in Europe – Technology, Law, Politics and Society.* GI Lecture Notes in Informatics, Bregenz (2004).
33. Ryan, P., Schoenmakers, B.: *E-Voting and Identity* Second International Conference, VOTE-ID 2009, Luxembourg, September 7-8, 2009, Proceedings. Springer-Verlag Berlin Heidelberg, Luxembourg (2009). <https://doi.org/10.1007/978-3-642-04135-8>.
34. Wang, R. et al.: Discover community leader in social network with PageRank. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics).* (2013). <https://doi.org/10.1007/978-3-642-38715-9-19>.
35. Wasserman, S., Faust, K.: *Social Network Analysis: Methods and Applications.* Cambridge University Press, Cambridge (1994).

Invited Speaker

The Academic Debate on Electronic Voting in a Socio-Political Context

Anne-Marie Oostveen and Peter van den Besselaar

Vrije Universiteit Amsterdam, de Boelelaan, Amsterdam, The Netherlands
arie.oostveen@gmail.com - p.a.a.vanden.besselaar@vu.nl

Abstract. Computer-based voting as a field of research and societal debate emerged in the early 2000s. Starting in the ‘old democracies’ in Europe and North America, it has spread to other parts of the world. The question is whether research and the academic debate on electronic voting is related to the socio-political context in which it takes place. In order to examine this, we retrieved from the Scopus database all papers that relate to internet voting to answer the following research questions: Is there an increased scientific interest for e-voting in emerging democracies? Is the approach towards e-voting different between ‘old’ and ‘emerging’ democracies (i.e. technical, political, economical, social) and in terms of evaluation of e-voting (i.e. positive, negative)? We find that developed democracies have a more balanced approach in terms of disciplinary attention and in terms of evaluation of e-voting than the emerging democracies and the hybrid and authoritarian regimes. Africa deviates from this, with comparable substantial social science research being conducted on e-voting.

Keywords: Electronic Voting, Computer Voting, Risks, Benefits, Level of Democracy, Research System, Academic Discourse, Linguistic Analysis.

1 Introduction

The OSCE Office for Democratic Institutions and Human Rights (ODIHR) established in 1991, provides support, assistance, and expertise to promote democracy, rule of law, and human rights. One of its most important tasks is to observe elections to ensure that votes are cast by secret ballot or an equivalent free voting procedure, and that they are counted and reported honestly with the official results made public. “While holding an election does not equate to instant democracy, genuine and periodic elections that permit fair competition are fundamental to the democratic process” [28]. The ODIHR office initially concentrated on countries in transition, i.e. those countries that are emerging from a non-democratic past. But increasingly there have been invitations from longer-standing democracies to observe their elections as they face new electoral challenges: “For example, the introduction of new voting technologies poses potential challenges for transparency and accountability in any country where such technologies are being used or considered” [28]. Examples of such new

voting technologies are ballot-scanning technology, direct recording electronic voting systems (DRE), and remote electronic voting, also known as internet voting [29].

In the early years after the introduction of electronic voting, there was a clear tendency by politicians and civil servants to see the new technology mainly as a panacea to problems related to voting in democratic countries. E-voting was introduced as a means to reduce costs, increase voter turnout (especially amongst the young), potentially renew interest in the political system, and speed up the counting process as no more ballots needed to be counted manually. The benefits for voters were seen as a more efficient and user-friendly way to cast a ballot, with increased accessibility if the e-voting could be done remotely. It also made it more unlikely that people would accidentally spoil their ballot by for instance ticking two candidates instead of one.

According to a recent article by the BBC around 33 countries use some form of electronic voting [5]. The integrity of the machines has been questioned in some of them [37]. It was noted early on by computer-security experts that e-voting technologies are not as secure, accurate, reliable, or intuitive as suppliers promised [35, 20]. Furthermore, there are also concerns about the complexity of the processes (far from the comprehension of common voters), the actual effects on turnout rates (there is very little evidence of an increase), and higher costs [30]. Finally, there is a fundamental conflict between verification and keeping votes anonymous.

One of the most recent discussions about the security and accuracy of electronic voting machines is being held in the USA. For the past 17 years, many states and counties have conducted elections on machines that have been repeatedly shown to be vulnerable to hacking, errors, and breakdowns. Due to the suspicion over Russian interference in the 2016 Presidential election, there is a drive to replace the old outdated machines with new ballot-marking devices (BMD) that provide a paper trail in order to make the 2020 elections more secure and verifiable. But the new equipment built by the same companies is under scrutiny and there are concerns with the integrity of the paper trail [37]. Even in a country like Switzerland, where the government has always championed the use of e-voting to aid direct democracy and where citizens have been broadly in favour of e-voting, there are now opposition campaigners highlighting the vulnerabilities of the technology [30].

While currently the majority of e-voting countries use DRE voting machines, there is an increased interest in remote electronic voting offering people the chance to cast their vote online from a computer or mobile device. Compared to electronic voting machines used in the controlled environment of polling stations, the opportunity for attacks on the internet is much broader. As long as there have been high-stakes elections, there has been an interest to influence their outcome, be it in a legal manner through campaign contributions, advertising, or using political consultants, or in an illegal manner through ballot stuffing, voter intimidation, or bribery. The people interested in controlling the outcome of an election “could include political zealots or campaigns, but they might also include organised crime or even other countries with huge resources” [40]. Computer scientist David Dill points out that “there are many other threats, including voters who are not experts in computer security and may be easily fooled, and potential for corrupt insiders at companies that produce the internet voting software” [40]. Remote voting also increases the possibilities of coercion and vote buying.

Not only long-standing full democracies have adopted electronic voting. Newer, less established democracies and even authoritarian states are using or considering e-voting. For instance, the five major emerging national BRICS economies have all shown interest in e-voting, with Brazil and India leading the way for Russia, China, and South Africa, having both used e-voting for two decades. A growing trend towards electronic voting in developing democracies has previously been noted [1, 2] with some researchers arguing that the speed of implementation has been higher in the developing world than in established democracies, “especially in Latin America, with several countries such as Brazil, Venezuela, Argentina and Ecuador implementing e-voting methods” [33]. Different motivations can underpin the introduction of e-voting in emerging democracies. Some governments see it as a way to significantly reduce electoral fraud, human error, or a general lack of confidence in the electoral process, while others aim to increase turnout, reduce costs, or address the problematic distribution of electoral materials in more difficult to access rural areas. In authoritarian regimes (with often high levels of corruption), e-voting could possibly and relatively easily be used to reinforce control by the ruling party.

With the global adoption of e-voting technologies in countries with vastly different political systems and levels of democracy, we are interested in the question whether research and the academic debate on electronic voting is related to the political and social context in which it takes place. In order to examine this, we retrieved from Scopus all papers that relate to internet voting to answer the following research questions: Is there an increased scientific interest for e-voting in emerging democracies? And is the approach towards e-voting different between ‘old’ and ‘emerging’ democracies in terms of topics addressed (i.e. technical, political, economic, social) and in terms of evaluation of e-voting (i.e. positive, negative)? We will investigate whether there is a balanced research agenda into the benefits and risks of electronic voting in emerging democracies and more authoritarian regimes.

2 Background

2.1 E-voting in Old Democracies

A good example of a critical approach towards electronic voting was observed in the Netherlands. The country was one of the first to introduce electronic voting in polling stations in the late 1980s without much public debate. By the turn of the century almost the entire population used voting computers to cast their ballots in polling stations. Although at times, a few citizens, scholars, and politicians posed critical questions about the security, transparency, and verifiability of the e-voting systems, the government always dismissed these concerns.

This changed in 2006 when concerned citizens organized themselves and started a grassroots campaign named *Wij vertrouwen stemcomputers niet* (We do not trust voting computers). The core group of activists grew over the course of 18 months from 4 to 13 members, with different skills and backgrounds (e.g. software engineers, a social scientist, ICT consultants, hackers, a Freedom of Information specialist). The goal of the campaign was to promote, defend, and examine verifiable and transparent elections, with particular emphasis on the obstacles posed by electronic voting.

After hacking and examining the two electronic voting systems in use in the Netherlands and finding many security and other flaws the activists decided that one of their main objectives was to have the government abandon these unverifiable e-voting systems [16]. In order to accomplish this objective the group needed to convince both the general public and Members of Parliament of the shortcomings of the e-voting computers in use. They developed a new media and traditional media strategy to change public attitude towards e-voting systems and to influence public policy [26]. Within weeks of setting up the campaign the activists had put the security and verification problems of e-elections firmly on the political agenda, ultimately leading to the decertification of all voting computers in the Netherlands in 2007 [27]. In May 2008, the Dutch government decided that future elections in the Netherlands would only use paper ballots and red pencils, but software has since been used to count votes electronically. However, in 2017 the Netherlands also renounced the use of electronic ballot counting software due to an elevated threat of foreign interference [10].

The work done by the activists helped to create awareness in many other European countries (e.g. Germany, Belgium, Ireland, Italy). This resulted in heated political and public debate about the use of e-voting systems (in polling stations and remote) and led in some cases to the abolishment of voting computers and changes in electoral law. Not only in Europe could critical voices be heard more strongly. In the US scholarly work was conducted that similarly pointed out that the voting basics of ensuring one vote per voter, maintaining voter anonymity, accuracy of the vote, security of the system, and prevention of fraud could not all be maintained in an electronic voting system. In America, where voting machines were introduced some 15 years ago (with approximately 35,000 of them in use) there have been concerns over machines without a back-up paper trail (VVAT) misreading the vote. Machines used to tally results and program voting machines were found to be carrying software allowing remote access to system administrators; a massive security breach [39].

Despite persisting concerns, the possible introduction of e-voting is a recurring theme, even in countries that have decided against it over a decade ago. For instance, in January 2015 the desirability of e-voting was put back on the British political agenda when the Speaker of the House of Commons published a report on digital democracy, which concluded that “online voting has the potential greatly to increase the convenience and accessibility of voting” [13, 18]. The report stated that in the 2020 general election, secure online voting should be an option for all voters. It is clear that with technological change taking place at a rapid pace, some will argue that new developments such as biometric verification methods, stronger encryption, or blockchain technology make e-voting a more viable option.

2.2 E-voting in Emerging Democracies

While stable democracies started to become more reluctant in using e-voting, many incipient democracies have shown an increased interest in e-voting technologies. Especially ‘Third Wave Democracies’ have embraced e-voting, with either systems already implemented (e.g. Estonia, Nigeria) or being piloted on a smaller scale (Jordan, Venezuela). The third wave of democratization describes the global trend that

has seen more than 60 countries throughout Europe, Latin America, Asia, and Africa undergo some form of democratic transitions since Portugal's "Carnation Revolution" in 1974. Sometimes these democratic transitions are little more than transitions to semi-authoritarian rule.

These fairly 'new' or 'emerging' democracies often witness(ed) electoral malpractice and manipulation (ballot stuffing, fabrication of results), voter intimidation, and violence at polling stations. Many elections from around the world have been characterised by cases of rigging and fraud. Vote rigging is the process of interfering with the elections either to win as a candidate or to make an opponent lose. According to [6] "democracy has established itself as the dominant political system on the African continent; and as an integral part of this process, multi-party elections have emerged as the most legitimate route to political office. Yet, in recent years violence has increased in such elections". A recent study by Kewir and Gabriel (2018) using data from more than 50 African elections from 2011 to 2017, showed that almost all these elections had cases of electoral violence at some stage of the election [6]. Bjarnesen and Söderberg Kovacs point out that electoral violence is not limited to general and national elections: "In Sierra Leone, for example, several parliamentary by-elections at constituency level have generated high levels of violence, intimidation and insecurity" [6]. Other examples are elections in Uganda, which are often characterised by controversies with the government accused of intimidating opposition leaders and their supporters, including arrests and detention [8], and the 2008 elections in Kenya where systematic electoral fraud including vote-rigging in a third of all constituencies, stuffed ballot boxes (leading to a turnout of 115%), and election officials changing results had a decisive impact on the outcome of the elections. Immediately after the results were announced, violence broke out, which evolved into ethnic clashes leaving more than 1300 people dead and 600,000 displaced [7]. But electoral misconduct or violence can be seen all over the world. Romanian elections have been characterized by allegations of electoral bribes using food. In 2016 an MP went to prison for two years for bribing voters with 60 tons of roasted chicken [11]. In 2014, food was distributed to more than 6.5 million people during campaigns. In 2015 a senior Romanian minister was convicted of electoral fraud over a 2012 attempt to use bribes and forged ballot papers to swing a vote [19]. At the 2014 Iraqi parliamentary election, six different polling stations were hit by suicide bombers, leading to at least 27 deaths. Insurgent group Islamic State of Iraq and Syria threatened violence against Sunni Muslims who vote in the election [3]. In Turkey in 2015, President Erdoğan, was accused of planning to commit election fraud and several irregularities. According to the OSCE an increase in violence with physical attacks on party members, particularly in the southeast, "restricted some contestants' ability to campaign freely"[4].

People are bribed at polling stations with money, seeds, or medicine. Voter intimidation in Africa is also widespread, with 44% of those surveyed saying they are sometimes, often or always threatened with violence when voting [31]. Penar explains: "This makes people afraid of going to the polls and expressing an autonomous and personal choice of who they want to win" [23].

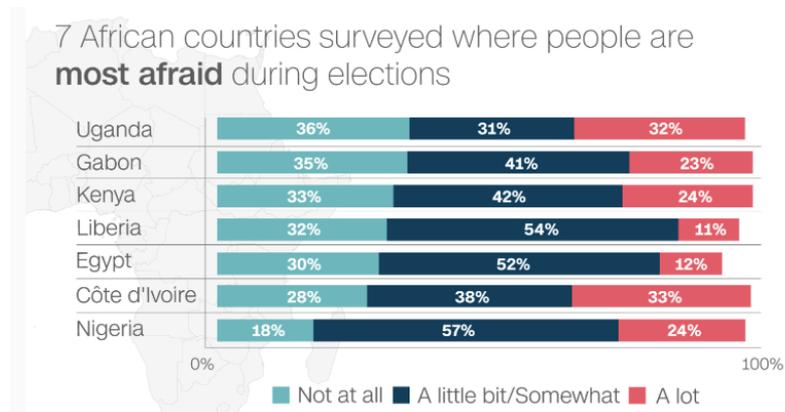


Fig. 1. Fear during elections. Source: Afrobarometer/CNN, 2016

Due to the high levels of violence, intimidation, and fraud, the turnout at elections is often low. E-voting systems are thought to provide a solution to these defective elections. However, the same inherent problems currently experienced in these repressive or corrupt regimes still remain with e-voting, such as the difficult-to-realise fundamental security requirements of privacy, verifiability, accountability, transparency, and coercion-resistance [22]. There have been plenty of real-life instances where there were concerns about the e-voting technologies in use. For instance, machines used in the 2017 elections in Venezuela allegedly inflated the actual turnout by at least a million votes, a claim rejected by the government [21]. Furthermore, in 2018 Iraq’s election commission ignored an anti-corruption body’s warnings about the credibility of electronic vote-counting machines used in the countries’ parliamentary election. Fraud allegations led to a manual partial recount of the ballots. “Concerns about the election count focus on discrepancies in the tallying of votes by the voting machines, mainly in the Kurdish province of Sulaimaniya and the ethnically-mixed province of Kirkuk, and on suggestions that the devices could have been tampered with or hacked into to skew the result” [34].

Another example of e-voting concerns are the machines used in India. For decades traditional voting in India had been blighted by the stuffing of ballot boxes by mobs hired by political parties. This changed with the introduction of electronic voting machines at the turn of the century [5]. Data from state elections show that e-voting machines used in polling stations had significantly reduced electoral fraud, had helped the poor and the weak to come out and vote, and made elections more competitive resulting in a decline of the winning margin and the vote share of the winning party [12]. However, not all analysis of the Indian voting system are positive. In 2010 a group of international researchers conducted the first independent, rigorous assessment of the security risks associated with Indian electronic voting machines and they found many vulnerabilities that the Election Commission of India’s (ECI) “Technical Experts Committee” had failed to pick up on [38]. As a final example we address the Democratic Republic of Congo, which used e-voting machines for the first time in 2018. These voting machines quickly became a source of contention in the Presiden-

tial election amid reports that they had not been thoroughly tested with a potential for misuse and serious security vulnerabilities [15].

Introducing e-voting in ‘new’ or ‘emerging’ democracies could have an additional barrier as the populations of these countries often have very low trust in their political and administrative systems. While “in most cases, established democracies have well-tested electoral practices that enjoy the overall confidence of their electorates, as well as pluralistic media that identify electoral shortcomings for public debate, independent judicial organs, and a generally robust civil society” this is different for emerging democracies [28]. Alvarez et al. [2] note that in many countries in Latin America public trust in elections and electoral authorities is very low: “the percentage of Latin Americans reporting that elections are free and fair in their countries is only 41 percent”. Similarly, a 2016 Afrobarometer study showed that in many African countries, corruption and a lack of transparency of election monitoring bodies all influence people's trust in the system. Many citizens are afraid when casting their vote. And 38% of Africans believe that votes are only sometimes, or even never counted [31]. According to Caarls [9] this will influence the implementation of e-voting systems. She argues that an e-voting system cannot be successfully adopted unless citizens trust their current (paper-based) political and administrative systems. In contrast to this, Penar [31] argues that electronic voting systems might actually be the way to increase trust in current political and administrative systems. He found that between 2011 and 2015 some African countries observed increased levels of trust in national electoral commissions, which according to him might have been a reflection of the reforms and technological advances in these countries. Namibia (the first African country to use e-voting), for example, introduced electronic voting in its 2014 general elections and a survey completed in the run-up to these elections showed a period marked by optimism about the implementation of electronic voting machines [31]. “In Namibia, the new electronic voting system appears to have boosted voters' trust, [...] with trust in their election commission at 74% [23]. This is in line with findings by Alvarez et al. [1] in Columbia where the proportion of respondents who declared to trust electronic voting was unusually high when compared to other international experiences. The authors explained that this “is probably related to the comparatively low degree of public confidence in elections in many countries in Latin America”.

Nevertheless, trust is a complicated concept and voters often base their trust in elections on perceived instead of provable security, privacy, and verifiability [25]. Küsters and Müller [22] argue: “In addition to the provable security a system provides, the level of security perceived by regular voters might be just as important and even more important for a system to be accepted”. We have pointed out in previous work that because of transparent election procedures, traditional paper-based voting in democratic countries satisfies security, privacy, and verifiability requirements, and are trusted by citizens [25]. We argued that: “In contrast, electronic voting systems are not transparent for the user, as the steps in the processing of the information cannot be observed. With electronic voting systems, public confidence in the election relies on trust in the organisers of the ballot, and in the technology (and technical experts) instead of on a transparent procedure”. The Office for Democratic Institutions and Human Rights [28] similarly noted that electronic voting technologies may pose perceived or real challenges to the transparency and accountability of an election process. Furthermore, they may influence perceptions of the security of the vote, with a poten-

tial impact on voter confidence. The level of trust may influence the decision to use e-voting systems and is therefore a crucial factor for e-voting to be successful. Küsters and Müller recommend keeping e-voting systems simple and comprehensible.

But while new digital tools can make it harder to tamper with votes, they can also create problems as they require electricity and technical know-how, something not to be taken for granted in rural areas throughout developing nations. This would possibly exclude certain groups within a given population from having their voices heard, an issue also emphasized by Caarls [9]. And while election processes could be improved by storing and counting votes digitally, issues with corruption may remain. "The servers on which the results data is stored may be owned by a friend of the president or a relative of an election commission. That happens" [23].

3 Methodology

We aim at comparing the reception of e-voting in relation to the level of democracy: do new and old democracies discuss it in different ways? As a proxy for this, we use the scholarly literature about e-voting, and with hindsight this is feasible: most papers are only nationally co-authored, implying that we do have national academic discourses on e-voting.

In order to do a meaningful comparison, we need to classify countries in homogeneous groups. A first attempt is using the Democracy Index compiled by the Economist Intelligence Unit (EIU), which uses a set of variables to produce a democracy ranking for 167 countries [14]. The Democracy Index is based on the following five categories: electoral process and pluralism; civil liberties; the functioning of government; political participation; and democratic political culture. Based on its scores on 60 different indicators within these categories, each country is then itself classified as one of four regime types: 'full democracy', 'flawed democracy', 'hybrid regime' and 'authoritarian regime' (Table 1).

Table 1. The Democracy Index compiled by the Economist Intelligence Unit (EIU)

Type of regime	Democracy score	Number of countries	% of countries	% of world population
Full democracies	$8 < s$	20	12	4.5
Flawed democracies	$6 < s \leq 8$	55	32.9	43.2
Hybrid regimes	$4 < s \leq 6$	39	23.4	16.7
Authoritarian regimes	$s \leq 4$	53	31.7	35.6

Several problems emerge, especially the fact that the boundaries between the types are rather arbitrary: countries with scores that differ only marginally are classified in different groups. For example, South Korea counts as full democracy with a score of 8.00, while Japan is a flawed democracy with a score of 7.99. The EIU acknowledges this, as it also uses a geographical classification (Table 2). The disadvantage of the geographical approach is that it may include countries of very different levels of democracy. As an example, the 'flawed democracies' include the USA as well as the Philippines and Mexico, which intuitively do not fit well in one class.

Table 2. Geographical classification. (Source: Economist Intelligence Unit, EIU)

Rank	Region	Countries	2006	2008	2010	2011	2012	2013	2014	2015	2016	2017	2018
1	N Am	2	8.64	8.64	8.63	8.59	8.59	8.59	8.59	8.56	8.56	8.56	8.56
2	W EU	21	8.6	8.61	8.45	8.4	8.44	8.41	8.41	8.42	8.4	8.38	8.35
3	Lat Am	24	6.37	6.43	6.37	6.35	6.36	6.38	6.36	6.37	6.33	6.26	6.24
4	Asia	28	5.44	5.58	5.53	5.51	5.56	5.61	5.7	5.74	5.74	5.63	5.67
5	C-E EU	28	5.76	5.67	5.55	5.5	5.51	5.53	5.58	5.55	5.43	5.4	5.42
6	SSA	44	4.24	4.28	4.23	4.32	4.33	4.36	4.34	4.38	4.37	4.36	4.36
7	MENA	20	3.54	3.48	3.52	3.62	3.73	3.68	3.65	3.58	3.56	3.54	3.54
	World	167	5.52	5.55	5.46	5.49	5.52	5.53	5.55	5.55	5.52	5.48	5.48

N Am = North America; W EU = Western Europe; Lat Am = Latin America & the Caribbean; Asia = Asia & Australasia; C-E EU = Central & Eastern Europe; SSA = Sub-Saharan Africa; MENA = Middle East & North Africa

We combine both approaches: the democracy score and the regional score. This leads to (geographical) groupings, which are relatively homogeneous in terms of their democracy score. Table 3 gives the country groupings, their average democracy score, and the variation of scores within the group. The Coefficient of Variation is about 0.2 in most of the groups. Some outliers were not included in the calculation, also because they had only a very low number of e-voting publications over the almost 20 year period.

Table 3. Country grouping

	# papers	share	Democracy score	
			Average	Variation
EU (North, West, South)	1352	50.6%	8.62	0.09
USA	517	19.4%	8.15	-
Japan, Singapore, Taiwan, S Korea	252	9.4%	7.70	0.04
China	249	9.3%	3.51	-
India & Pakistan	209	7.8%	4.68	0.22
Middle East and North Africa	186	7.0%	3.81	0.27
Australia, N.Z., Canada	178	6.7%	9.17	0.24
East EU and East Europe	137	4.9%	6.16	0.15
South America (excl. Chile)	104	3.9%	5.13	0.20
Indonesia & SE Asia	69	2.6%	5.02	0.19
Sub-Saharan Africa	53	2.0%	4.65	0.33
Russia	14	0.5%	3.55	0.19

The relevant publications were collected through a set of search terms: electronic voting, e-voting/evoting, e-vote, remote electronic voting, internet voting, online voting, voting online, electronic elections, internet elections, online elections, electronic ballots, digital ballot, direct-recording electronic voting machine, voting machine. This set was used to retrieve from the Scopus database all journal articles, conference proceedings, and book chapters that have one of these search terms in the title, abstract, or keyword list. We found a total of 2928 relevant publications, mostly published after 2000. We then used the author affiliation addresses to classify the retrieved documents in terms of the country groupings described above. The following analyses were conducted:

1. The growth of papers on e-voting in each of the country-groupings was calculated: what is the relative share of each group of countries, and which are the fast and slow growing groups?
2. What disciplines are involved in e-voting research? In other words, what kind of issues are addressed (technical, political, economic, social, legal)?
3. What is the attitude and sentiment towards e-voting: what is the relative frequency of positive and negative evaluation words, and what is the relative frequency of positive and negative emotion words? This analysis is based on the abstracts of the retrieved papers.

The first two questions can be answered by counting the publications, and by classifying them in subject areas. For this, standard Scopus tools were used. Somewhat more complex is the approach for the third question. The Linguistic Inquiry and Word Count (LIWC) software was used to analyse the abstracts of the papers [32]. LIWC contains dictionaries that linguistically classify words. Many linguistic categories are available, and for each a list of words is provided that cover that linguistic category. For the categories we use, we give some example words below.

For our analysis we use the following linguistic categories: positive emotions, negative emotions, superlatives, and negation words (all for measuring the general attitude towards e-voting), positive evaluation words and negative evaluation words (both for measuring the way e-voting is evaluated). Details about these categories that partly come from LIWC and were extended by us can be found in [36]. The asterisk (*) after a word means that this is a stem word, and the linguistic category includes all variants of that stem.

- Positive emotions: words such as agreeable*, benefit, helpful (LIWC).
- Negative emotions: words such as abuse*, bitter*, bad*.
- Superlatives: words such as outstanding, exceptional*, groundbreaking, great potential, high gain.
- Negating words such as hasn't, don't, can't.
- Negative evaluation: words such as naïve, defect*, lack*.
- Positive evaluation: words like intriguing, compelling, commit*.

After describing the relative share of the different regions in the scholarly e-voting literature, we restrict the other analyses to the country groupings with more than a 5% share in the total e-voting papers, with the exception of Sub-Saharan Africa (2% of total publications) as we have some special interest in it.

4 Results

4.1 Growth of e-voting research by region

Overall, e-voting research started with the new century, and the output has increased over the years (Fig. 2 and 3). The EU15+ region is responsible for the largest part, and its output has steadily increased. The US showed a small increase in the beginning, but has stabilized over the last decade.

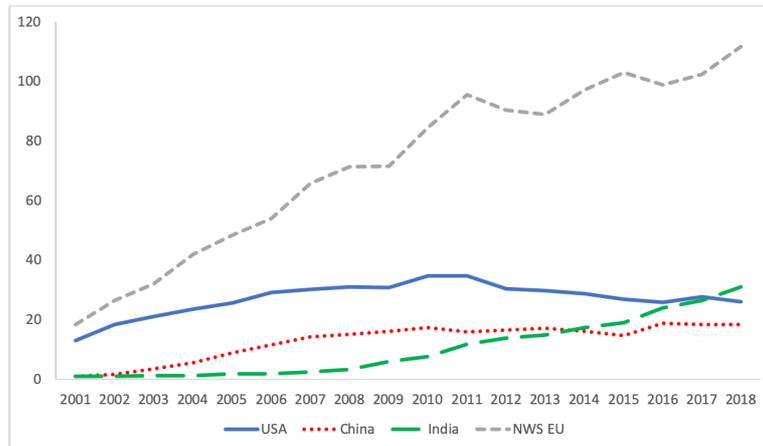


Fig. 2. Number of publications by year and region (selected regions).

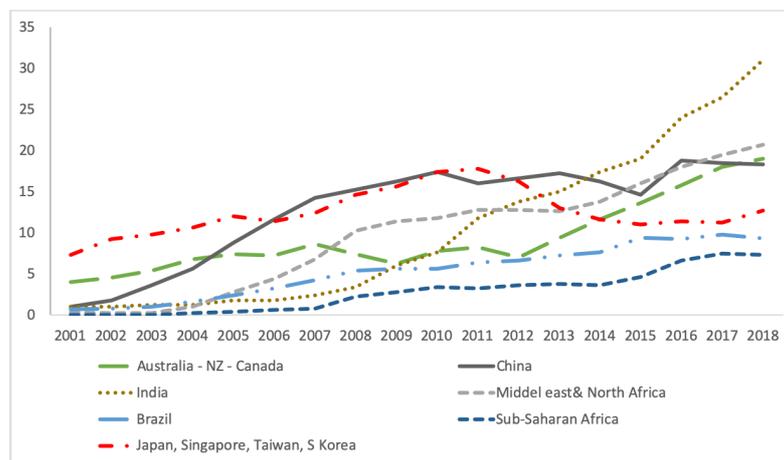


Fig. 3. Number of publications by year and region (selected regions).

In China we also see a small and more or less constant output, but one should be aware that this only covers the output in journals indexed in Scopus, so there may be an unobserved local output in the Chinese language. In India we see an emerging and growing activity. The other regions have a small output, that in most cases show an increase.

4.2 Distribution of papers over disciplines: regional differences?

Probably more interesting than the size of e-voting research may be the topical distribution of the research. For this, the field of the publications is a useful indicator. Figure 5 shows the distribution of papers over subject areas (disciplines). Computer Science is the largest, with 42.6%, followed by Mathematics with 15.3% and Engineer-

ing with 14.5%, and together they cover almost three quarters of all publications. This is in line with a study by [17] who systematically examined sixty-seven articles on e-voting in developing countries and found that the literature focused mainly on how to practically put the technology into effect.

There is also a set of Medical papers (2.2%), and inspecting those in more detail shows that they describe the use of e-voting technologies in for example consultation meetings on medical issues, where reaching a consensus is the aim. So these medical papers are in fact of a social science nature. The other social science related papers are Business and Management (3.0%) and Decision Sciences (3.6%). Overall, Social Sciences are good for 19.6% of the papers.

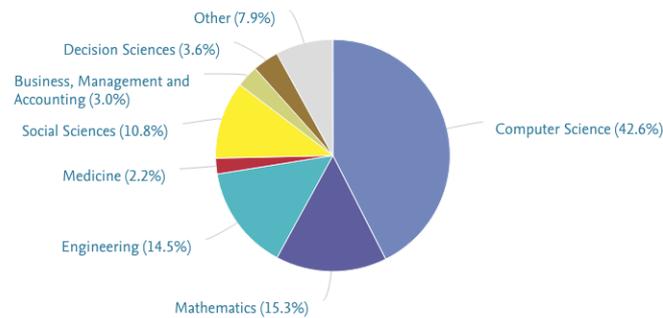


Fig. 4. Worldwide e-voting research by subject area 2009-2018. (Source: Scopus)

The regions clearly differ in the disciplinary (and related topical) emphasis. Overall, Computer Science, Mathematics, and Engineering cover 72.4% of e-voting research, but that differs between the regions.

Table 4. Percentage of e-voting publications by discipline (2009-2018).

	EU	USA	A/N/C	MENA	Japan	China	India	SSA
Computer Science	43.1	40.2	38.1	42.1	49.4	42.5	42.6	40
Mathematics	18.5	14.6	20.5	8.9	14.4	12.7	9.4	4
Engineering	10.2	12.1	11.2	18.6	18.7	23.9	24.3	19
Physics				2.1		5.9	3.9	
<i>Subtotal</i>	<i>71.8</i>	<i>66.9</i>	<i>69.8</i>	<i>71.7</i>	<i>82.5</i>	<i>85</i>	<i>80.2</i>	<i>63</i>
Social Sciences	12.1	16	11.2	7.5	5.2	3.9	3.4	17
Management	3.5	2.7		4.3	2.3	2.2		6
Decision Sciences	3.7	3.1		4.3	3.5	2.4	2.6	6
Medicine	2.7	4	9.3	3.2	2.3			2
<i>Subtotal</i>	<i>22</i>	<i>25.8</i>	<i>20.5</i>	<i>19.3</i>	<i>13.3</i>	<i>8.5</i>	<i>6</i>	<i>31</i>
Other	6.3	7.4	9.7	8.9	4.2	6.4	13.6	7

EU = North, West, South EU; A/N/C = Australia, New Zealand, Canada;
MENA = Middle East & North Africa; Japan = incl. Singapore, Taiwan, S Korea;
SSA = Sub-Saharan Africa

As Table 4 shows, India (80.2%), Japan (82.5%) and China (85%) are at the higher end, whereas Europe (71.8%), Canada & Australia (69.8), and the USA (66.9%), but also the Middle East and North Africa (MENA) are at the lower end (71.7%). The latter region is dominated by Iran. The opposite is the case for social research. MENA (19.3%), the EU (22%), Canada & Australasia (20.5%), and especially the USA (25.8%) have a relatively large share of their research on e-voting within the Social Sciences, while this is much smaller in India (6%), China (8.5%), and Japan (13.3%). We added Sub-Saharan Africa, and interestingly, in terms of disciplinary distribution it seems to be similar to the full democracies, with its relatively small share of Computer Science, Mathematics and Engineering papers, and a large share of Social Science related papers.

4.3 Perception and evaluation of e-voting

We use the abstract of the papers for a linguistic analysis. We are especially interested in whether e-voting is positively or negatively perceived and evaluated. As mentioned in the methodology section, we measure this through positive and negative emotion words, through negation words and superlatives, and through positive and negative evaluation words.

If the discourse about electronic voting is generally positive in a country, one would expect more positive evaluation words, more positive emotion words, and more superlatives, whereas in a negative discourse the opposite would be expected. In the next Table 5, we compare the full democracies with the hybrid and authoritarian regimes. As we discussed in the background section, Africa may be an interesting case, and therefore we added to this analysis also Sub-Saharan Africa. As the number of papers is small, the results should be interpreted with care.

Table 5. Results linguistic analysis

	Democracy	Positive emotions	Positive evaluation	Negative evaluation
Australia, NZ, Canada	9.17	3.05	2.28	0.83
North, West, South EU	8.62	2.95	2.08	0.74
USA	8.15	2.79	2.14	0.80
Japan, Singapore, Taiwan, S Korea	7.70	3.13	1.58	0.84
India	4.68	3.35	1.79	0.59
Sub-Saharan Africa	4.65	2.70	1.99	0.57
Middle East & North Africa	3.81	3.16	1.93	0.68
China	3.51	4.13	1.84	0.64

Several linguistic categories (negative emotions, superlatives) did not show much difference between the regions hence we exclude them here. Table 5 shows that the full democracies have a higher percentage of both positive and negative evaluation words than the hybrid and authoritarian regimes. On the other hand, the latter group has higher scores on positive emotion words. Together this suggests a somewhat more neutral (smaller role for emotions) and balanced (in terms of pro and con voices) aca-

ademic debate on e-voting technologies in the full democracies. Japan, Singapore, Taiwan, and South Korea are in between in terms of the development of democracy, and score similar to hybrid/authoritarian regimes on positive emotions and positive evaluations but similar to the full democracies on negative evaluation. Finally, Sub-Saharan Africa is in positive emotions (and to a lesser extent in positive evaluation) similar to the full democracies, and scores for negative evaluation similar to the hybrid democracies/authoritarian regimes group. Despite its low score on democracy, it has an in between position for the linguistic variables.

5 Discussion and conclusions.

The decision to implement electronic voting depends on many different factors, and governments and civil society have to weigh up the opportunities and risks involved. Several old democracies have concluded that e-voting systems are too insecure and that they need an election system that is resilient to both insider and outsider threats. For many this means sticking with, or returning to, paper. Election meddling – internal or external – is a big worry for those concerned with the integrity of democracies. As we have seen, election fraud scenarios are not impossible under paper-based practices, but e-voting would put them on an infinitely wider scale. While there are clear possible benefits and drawbacks of e-voting in well-established democracies, the situation in emerging democracies is more complex. Inadequate transparent mechanisms are a problem of many existing voting systems in new democracies, with a danger of being prone to human error and deliberate manipulations. We wanted to investigate whether the different political contexts in countries lead to different national discourses on computer-based voting.

First of all, we see an increased activity related to e-voting research in emerging democracies. This research mainly takes place in the fields of Computer Science, Mathematics, and Engineering indicating an emphasis on technological issues. We calculated the share in Social Sciences and Humanities research within e-voting. The regional differences are significant (Table 4) with the old democracies including the USA showing a large share of SSH research, whereas it is almost absent in the upcoming e-voting research areas (with the exception of Sub-Saharan Africa and to a lesser extent in North Africa and the Middle East). When looking at the linguistic analysis of the publications we also find that long-standing democracies have a more neutral and balanced academic debate on e-voting technologies, addressing both the advantages and disadvantages, followed by Africa and the Middle East.

In the last 15 years there has been an increasing awareness that the implementation of e-voting could present a number of social, political, legal, and economic challenges, and these aspects should be considered thoroughly. Nonetheless, emerging democracies seem less critical about e-voting with overall fewer national academic debates taking place related to these vulnerabilities and challenges. In the cases where vulnerabilities do get addressed it is often by international researchers, as we have shown for India [38].

Why is it that we can mainly detect a substantial share of papers focusing on societal context in old democracies? Partly it can be explained by the scientific specialisation of countries. In Western democracies the share of Social Sciences and Humanities in the total academic output is much higher than in most other regions. For example, in the US it is about 20%, in the Netherlands about 22%, and in India and China only around 4%. Interestingly, the high score of social science papers on e-voting (and the related evaluation and emotion scores for e-voting) in Sub-Saharan Africa is in this perspective not unexpected. For example, in South Africa, the share of Social Sciences and Humanities in total research output is higher than in the USA and the Netherlands: 23%. This suggests that the structure of the academic system itself is critical for a balanced assessment of the pros and cons of e-voting technologies.

But there may be other factors. Do the often-troubled elections in newly emerged democracies instigate a less critical view as they can bring important benefits to a country and its people? In other words, does the context of elections justify the different approaches and assessment of e-voting systems? Do the benefits outweigh the negatives in some countries? As Alvarez et al. [1] note: “In countries where there is widespread disbelief in the freeness and fairness of elections and where the complexity of voting procedures can actually prevent important segments of the electorate from exercising their right to vote, the introduction of e-voting systems poses both a difficult challenge and an interesting opportunity”.

When we started our e-voting research back in 2002, we piloted computer-based voting in 4 different European countries and we found that there is a correlation between the location (country) of the respondents and their trust in the verifiability and security of the system. We saw that the respondents from Italy, which was then under the rule of Berlusconi - a Prime Minister surrounded by many controversies (e.g. economic conflicts of interest, media control, alleged links to the Mafia) - had a lower trust in the verifiability of the system than the Finnish respondents [24]. This suggested that people judged the appropriateness of the implementation of e-voting technology on how democratic they considered a country to be. Our respondents felt that while it was okay to adopt electronic voting in full stable democracies, it was inconceivable to them to use it in emerging democracies or authoritarian states with high levels of corruption.

One could however argue that while there is a danger that electronic voting systems can be used as a controlling mechanism in regimes where elections are integrally tied to the regime’s legitimacy (i.e. primarily one-party regimes), it can also be used to combat electoral fraud and violence in a corrupt system. It could be reasoned that while accepting that e-voting security is a concern, the challenges are outweighed by the benefits. Systematic irregularities and violations, such as ballot box stuffing, voter intimidation, violence, dishonest counting of votes, or insecure transport of ballots after the vote could be avoided by the implementation of e-voting. Adopting electronic voting systems has the potential to enhance the quality of electoral processes in less developed democracies [1]. According to [17] “The shortcomings experienced during previous democratic practices might have resulted in technological determinism shown by countries such as Nigeria and India”. Many countries are accustomed to much more old-fashioned kinds of manipulation and corruption and therefore the

adoption of electronic voting systems has helped strengthen public confidence in the legitimacy of elections in many places, and has perhaps created a reluctance to conduct research on the negative implications of this technology.

This paper contributes to the analysis of how electronic voting is received, and whether there are differences related to region and level of democracy. In a follow-up paper, we plan to go more in detail and we will analyze the academic discourse at a more fine-grained topical level: what technical questions are discussed, and what political questions. Furthermore, there is an urgent need to study the benefits and risks in those countries where e-voting is practiced. Are the benefits really worth the risks? And how do ethical and political decisions about e-voting reflect this? Is there evidence that the new electoral procedure does more good than harm, and does that depend on the level of democracy in the countries that deploy e-voting? In this context, one has to be critical on the motives behind the introduction of e-voting [17].

6 References

1. Alvarez, R. M., Katz, G., Llamasa, R., & Martinez, H. (2009) Assessing voters' attitudes towards electronic voting in Latin America: Evidence from Colombia's 2007 e-voting pilot. *International Conference on E-Voting and Identity* (pp. 75-91). Springer, Berlin.
2. Alvarez, R. M., Katz, G., & Pomares, J. (2011). The impact of new technologies on voter confidence in Latin America: Evidence from e-voting experiments in Argentina and Colombia. *Journal of Information Technology & Politics*, 8(2), 199-217.
3. Arango, T. and D. Adnan (2014) Militants Pose Threat on Eve of National Elections in Iraq. *The New York Times*, 28 April 2014. <https://www.nytimes.com/2014/04/29/world/middleeast/iraq-prepares-for-national-elections-in-the-shadow-of-militant-threats.html>
4. BBC (2015) Turkey election: OSCE says 'serious concerns' over vote. <https://www.bbc.co.uk/news/world-europe-34704834>
5. Biswas, S. (2019) India election 2019: Are fears of a mass hack credible? BBC, 25 January 2019. <https://www.bbc.co.uk/news/world-asia-india-46987319>
6. Bjarnesen, J. and M. Söderberg Kovacs (2018) Violence in African Elections. NAI Policy Note No 7:2018. Nordiska Afrikainstitutet/The Nordic Africa Institute. https://reliefweb.int/sites/reliefweb.int/files/resources/Violence%20in%20African%20Elections_Policy%20Note_Final%20version.pdf
7. Bloomfield, S. (2008) Kibaki 'stole' Kenyan election through vote-rigging and fraud. *The Independent*, 23 January 2008. <https://www.independent.co.uk/news/world/africa/kibaki-stole-kenyan-election-through-vote-rigging-and-fraud-772349.html>
8. Butagira, T. (2016) Ugandans are desperate for democracy, Yoweri Museveni only gives them tyranny, *The Guardian*, 16 February 2016. <https://www.theguardian.com/commentisfree/2016/feb/22/uganda-democracy-yoweri-museveni-tyranny-president-force-intimidation>
9. Caarls, S. (2010) E-voting Handbook: Key Steps in the Implementation of E-enabled Elections. Council of Europe Publishing, Strasbourg.
10. Chan, S. (2017) Fearful of Hacking, Dutch Will Count Ballots by Hand. *New York Times*, 1 February 2017. <https://www.nytimes.com/2017/02/01/world/europe/netherlands-hacking-concerns-hand-count-ballots.html>

11. Day, M (2016) Romanian MP jailed for bribing voters with 60 tonnes of roasted chicken. The Telegraph, 16 March 2016. <https://www.telegraph.co.uk/news/worldnews/europe/romania/12195737/Romanian-MP-jailed-for-bribing-voters-with-60-tonnes-of-roasted-chicken.html>
12. Debnath, S., M. Kapoor and S. Ravi (2017) The Impact of Electronic Voting Machines on Electoral Frauds, Democracy, and Development. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3041197
13. Digital Democracy Commission (2015). Open up! Report of the Speaker's Commission on Digital Democracy. <http://www.digitaldemocracy.parliament.uk/documents/Open-Up-Digital-Democracy-Report.pdf>
14. EIU (2019) Democracy Index 2018: Me too? Political participation, protest and democracy. London: The Economist Intelligence Unit. https://www.eiu.com/public/topical_report.aspx?campaignid=Democracy2018
15. Giles, C. (2018) DR Congo elections: Why do voters mistrust electronic voting? BBC Reality Check, 28 December 2018. <https://www.bbc.co.uk/news/world-africa-46555444>
16. Gonggrijp, R., & Hengeveld, W. J. (2007). Studying the Nedap/Groenendaal ES3B voting computer: A computer security perspective. In *Proceedings of the USENIX workshop on accurate electronic voting technology*. USENIX Association. http://static.usenix.org/event/evt07/tech/full_papers/gonggrijp/gonggrijp.pdf
17. Hapsara, M., Imran, A., & Turner, T. (2016). E-Voting in Developing Countries. In *International Joint Conference on Electronic Voting* (pp. 36-55). Springer, Cham. https://www.researchgate.net/publication/312923714_E-Voting_in_Developing_Countries
18. Hern, A. (2015) Should Britain introduce electronic voting? The Guardian, 26 February 2015. <https://www.theguardian.com/technology/2015/feb/26/should-britain-introduce-electronic-voting>
19. Ilie, L. (2015) Romanian minister found guilty of vote-rigging in referendum. Reuters, 15 May 2015. <https://www.reuters.com/article/us-romania-corruption/romanian-minister-found-guilty-of-vote-rigging-in-referendum-idUSKBN000J820150515>
20. Kitcat, J. (2007) Electronic voting: A challenge to democracy? Report by the Open Rights Group (ORG). <https://www.openrightsgroup.org/uploads/org-evoting-briefing-pack-final.pdf>
21. Kohut, M. (2017) Venezuela Reported False Election Turnout, Voting Company Says. The New York Times, 2 August 2017. <https://www.nytimes.com/2017/08/02/world/americas/venezuela-election-turnout.html>
22. Küsters, R. and J. Müller (2017) Cryptographic Security Analysis of E-voting Systems: Achievements, Misconceptions, and Limitations. In: R. Krimmer et al. (Eds.): *E-Vote-ID 2017*, LNCS 10615, pp. 21–41, 2017.
23. Morlin-Yron, S. (2016) Why nearly half of Africans don't trust elections, CNN September 26, 2016. <https://edition.cnn.com/2016/09/25/africa/africa-view-election-distrust/index.html>
24. Oostveen, A. and P. van den Besselaar (2004) Security as Belief. Users Perceptions on the Security of Electronic Voting Systems. In: A. Prosser and R. Krimmer (Eds.) *Electronic Voting in Europe: Technology, Law, Politics and Society*. Lecture Notes in Informatics. Vol.47, pp 73-82. Bonn: Gesellschaft für Informatik. <http://www.social-informatics.net/ESF2004.pdf>
25. Oostveen, A. and P. van den Besselaar (2005) Trust, Identity, and the Effects of Voting Technologies on Voting Behavior. *Social Science Computer Review* 23 (3) pp. 304-311. <https://pdfs.semanticscholar.org/14db/661dc1d39c231278457ca0959ca4536ccc15.pdf>

26. Oostveen, A. (2010a) Citizens and Activists: Analyzing the Reasons, Impact and Benefits of Civic Emails Directed at a Grassroots Campaign. *Information, Communication and Society* 13 (6) pp.793-819.
<https://www.tandfonline.com/doi/abs/10.1080/13691180903277637>
27. Oostveen, A. (2010b) Outsourcing Democracy: Losing Control of e-Voting in the Netherlands. *Policy and Internet* 2 (4) pp. 201-220.
<https://onlinelibrary.wiley.com/doi/abs/10.2202/1944-2866.1065>
28. OSCE (2005) Election Observation. A decade of monitoring elections: the people and the practice. OSCE Office for Democratic Institutions and Human Rights (ODIHR). Poland: Warsaw. <https://www.osce.org/odihr/elections/17165?download=true>
29. OSCE (2013) Handbook for the Observation of New Voting Technologies. OSCE Office for Democratic Institutions and Human Rights (ODIHR). Poland: Warsaw. <https://www.osce.org/odihr/elections/104939?download=true>
30. O'Sullivan, D. (2019) 'Make Swiss Democracy Safe Again'. How e-voting became a fight for democracy. https://www.swissinfo.ch/eng/technology_how-e-voting-became-a-fight-for-democracy-/44838246
31. Penar, P., R. Aiko, T. Bentley, and K. Han (2016) Election quality, public trust are central issues for Africa's upcoming contests. Afrobarometer Policy Paper No. 35. http://afrobarometer.org/sites/default/files/publications/Policy%20papers/ab_r6_policypapero35_electoral_management_in_africa1.pdf
32. Pennebaker, J.W., Boyd, R.L., Jordan, K., & Blackburn, K. (2015). The development and psychometric properties of LIWC2015. Austin, TX: University of Texas at Austin. http://liwc.wpengine.com/wp-content/uploads/2015/11/LIWC2015_LanguageManual.pdf
33. Pomares, J., Levin, I., Alvarez, R. M., Mirau, G. L., & Ovejero, T. (2014). From piloting to roll-out: voting experience and trust in the first full e-election in argentina. In *6th International Conference on Electronic Voting (EVOTE)*, pp. 1-10. IEEE.
34. Rasheed, A., R. Jalabi, A. Aboulenein (2018) Exclusive: Iraq election commission ignored warnings over voting machines - document. Reuters, 5 August 2018. <https://www.reuters.com/article/us-iraq-election-exclusive/exclusive-iraq-election-commission-ignored-warnings-over-voting-machines-document-idUSKBN1KQ0CG>
35. Schneier, B. (2004) Voting Security. IEEE Security & Privacy. July/August 2004. https://www.schneier.com/essays/archives/2004/07/voting_security.html
36. Van den Besselaar, P., U. Sandström, H. Schiffbaenker (2018) Using linguistic analysis of peer review reports to study panel processes. *Scientometrics* 117, 313-329. <https://link.springer.com/article/10.1007/s11192-018-2848-x>
37. Wilkie, J. (2019) America's new voting machines bring new fears of election tampering. The Guardian, 22 April 2019. <https://www.theguardian.com/us-news/2019/apr/22/us-voting-machines-paper-ballots-2020-hacking>
38. Wolchok, S., E. Wustrow, A. Halderman, H.K. Prasad, A. Kankipati, S. Krishna Sakhamuri, V. Yagati, and R. Gonggrijp (2010) Security analysis of India's electronic voting machines. In Proceedings of the 17th ACM conference on Computer and communications security (CCS '10). ACM, New York, NY, USA, 1-14. https://www.researchgate.net/publication/221609811_Security_Analysis_of_India's_Electronic_Voting_Machines
39. Zetter, K. (2018) The Myth of the Hacker-Proof Voting Machine. The New York Times Magazine, 21 February 2018. <https://www.nytimes.com/2018/02/21/magazine/the-myth-of-the-hacker-proof-voting-machine.html>

40. Chipman, I. (2016) David Dill: Why Online Voting is a Danger to Democracy. Stanford Engineering. <https://engineering.stanford.edu/magazine/article/david-dill-why-online-voting-danger-democracy>

Risk Limiting Audit and its Applications

Auditing Indian Elections

Vishal Mohanty¹[0000-0002-1556-1113], Chris Culnane²[0000-0002-9543-1342],
Philip B. Stark³[0000-0002-3771-9604], and Vanessa Teague²[0000-0003-2648-2565]

¹ Indian Institute of Technology Madras vishalmohanty97@gmail.com

² University of Melbourne {christopher.culnane,vjteague}@unimelb.edu.au

³ University of California, Berkeley
stark@stat.berkeley.edu

Abstract. Electronic Voting Machines (EVMs) used in the 2019 General Elections in India were fitted with printers to produce Voter-Verifiable Paper Audit Trails (VVPATs). VVPATs allow voters to check whether their votes were recorded as they intended. However, confidence in election results requires more: VVPATs must be preserved inviolate and then actually used to check the reported election result in a trustworthy way that the public can verify. A full manual tally from the VVPATs could be prohibitively expensive and time-consuming; moreover, it is difficult for the public to determine whether a full hand count was conducted accurately. We show how *Risk-Limiting Audits* (RLAs) can provide high confidence in Indian election results. Compared to full hand recounts, RLAs typically require manually inspecting far fewer VVPATs when the outcome is correct, and are much easier for the electorate to observe in adequate detail to determine whether the result is trustworthy. We show how to apply two RLA strategies, *ballot-level comparison* and *ballot polling*, to General Elections in India. Our main result is a novel method for combining RLAs in constituencies to obtain an RLA of the overall parliamentary election result.

Keywords: Risk-Limiting Audit, Ballot-Polling Audit, Ballot-Level Comparison Audit, Transitive Audit, Multi-level Election Audit, Fisher’s Combining Function

1 Introduction

Since Electronic Voting Machines (EVMs) were introduced in India in the 1999 elections, there have been concerns about their transparency and trustworthiness; a number of security vulnerabilities have been documented [19]. In 2013, the Indian Supreme Court ruled that all EVMs in Indian General Elections must be equipped with printers providing Voter-Verifiable Paper Audit Trails (VVPATs, [17]). The Election Commission of India used VVPAT-equipped EVMs in the 2019 General Elections in all constituencies across the nation.

VVPATs allow each voter to verify that his or her intended selections are correctly printed on a paper record, collected in a separate container called the

VVPAT box. VVPATs provide a way to check and correct election results, for instance, if there is a legal demand by a candidate, or for routine checks of election tabulation accuracy—audits. VVPATs could be manually recounted to check the electronic results, but that is labor-intensive and time-consuming. We show how auditing a random sample of VVPAT records can justify confidence in election results without a full manual tally. Auditing a VVPAT means manually inspecting the paper record to see the voter preferences it shows.

The Election Commission (EC) of India is increasing the transparency of the Indian elections. In a recent report,⁴ the EC decided to tally the paper trail slips and compare them with the electronic result provided by the EVMs in 5% of the booths in each Assembly seat district, selected randomly. This effort, while well-intentioned, does not suffice to give strong evidence that election results are correct. In this paper, we show rigorous ways of attaining well-defined confidence levels.

Suitable post-election audits may justify confidence of voters, candidates, and parties that election results are correct. One type of post-election audit is a Risk-Limiting Audit (RLA), which either develops strong statistical evidence that the reported outcome is correct, or corrects the results (by conducting a full manual tally of a reliable paper trail). Here, “outcome” means the set of reported winners of the contests, not the exact vote tallies. To ensure that the tallies are correct to the last vote is prohibitively expensive, if not impossible; conversely, to ensure that the reported winners really won seems like the lowest reasonable standard for accuracy.

Before an RLA commences, the *risk limit* α must be chosen; ideally, it is set in legislation or regulation, so that auditors cannot manipulate the level of scrutiny a contest gets by adjusting the risk limit. The risk limit is the maximum probability that the audit will fail to correct the reported election outcome, on the assumption that the reported outcome is wrong. The risk limit is a worst-case probability that makes no assumption about *why* the outcome is wrong, *e.g.*, it could be because of accidental error, procedural lapse, bugs, misconfiguration, or malicious hacking by a strategic adversary who knows how the audit will be conducted. RLAs assume that the paper ballots reflect the correct outcome, *i.e.* that a full manual tally of the paper trail would show who really won. An RLA of an unreliable paper trail is “security theater.” Hence, there need to be procedures (called *compliance audits* by [2,6,16,14]) to ensure that the paper trail is complete and intact before the RLA begins.

This paper shows how two types of RLAs can be used with Indian elections: transitive ballot-level comparison audits and ballot-polling audits. Ballot-level comparison audits are more efficient in the sense that they generally involve inspecting fewer ballots to attain the same risk limit when the reported outcome is correct. However, they require more setup. As discussed in Section 3, they may require a voting system that can export its interpretation of individual ballots in a way that can be matched to the corresponding paper, or may require sorting

⁴ <http://indianexpress.com/article/india/ec-to-tally-paper-trail-slips-with-evms-in-5-pc-booths-in-each-assembly-seat-4737936/>

the physical ballots or VVPATs before the audit, according to the votes they (reportedly) show.

Our main contribution is to develop RLAs for a new social choice function—Indian parliamentary majorities—with procedures suited to the logistics of Indian elections. To verify the overall election outcome we need to verify that the party/coalition reported to have been elected to form the government actually won. That generally requires less auditing than confirming the winner in every constituency. The method we develop splits the responsibility of the auditing among various constituencies in a way that the combined result gives higher confidence in the correctness of the overall parliamentary outcome than each constituency would have in its results alone. This procedure is discussed in Section 4. Our methods apply to any parliamentary democracy, but the computations are particularly simple when all constituencies have equal weight.

2 Background

RLAs are procedures that guarantee a minimum chance of conducting a full manual tally of the voter-verifiable records when the result of that tally would belie the reported outcome. They amount to a statistical test of the *null hypothesis* that the election outcome is wrong, at significance level α , the *risk limit*. An RLA continues to examine more ballots until the null hypothesis is rejected at significance level α , or until there has been a complete manual tally to set the record straight. The risk limit is the largest chance that the audit will *not* require a full manual tally of the paper records if the electoral outcome according that tally would differ from the reported electoral outcome.

2.1 Indian Elections

Indian General elections are held Quinquennially to elect the Lok Sabha (Lower House of the Parliament). The country is divided into 543 constituencies, each represented by one person elected to the Lok Sabha. Elections at the constituency level are *plurality* contests: the person who gets the most votes wins. Candidates at the constituency level typically belong to some political party, but can be unaffiliated with any party. At the parliamentary level, the party that gets the *majority* (at least 272) of the seats forms the government. If no party has a majority, parties may form coalitions to attain a majority. Coalitions can be formed before or after elections, although before is more common. Elections are conducted in phases spread over a month. Each phase consists of single-day elections in a subset of constituencies, typically grouped by geography.

2.2 Related work on Election Auditing

RLAs were introduced by [11], but were not so named until [12]. The first RLAs were conducted in California in 2008 [4]. RLAs have been conducted in California, Colorado, Indiana, Michigan, New Jersey, Ohio, Rhode Island, Virginia,

and Denmark. RLAs have been developed for a variety of social choice functions, a variety of sampling strategies (unstratified sampling of individual ballots or batches, with or without replacement, with or without weights; stratified sampling with and without replacement and with and without weights, Bernoulli sampling, weighted random sampling) and auditing strategies (*batch-level comparisons*, *ballot-level comparisons*, *ballot polling*, and mixtures of those strategies).

Ballot-polling audits [7,6,8] do not require knowing how the system interpreted individual ballots nor how it tallied the votes on subsets of ballots. They directly check whether the reported winner(s) received more votes than the reported loser(s) by sampling and manually interpreting individual ballots. To draw a random sample of ballots typically involves a *ballot manifest*, which describes how the physical ballots are organized: the number of bundles, the labels of the bundles, and the number of ballots in each bundle (However, see [8]).

The BRAVO ballot-polling method [7] uses Wald’s sequential probability ratio test [18] to test the hypotheses that any loser in fact tied or beat any winner. The audit stops short of a full hand count if and only if there is sufficiently strong evidence that every winner beat every loser.

Comparison audits involve manually checking the voting system’s interpretation of the votes on physically identifiable subsets of ballots. They require the voting system to export vote tallies for physically identifiable subsets of ballots, so that the votes on those ballots can be tallied by hand and compared to the voting system’s tallies. They also require checking that the reported subtotals yield the reported contest results, and that the subtotals account for all ballots cast in the contest. They generally also require ballot manifests.

A comparison audit that checks the voting system’s interpretation of individual ballots is a *ballot-level comparison* RLA. Ballot-level comparison RLAs are more efficient than batch-level comparison RLAs and ballot-polling RLAs in that they generally require examining fewer ballots when the reported outcome is correct. However, they have higher set-up costs and require more data export from the voting system: they need a *cast vote record* or *CVR* for each physical ballot, a way to locate the CVR for each physical ballot, and *vice versa*. (A CVR is the voting system’s interpretation of voter intent for a given ballot.) Relying on more general results in [12] for batch-level comparison audits, [13] developed a sequential ballot-level comparison RLA method that results in particularly simple calculations.

Hybrid audits combine different approaches to using audit data in different strata, for instance, ballot-polling in some strata and ballot-level comparisons in other strata; See [9].

Transitive audits [3,6] involve auditing an unofficial system. If that system reports the same winner(s) as the official system, an audit that provides strong evidence that the unofficial system found the correct winner(s) transitively provides strong evidence that the official system did also; and if the audit of the unofficial system leads to a full manual tally, the outcome of that tally can be used to correct the official result. A transitive audit does not confirm that the

official system tallied votes correctly: the two systems might disagree about the interpretation of every ballot, but still agree who won.

Indian EVMs do not create CVRs, but organizing VVPATs appropriately makes it possible to audit EVMs using a transitive ballot-level comparison audit. CVRs can be constructed for EVMs by sorting the VVPATs into bundles that (purportedly) show the same voter preferences, counting the number of VVPATs in each batch, and labeling each bundle with the number of ballots and the voter preferences it purports to contain. A report of the bundle labels, the number of VVPATs in each bundle, and the reported voter preference for the bundle amounts to a CVR for every VVPAT. Such a report in effect combines a *ballot manifest* [6] and a commitment to a cast vote record for every ballot, implied by the label of the bundle the ballot is in. We shall call such a report a *preference manifest*.

If ballots are sufficiently simple—*e.g.*, if each contains only one contest, as in India—sorting ballots by voter preference can be practical. Indeed, this is how ballots are tallied in Denmark: on election night, ballots are sorted within polling places according to the voter’s party preference. The following day, ballots are sorted further according to the voter’s candidate preference, to produce homogeneous bundles of ballots, each labeled with the number of ballots and the voters’ preference.

Such sorting-based CVRs were the basis of an RLA in Denmark [10]. The sorting might be manual, as it is in Denmark, but it could be automated partly or entirely. (Sorting may also increase vote anonymity by breaking any link between voter and ballot.) When the official tallying process itself is based on creating and counting the homogeneous bundles, as it is in Denmark, the audit is a direct audit of the voting system. If the sorting is conducted independently of the tabulation, as it would be if India were to sort the paper ballots to produce a preference manifest, the resulting audit is a transitive audit.

The first step of a ballot-level comparison RLA is to verify that the CVRs produce the reported results: that applying the social choice function to the vote subtotals implied by the sizes of the bundles and the votes they purport to contain produces the same set of winners. If the preference manifest does not produce the reported set of winners, the audit should not continue: there is a serious problem. The audit should also check that the number of CVRs for each contest does not exceed the number of ballots cast in the contest, which should be determined without reliance on the voting system [1]. If the preference manifest passes these checks, the audit can check the accuracy of the CVRs implied by the preference manifest against a manual reading of voter intent from randomly selected paper ballots.

Kroll *et al.* [5] present a method for reducing the workload in auditing multi-level elections, inspired by the US Electoral College. They show that to achieve an overall confidence that a party or coalition secured the majority of seats, the individual constituencies can sometimes be audited to lower confidence levels. They provide a constrained optimization program describing the set of feasible solutions (*i.e.* those that constitute a sufficient audit) and a number of methods

for finding the optimal solution. In India’s electoral system, as in many other parliamentary democracies, every constituency has equal weight.

3 Auditing Individual Constituencies using Extant Methods

This section discusses how existing methods for RLAs apply to Indian elections. We consider auditing individual constituencies rather than the entire election; Section 4 shows how to combine audits of constituencies to audit an entire contest.

India’s voting system currently does not support ballot-level comparison audits, but, as described above, if procedures were added to sort the paper ballots and produce a preference manifest, transitive ballot-level comparison audits would be possible. Because ballots in India are simple—a single selection in a single contest—such sorting is feasible.

Ballot-polling audits could be used in India without sorting the ballots or modifying the voting system, if ballot manifests were available (see section 3.2). The calculations for BRAVO [7] and the ballot-polling method in [6] are simple enough to do with a pencil and paper or hand calculator, and are implemented in open-source online tools by Stark.⁵

When the election outcome is correct, ballot-level comparison audits generally require inspecting fewer ballots than ballot-polling audits. (Because they are RLAs, when the outcome is incorrect, both have a large chance of requiring a full manual tally.) The advantage grows as the margin shrinks: as a rule of thumb, workload increases inversely with the reported margin for ballot-level comparison audits, and increases inversely with the square of the actual margin for ballot-polling audits. However, preparing for a ballot-level comparison audit is harder, because it requires CVRs linked to the corresponding physical ballots. The simplicity of ballot-polling audits may offset the work of examining more paper ballots, unless the margin is very small.

3.1 Transitive Ballot-Level Comparison RLA

Ballot-level comparison RLAs were introduced by [13] who provides online tools at the link ⁶; see also [6]. Ballot-level comparison audits require a way to find the CVR corresponding to each paper ballot, and vice versa. The EVMs currently used in India do not provide CVRs at all.

However, as shown by [10], sorting ballots into groups according to the vote (if any) that they are reported to show in effect provides a CVR for each ballot through a preference manifest that lists the bundles of ballots, the number of ballots in each bundle, and the (single) preference that every ballot in the bundle is supposed to show. In Denmark, ballots are manually sorted into bundles with

⁵ <https://www.stat.berkeley.edu/~stark/Vote/ballotPollTools.htm>

⁶ <https://www.stat.berkeley.edu/~stark/Vote/auditTools.htm>

homogeneous voter intent, but sorting could be automated with relatively simple equipment, possibly something similar to the system used in South Korea.

Whether it is worth the effort to sort the ballots depends on the margin: if the margin is wide, it will be less expensive to use ballot polling, but if the margin is very narrow, the cost of sorting—whether manual or automated—may reduce the sample size required to confirm the outcome by orders of magnitude.

Classifying CVR Errors Stark [12] reviews a number of methods to test the hypothesis that any loser received more votes than any winner by comparing hand counts of votes in randomly selected batches of ballots to the machine counts of the votes on the same ballots. The methods apply to arbitrarily small batches, including batches consisting of a single ballot; that is, to ballot-level comparison audits. [13] elaborated one of those methods, which relies on the Kaplan-Markov inequality. By introducing a taxonomy of discrepancies, the arithmetic can be simplified to the point that a pencil and paper suffice, while rigorously controlling the risk. That “super-simple simultaneous single-ballot” method was further simplified by [6], and is the basis of pilot audits in Denmark, California, Colorado, Indiana, Michigan, New Jersey, Rhode Island, and Virginia, and of the statutory risk-limiting audits in most Colorado counties.

Stark and Teague [15] presented a ballot-level comparison RLA method based on the Kaplan-Wald inequality, which has some advantages over the Kaplan-Markov inequality. In this paper, we use the method of [6], because it has been used more widely. We shall refer to it as the *LSKM method*. It is straightforward to modify the procedures below to use the method of [15] or any other ballot-comparison RLA method instead.

The LSKM method is *sequential*: it examines more and more ballots selected at random until either there is strong evidence that the reported winners really won, or until there has been a full hand count and the correct outcome is known. Conceptually, after examining one or more ballots, one calculates a sequentially valid⁷ P -value of the hypothesis that the outcome is wrong. If that P -value is less than or equal to the risk limit α , the audit stops; otherwise, more ballots are audited and the sequential P -value is updated. The method presented in section 4 to check the overall electoral outcome involves combining the P -values for individual constituencies.

If the audit does lead to a full hand tally in a constituency, the reported results are replaced by the results according to that full hand tabulation. Election officials may elect to terminate the audit and conduct a full hand count at any time, for instance, if they estimate that the cost of additional sampling will exceed the cost of a full manual tally.

⁷ *Sequentially valid* means that the chance that the infimum of the P -value over all sample sizes is less than or equal to α is itself less than or equal to α if the null hypothesis is true. In contrast, standard hypothesis tests are designed for sample sizes that are fixed ahead of time: expanding the sample and re-calculating the P -value for such tests generally produces type I error rates far larger than the nominal significance level, because it does not account for multiplicity.

The LSKM method involves classifying discrepancies between the CVR and a manual reading of voter intent from the paper ballot:

- If correcting the CVR would reduce the margin between any (reported) winner and any (reported) loser by two votes, the discrepancy is a *2-vote overstatement* (the number of 2-vote overstatements is denoted o_2).
- If not, but if correcting the CVR would reduce the margin between any winner and any loser by one vote, the discrepancy is a *1-vote overstatement* (the number of 1-vote overstatements is denoted o_1).
- If not, but if correcting the CVR would not increase the margin between every winner and every loser, the discrepancy is a *neutral error*. (Neutral errors do not enter the stopping rule explicitly.)
- If not, but if correcting the CVR would increase the margin between every winner and every loser by at least one vote, and increase the margin between some winner and some loser by exactly one vote, the discrepancy is a *1-vote understatement* (the number of 1-vote understatements is denoted u_1).
- If correcting the CVR would increase the margin between every winner and every loser by two votes, it is a *2-vote understatement* (the number of 2-vote understatements is denoted u_2).

Two-vote overstatements should be rare if the voting system is working correctly: they involve mistaking a vote for a loser as a vote for a winner. Two-vote understatements should be even rarer—and are typically mathematically impossible. For instance, in a plurality, vote-for-one contest with three or more candidates, two-vote understatements are impossible, because they would require having mistaken a valid vote for the winner as a valid vote for every losing candidate.

We assume that there is a trustworthy upper bound on the total number of ballots cast, for instance, from pollbooks or from information about the number of eligible voters. A preliminary check should ensure that the preference manifest does not list more ballots than that upper bound: if there are more ballots listed than can exist, there is a serious problem that the audit cannot address by itself.⁸

In the sorted-ballot method described above,

- a *2-vote overstatement* occurs if we find a vote for a reported loser in the reported winner’s pile;
- a *1-vote overstatement* occurs if we find a vote for a different reported loser in a reported loser’s pile;
- *neutral errors* don’t occur;⁹

⁸ Prof. Sandeep Shukla of IIT Kanpur has pointed out that the current Indian VVPAT design does not protect against the EVM adding electronic votes and corresponding VVPATs when voters are not looking, because there is no publicly observable mechanism to ensure that at most one VVPAT is inserted into the box per voter. This needs to be addressed by improving the physical design in a way that is beyond the scope of this paper.

⁹ Indian EVMs (as far as we know) do not produce blank votes. However, if they did they could be accommodated easily. A 1-vote overstatement occurs if we find

- a *1-vote understatement* occurs when there are at least three candidates and we find a vote for the reported winner in a reported loser’s pile;
- a *2-vote understatement* occurs only when there are exactly two candidates, and we find a vote for the reported winner in the reported loser’s pile;
- if a pile turns out to be *smaller* than reported, the discrepancy can be addressed using the “phantom to zombie” approach of [1].
- if a pile turns out to be *larger* than reported, then some other pile must be *smaller* than reported, and the “phantom to zombie” approach of [1] will still ensure that the risk is controlled conservatively.

There are sharper ways to treat discrepancies than to use these categories (in particular, keeping track of *which* margins are affected by each discrepancy can reduce the number of ballots the audit inspects; see [13]). However, the bookkeeping is more complex. Categorizing discrepancies this way makes the calculations simple enough to do with a pencil and paper (aside from calculating 5 constants involving logarithms, which can be done once and for all and verified by anyone).

Calculations Let n denote the current sample size and α the risk limit. Fix $\gamma \geq 1$. The LSKM method stops auditing (and concludes that the reported winners really won) if

$$n \geq \frac{2\gamma}{\mu} \left(o_1 \log \left(\frac{1}{1 - \frac{1}{2\gamma}} \right) + o_2 \log \left(\frac{1}{1 - \frac{1}{\gamma}} \right) - u_1 \log \left(1 + \frac{1}{2\gamma} \right) - u_2 \log \left(1 + \frac{1}{\gamma} \right) - \log(\alpha) \right) \quad (1)$$

In this expression, μ is the *diluted margin*, the smallest difference in votes between any winner and any loser, divided by the total number of ballots in the population from which the sample is to be drawn, including ballots with invalid votes. The constant $\gamma \geq 1$ is the *error inflation factor*, which controls the operating characteristics of the LSKM method: the larger γ is, the fewer additional ballots need to be audited if a 2-vote overstatement is observed, but the smaller γ is, the fewer ballots need to be audited if no 2-vote overstatements are observed. Because two-vote overstatements should be rare, taking γ slightly larger than 1 should suffice. For γ exactly equal to 1, then if the audit finds even one 2-vote overstatement, the audit will not terminate without a full hand count. [6] suggest using $\gamma = 1.03905$, which makes the “cost” of a 2-vote overstatement 5 times larger than the “cost” of a 1-vote overstatement, where “cost” means the number of additional ballots that must be audited to attain the risk limit. Any value of $\gamma \geq 1$ gives a risk-limiting audit, but γ must be chosen before inspecting any ballots.

a blank vote in the reported winner’s pile. A neutral error would occur when there were at least three candidates and we found a blank vote in a reported loser’s pile. A one-vote understatement would occur when there were exactly two candidates and we found a blank vote in the reported loser’s pile.

3.2 Ballot-Polling RLA using BRAVO

The BRAVO ballot-polling RLA by [7] can be applied immediately to constituencies in India. In the Indian scenario, we have only one winner per constituency and one candidate per ballot. For each loser ℓ , the null hypothesis $H_{0w\ell}$ states that w did not get more votes than ℓ , that is, that ℓ actually tied or beat w . BRAVO uses the Sequential Probability Ratio Test by [18] to test all the null hypotheses simultaneously.

The audit begins by choosing the risk limit α . It also requires the reported vote totals for each candidate,¹⁰ but no other data from the voting system.

For every apparent loser ℓ , define the *conditional vote share* $s_{w\ell}$

$$s_{w\ell} \equiv \frac{v_w}{v_w + v_\ell} \tag{2}$$

Here, v_w and v_ℓ are the reported vote totals for the winner w and the loser ℓ respectively. If the reported vote tally is correct, the chance that a randomly selected ballot shows a vote for w , given that it shows a vote for either w or ℓ , is $s_{w\ell}$.

BRAVO maintains a test statistic $T_{w\ell}$ for each reported (winner, loser) pair. In Indian elections, there is only one reported winner w per constituency, so this amounts to a test statistic for each reported loser ℓ . Null hypothesis $H_{0w\ell}$ is rejected if

$$T_{w\ell} \geq \frac{1}{\alpha}. \tag{3}$$

If the null hypotheses $\{H_{0w\ell}\}$ for all apparent losers $\ell \in L$ are rejected, the audit stops and the reported outcome becomes final.

At any time, for example if the audit is expected to take more time than simply counting the ballots, auditors can stop sampling and perform a full manual recount. The algorithm runs as follows:

¹⁰ There are other ballot-polling methods that do not use the reported results at all.

Algorithm 1 BRAVO with protection against manifest errors.

This is a simplified version of BRAVO that assumes the contest has only one winner, and that there can be at most one valid vote per ballot. It incorporates the “phantom to zombie” method of [1] for dealing with errors in the ballot manifest.

Input: Risk Limit α ; ballot manifest, announced winner w , losers set L and corresponding weighted vote shares $s_{w\ell}$ for each $\ell \in L$. Upper bound N on the number of ballots, where N is at least as large as the number of ballots listed in the manifest. Work threshold $K \leq N$.

- 1: Initialize probability ratios: $\forall \ell \in L : T_{w\ell} \leftarrow 1$
 - 2: Number of audited ballots: $n \leftarrow 0$
 - 3: $\mathcal{L} \leftarrow L$
 - 4: **while** $\mathcal{L} \neq \emptyset$ **do**
 - 5: Generate a random number i between 1 and N
 - 6: Look up the i th ballot in the ballot manifest and (attempt to) retrieve it
 - 7: **if** Ballot i is not in the ballot manifest or cannot be found **then**
 - 8: For every $\ell \in \mathcal{L}$, $T_{w\ell} \leftarrow T_{w\ell} * 2(1 - s_{w\ell})$
 - 9: **else if** Ballot i shows a vote for the winner w **then**
 - 10: For each $\ell \in \mathcal{L}$, $T_{w\ell} \leftarrow T_{w\ell} * 2s_{w\ell}$
 - 11: **else if** Ballot i shows a vote for loser $\ell \in \mathcal{L}$ **then**
 - 12: $T_{w\ell} \leftarrow T_{w\ell} * 2(1 - s_{w\ell})$
 - 13: **if** $T_{w\ell} \geq \frac{1}{\alpha}$ for any loser $\ell \in \mathcal{L}$ **then**
 - 14: $\mathcal{L} \leftarrow \mathcal{L} \setminus \{\ell\}$
 - 15: **if** the number of ballots inspected exceeds K , or optionally at any time **then**
 - 16: **STOP** the audit and perform a full manual recount.
 - 17: Declare election outcome correct—since all null hypotheses have been rejected.
-

At any stage, $P = \max_{\ell \in L} 1/T_{w\ell}$ is a conservative sequential P -value for the hypothesis that the reported winner w did not actually win the constituency.

Number of votes to be audited Consider an example of a 3-candidate contest with a single plurality winner. The candidates are Ram, Shyam and Janani. Their respective shares are recorded in the following table:

Ram	Shyam	Janani
20,000	30,000	50,000

In this case, the winner is Janani. Let us denote the winner-loser pairs as (j, r) for Janani and Ram and (j, s) for Janani and Shyam. The weighted vote shares are:

$$s_{jr} = \frac{v_j}{v_j + v_r} = \frac{50000}{50000 + 20000} = 0.714$$

$$s_{js} = \frac{v_j}{v_j + v_s} = \frac{50000}{50000 + 30000} = 0.625$$

We set the risk limit α at 5%. Every time the audit selects a ballot that shows a vote for Janani we multiply T_{jr} by $\frac{0.714}{0.5} = 1.428$ and T_{js} by $\frac{0.625}{0.5} = 1.25$. Therefore, the minimum sample size n to attain a risk limit $\alpha = 0.05$ satisfies

$$1.428^n \geq 20$$

and

$$1.25^n \geq 20$$

The smallest such n is $n = 14$. Hence, we need to audit at least 14 ballots—if they all show up votes for Janani, BRAVO will confirm the election outcome at risk limit 5%.

If the reported election results were accurate, on average we would see 50% of ballots for Janani, 30% for Shyam and 20% for Ram. [7] describe how to find the *Average sample number (ASN)*, the expected sample size necessary to reject all the null hypotheses, assuming the reported results are indeed correct. Stark’s online ballot-polling tool shows an ASN of 123 for this example. There have been numerous improvements to the efficiency of Risk-Limiting Audits, any of which could easily apply to India’s simple electoral system.¹¹

The next section explains how to audit the overall parliamentary winner by an efficient combination of single-constituency audits. It requires independent, sequentially valid P -values $\{P_i\}$ for the hypotheses that the reported outcome in constituency i is incorrect. It does not require the P -values to be obtained using the same method. For instance, some constituencies could use ballot polling and others could use transitive ballot-level comparison audits.

4 Auditing the Overall Parliamentary Winner

A party or a coalition needs a majority of the seats in the Lower House of Parliament to form a new government. The total number of seats is 543, so to win, a party or coalition needs at least 272 seats. The audit needs to confirm that the reported winning party or coalition truly won at least 272 seats. (The particular seats the reported winner won is immaterial to whether they won overall.) If party w supposedly won $M \geq 272$ constituencies, then for a different party to have won in fact, the reported outcome must be wrong in at least $m = M - 271$ of the constituencies that w supposedly won. This condition is necessary but not sufficient for the parliamentary outcome to be wrong: if w in fact won some constituencies it was reported to have lost, the outcome could be wrong in m constituencies w supposedly won and yet w could still be the overall winner.

Let W denote the set of constituencies w reportedly won. Then $|W| \geq 272$ and $m = |W| - 271$, where $|W|$ denotes the cardinality of the set W . If there is

¹¹ See for example <https://github.com/pbstark/S157F17/blob/master/kaplanWald.ipynb> and <https://github.com/pbstark/S157F17/blob/master/pSPRTnoReplacement.ipynb>.

no set of constituencies $C \subset W$ with $|C|=m$ for which w lost in every $c \in C$, w must have won overall.

Let α denote the overall risk limit, and let P_c denote a P -value for the hypothesis that the reported outcome in constituency c is wrong. We suppose that the audits in different constituencies rely on independently selected random samples of ballots, so the P -values $\{P_c\}$ are independent random variables. If the reported outcome in constituency c is incorrect, the probability distribution of P_c is stochastically dominated by a uniform distribution. That is, $\Pr\{P_c \leq p\} \leq p$ if the reported outcome in constituency c is wrong.

Let C denote a set of constituencies. *Fisher's combining function* for a set of P -values $\{P_c\}_{c \in C}$ is

$$\chi^2(C) \equiv -2 \sum_{c \in C} \ln P_c. \quad (4)$$

If the P -values $\{P_c\}$ are independent and all the null hypotheses are true, the probability distribution of $\chi^2(C)$ is stochastically smaller than a chi-square distribution with $2|C|$ degrees of freedom.¹² That is, if the reported outcome in every constituency $c \in C$ is wrong,

$$\Pr\{\chi^2(C) \geq \chi_{2|C|}^2(1 - \alpha)\} \leq \alpha, \quad (5)$$

where $\chi_{2|C|}^2(1 - \alpha)$ is the $1 - \alpha$ quantile of the chi-square distribution with $2|C|$ degrees of freedom.

Let W denote the set of constituencies the reported winning party allegedly won, and let $m \equiv |W| - 271$. Then m is the parliamentary “margin in constituencies: the reported winner really won overall unless there are at least m constituencies where the overall winner reportedly won, but in fact lost. For any set C of constituencies, let \mathcal{C}_m denote the set of all subsets of C with cardinality m . The overall auditing strategy is to test whether there is any subset of m constituencies in W where the reported parliamentary winner actually lost. If there is no such subset, the reported winner must actually have won overall. We test that hypothesis by examining all such subsets. For each such subset, we use the separate audits of the constituencies to construct, via Fishers combining function, a test of the hypothesis that the reported winner lost in all m constituencies.

If the audit of some constituency c leads to a full hand count that confirms the result, then all subsets of size m that contain c can be eliminated from further consideration: the reported winner cannot have lost in all m constituencies in such a subset, because it actually won in c .

For any collection \mathcal{U} of sets of constituencies, define $\mathcal{U}(c) \equiv \{U \in \mathcal{U} : U \ni c\}$, all sets of constituencies in \mathcal{U} that contain c . (These will be the collections eliminated from further consideration if a full count of c confirms the result in c .)

With these definitions, our audit procedure is as follows:

¹² See, e.g., [9].

1. Select an overall risk limit α for the parliamentary outcome and select an auditing method for each constituency $c \in W$, e.g., BRAVO or LSKM.
2. Perform the following audit.

Algorithm 2 Audit of overall election result.

Input: W , the constituencies reportedly won by w ; α , the risk limit; a sequential auditing method for each $c \in W$

- 1: Set $\mathcal{U} = \mathcal{W}_m$, the collection of all subsets of W of cardinality $m = |W| - 271$.
 - 2: **while** $\mathcal{U} \neq \emptyset$ **do**
 - 3: Augment the sample in one or more constituencies c ¹³
 - 4: Find the constituency-level P -values $\{P_c\}$ using the auditing method pre-specified for each c
 - 5: **if** $\exists C \in \mathcal{U}$ such that $\forall c \in C$ a full hand count revealed a different winner **then**
 - 6: Perform a full hand count of the entire election; report the resulting outcome, and stop.
 - 7: **for all** constituencies $c \in \cup \mathcal{U}$ **do**
 - 8: **if** c has been fully counted by hand, and the count confirmed the reported outcome in c **then**
 - 9: $\mathcal{U} \leftarrow \mathcal{U} \setminus \mathcal{U}(c)$
 - 10: $\mathcal{U} \leftarrow \mathcal{U} \setminus \{C \in \mathcal{U} : \chi^2(C) \geq \chi^2_{2|C|}(1 - \alpha)\}$
-

3. If the loop terminates with $\mathcal{U} = \emptyset$, the audit has confirmed the parliamentary outcome at risk limit α .

Proof that the algorithm above is a RLA of the parliamentary outcome. We show that if the parliamentary outcome is wrong, the chance that the

¹³ BRAVO and LSKM produce sequentially valid P -values. That is, the chance that the infimum of the P -value for constituency c over all sample sizes is less than or equal to p is itself less than or equal to p , if the outcome is incorrect in constituency c . For that reason, every rule for increasing sample sizes in constituency c results in a valid P -value for constituency c . Because the samples are selected independently in different constituencies (although sample *sizes* are dependent), the composite P -value from Fishers combining function remains valid. Thus, the rule for increasing sample sizes when the risk limit has not been attained could be as simple as “increase every n_c by 1 ballot, or it could be designed to minimize the expected total amount of auditing required, for instance, by preferentially increasing the sample size in constituencies with large margins and taking into account differences in auditing methods in different jurisdictions (ballot polling versus transitive ballot-level comparison). All else equal, when the outcome is correct, auditing an additional ballot is expected to decrease the P -value more, the larger the true margin is. Similarly, all else equal, auditing an additional ballot in a jurisdiction conducting a transitive ballot-level comparison RLA is expected to decrease the P -value more than auditing an additional ballot in a jurisdiction conducting a ballot-polling RLA. It is always permissible to perform a full hand count in any constituency rather than increase the sample size incrementally: anything that increases the chance of a full count cannot increase the risk.

audit stops without a full manual tally of every constituency is at most α . If the parliamentary outcome is wrong, the reported winner is wrong in every $c \in C$ for some $C \in \mathcal{W}_m$. Suppose there is such a C . If the audit leads to hand counting every $c \in C$, step (5) ensures that there will be a full manual tally of the entire election. Therefore, there will be a full manual tally unless C is removed from \mathcal{U} . There are two places that sets of constituencies can be removed from \mathcal{U} : step (9) and step (10). Step (9) cannot remove C from \mathcal{U} , because, by assumption, hand-counting any $c \in C$ would belie the reported outcome in c . Therefore, the chance that C is *not* fully hand counted is at most the chance that step (10) removes C from \mathcal{U} . But, by construction (through Fisher’s combining function applied to the independent constituency-level P -values), that chance is not larger than α . If there is more than one $C \in \mathcal{W}_m$ for which every reported outcome is wrong, the audit must erroneously remove *all* of them at step (10). But the chance of erroneously removing all of them cannot be larger than chance of removing any one of them individually, which is in turn at most α .

5 Conclusion and Future Work

We have presented an approach to conduct risk-limiting audits of the national outcome of Indian elections by combining audits conducted in different constituencies using independent samples. Within a given constituency, the audit could use ballot polling, or—with an initial step of sorting VVPATs—transitive ballot-level comparisons. The sets of constituencies are constructed in such a way that for the reported parliamentary outcome to be wrong, the reported outcome must be wrong in *every* constituency in at least one of the sets. If there is strong statistical evidence that there is no set of constituencies in the collection for which every reported outcome is wrong, that confirms the national parliamentary outcome. In future research we will address how to schedule increases in sample sizes in different constituencies to minimize the total expected workload, taking into account the reported margins in different constituencies and the auditing methods used in different constituencies.

Acknowledgments

Many thanks to Andrew Conway, Archanaa Krishnan, Chittaranjan Mandal, Sandeep Shukla, Nicholas Akinyokun, Peter Stuckey and Poorvi Vora for valuable suggestions on this work.

References

- [1] Bañuelos, J., Stark, P.: Limiting risk by turning manifest phantoms into evil zombies. arXiv preprint arXiv:1207.3413 (2012)
- [2] Benaloh, J., Jones, D., Lazarus, E., Lindeman, M., Stark, P.: Soba: Secrecy-preserving observable ballot-level audit. *EVT/WOTE* **11** (2011)
- [3] Calandrino, J., Halderman, J., Felten, E.: Machine-assisted election auditing. *EVT* **7**, 9–9 (2007)
- [4] Hall, J., Miratrix, L., Stark, P., Briones, M., Ginnold, E., Oakley, F., Peaden, M., Pellerin, G., Stanionis, T., Webber, T.: Implementing risk-limiting audits in california. *EVT/WOTE* (2009)
- [5] Kroll, Halderman, A., Felten, E.: Efficiently auditing multi-level elections (2014)
- [6] Lindeman, M., Stark, P.: A gentle introduction to risk-limiting audits (2012)
- [7] Lindeman, M., Stark, P., Yates, V.: Bravo: Ballot-polling risk-limiting audits to verify outcomes (2012)
- [8] Ottoboni, K., Bernhard, M., Halderman, A., Rivest, R., Stark, P.: Bernoulli ballot polling: A manifest improvement for risk-limiting audits. *Voting '19* (2019)
- [9] Ottoboni, K., Stark, P., Lindeman, M., McBurnett, N.: Risk-limiting audits by stratified union-intersection tests of elections (suite). In: *International Joint Conference on Electronic Voting*. pp. 174–188. Springer (2018)
- [10] Schürmann, C.: A risk-limiting audit in denmark: A pilot. Springer (2016)
- [11] Stark, P.: Conservative statistical post-election audits **2**(2), 550–581 (2008)
- [12] Stark, P.: Risk-limiting post-election audits: *P*-values from common probability inequalities. *IEEE Transactions on Information Forensics and Security* **4**, 1005–1014 (2009)
- [13] Stark, P.: Super-simple simultaneous single-ballot risk-limiting audits. USENIX Association Berkeley, CA, USA ©2010 (2010)
- [14] Stark, P.: An introduction to risk-limiting audits and evidence-based elections: Testimony to the Little Hoover Commission (2018)
- [15] Stark, P., Teague, V.: Verifiable european elections: risk-limiting audits for dhondt and its relatives. *USENIX Journal of Election Technology and Systems (JETS)* **1**(3), 18–39 (2014)
- [16] Stark, P., Wagner, D.: Evidence-based elections. *IEEE Security & Privacy* **10**(5), 33–41 (2012)
- [17] Supreme Court of India: Civil appeal no.9093 of 2013—Dr. Subramanian Swamy vs. Election Commission of India (10 2013)
- [18] Wald, A.: Sequential tests of statistical hypotheses. *Annals of Applied Statistics* **16**(2), 117–186 (1945)
- [19] Wolchok, S., Wustrow, E., Halderman, J., Prasad, H., Kankipati, A., Sakhamuri, S., Yagati, V., Gonggrijp, R.: Security analysis of india’s electronic voting machines. In: *Proceedings of the 17th ACM conference on Computer and communications security*. pp. 1–14. ACM (2010)

Risk-Limiting Tallies

Wojciech Jamroga^{1,2}, Peter B. Roenne¹, Peter Y. A. Ryan¹,
and Philip B. Stark³

¹ Interdisciplinary Centre for Security, Reliability, and Trust, SnT, University of Luxembourg

² Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland

³ Department of Statistics, University of California, Berkeley, US
{wojciech.jamroga,peter.roenne,peter.ryan}@uni.lu,
stark@stat.berkeley.edu

Abstract. Many voter-verifiable, coercion-resistant schemes have been proposed, but even the most carefully designed systems necessarily leak information via the announced result. In corner cases, this may be problematic. For example, if all the votes go to one candidate then all vote privacy evaporates. The mere possibility of candidates getting no or few votes could have implications for security in practice: if a coercer demands that a voter cast a vote for such an unpopular candidate, then the voter may feel obliged to obey, even if she is confident that the voting system satisfies the standard coercion resistance definitions. With complex ballots, there may also be a danger of “Italian” style (aka “signature”) attacks: the coercer demands the voter cast a ballot with a specific, identifying pattern.

Here we propose an approach to tallying end-to-end verifiable schemes that avoids revealing all the votes but still achieves whatever confidence level in the announced result is desired. Now a coerced voter can claim that the required vote must be amongst those that remained shrouded. Our approach is based on the well-established notion of Risk-Limiting Audits, but here applied to the tally rather than to the audit. We show that this approach counters coercion threats arising in extreme tallies and “Italian” attacks. We illustrate our approach by applying it to the Selene scheme, and we extend the approach to Risk-Limiting Verification, where not all vote trackers are revealed, thereby enhancing the coercion mitigation properties of Selene.

Keywords: End-to-end verifiability, risk-limiting audits, plausible deniability, coercion resistance.

1 Introduction

Many verifiable voting schemes have been proposed that are designed to give a high level of resistance against coercion or vote buying [4,12,8,20,22]. However, it is typically assumed that little can be done about coercion threats in case of extreme outcomes, e.g., no or few votes for some candidates. Unfortunately, such situations do happen in real elections. If there is a (perceived) risk that

candidate X will get no votes, and the coercer tells the voter to vote for X , then the voter may feel obliged to comply even if the voting scheme satisfies the standard definitions of coercion resistance. The possibility is more dangerous than it seems at the first glance. True, coercing for the low support candidate X is unlikely to get him or her win. However, the coercer can use X to construct what is effectively an abstention attack, and take away the votes from the main opponent of his preferred candidate. If the coercer prefers candidate A , he can help him win by coercing supporters of B to vote for X .

Another difficulty that may arise is that of so-called “Italian”-style attacks, also known as “signature” attacks: if the voting method allows for a large number of distinct ways of filling out the ballot, a coercer may require the voter to fill out the ballot with a distinctive pattern allowing it to be uniquely identified with high probability in the final tally. This is especially an issue with long, complex ballots and with preferential voting schemes. It can be countered by, for example, using homomorphic tallying techniques to compute the overall result without revealing the individual ballots, but this is computationally intensive and even then may leak some critical information [27].

Here, we show that risk-limiting audit techniques [16] can be adapted to achieve whatever level of confidence in the outcome is required while ensuring that a proportion of the ballots remain shrouded. This allows us to significantly enhance the coercion resistance of verifiable schemes. The Risk Limiting Tally (RLT) approach that we present here provides a simple way to guarantee voters plausible deniability against the above attacks: a coerced voter, who did not cast the ballot the way the coercer had demanded, can simply claim that the required ballot is amongst those unrevealed.

We also present a variant of the idea, *Risk Limiting Verification* (RLV), where we ensure that a proportion of verification tokens remain unrevealed. The basic version of the Selene scheme [20] has the drawback that the the coercer can claim that the fake tracker provided by a voter is his own. We describe how RLV mitigates this.

A possible objection to RLT is that it is “undemocratic” not to count all votes. However, the method allows the electoral outcome (i.e. the winner or winners) to be ascertained to any desired level of statistical certainty. Moreover, the sample of ballots will be drawn in such a way that every cast vote has an equal chance of being in the revealed sample, so there is no lack of fairness. RLTs are related to *random sample voting* (RSV), due to Chaum [7], except that there the sample is drawn from the set of eligible voters, rather than from the cast votes. If anything, RLTs seem to be more democratic in that in RSV voters who are not chosen might well feel excluded. Furthermore, in RLTs we are able to adjust the sample size after voting to achieve the desired confidence level. We note also that some tally algorithms, e.g. some forms of STV, intrinsically involve a probabilistic element.

The outline of this paper is as follows. In section 2 we discuss coercion-resistant and verifiable voting schemes in general, and Selene in particular. Section 3 briefly recalls Risk-Limiting Audits, and Section 4 introduces the techniques for RLTs.

We present the actual protocol in Section 5, and a brief security discussion in Section 6. Section 7 discusses the risk-limiting verification. Related work is presented in Section 8, and we conclude in Section 9.

1.1 Contribution

In this paper we present several contributions:

1. The use of using risk-limiting techniques to shroud a proportion of votes, improving coercion resistance while achieving whatever confidence level is required.
2. A novel extension of Risk-Limiting Audit (RLA) techniques to handle the situation in which we do not have an initial null hypothesis (reported outcome).
3. A new test statistic for RLAs with operating characteristics that do not depend on the reported votes, only on the reported winner(s).
4. Protocols to enable RLT for most end-to-end verifiable schemes, including strategies to ensure plausible deniability whatever the vote distribution is.
5. Extension of the approach to Risk-Limiting Verification: the shrouding of a randomly selected subset of verification tokens to improve coercion resistance, in particular for the Selene scheme.

2 Coercion-Resistant and Verifiable Voting

We set the scene by recalling the concepts of End-to-End Verifiability [21] and coercion-resistance [4,12], and showing an example scheme designed to balance the two requirements.

2.1 An Outline of End-to-End Verifiable Voting

The use of digital technologies to record and process votes might provide efficiency and convenience, but it can also bring serious new threats, in particular, virtually undetectable ways to manipulate votes on a large scale. These concerns motivated the development of *End-to-End Verifiable* (E2E V) voting schemes. Such schemes provide the voters with means to confirm that their vote is accurately included in the tally, without opening up possibilities of coercion or vote-buying. This is usually accomplished by creating an encryption of the vote at the time of casting, and posting this to a public *Bulletin Board* (*BB*). Voters can then confirm that their “receipt,” i.e., the encrypted vote, appears correctly on the *BB*.

Once we have consensus on the correct set of encrypted votes, these can be processed in a verifiable fashion to calculate the outcome in a way that does not compromise the privacy of the votes. For instance, the encrypted votes might be put through a sequence of verifiable re-encryption mixes and then verifiably decrypted, allowing anyone to compute the result. Alternatively, the encrypted votes might be tallied under encryption, exploiting the homomorphic properties of the encryption algorithm, and the final result decrypted. To complete the assurance argument we need some additional ingredients:

- The voter needs to be confident that her intended vote is correctly encrypted in her receipt.
- We need to prevent ballot stuffing, i.e., we need to ensure that only legitimately cast votes appear in the list of receipts on the *BB*, and only one per voter.
- We need to know that “enough” voters check that their intended votes are correctly encrypted, and that their encrypted votes appear on *BB*.
- We need dispute-resolution mechanisms in place to ensure that if voters detect (or claim to detect) problems, the culprit can be identified and appropriate action taken.

Typically the first point is addressed by some form of cut-and-choose protocol, e.g. *Benaloh Challenge* [5], or a more sophisticated approach such as Neff’s *MarkPledge* scheme [2]. Ballot stuffing is usually countered either by procedural measures in the polling station, or by requiring that receipts be digitally signed by the voters. The former does not provide *universal eligibility verifiability* while the latter can but requires infrastructure to equip voters with signing keys.

We will not delve deeper into how the various E2E V schemes work but rather assume that the correct set of encrypted votes is posted to the *Bulletin Board*.

2.2 Ballot Privacy, Receipt-Freeness and Coercion Resistance

Ballot privacy is often defined using anonymity style definitions as originally proposed in [23]. Informally, consider two instances of the system, one in which *A* votes for *X* and *B* for *Y*, and the other in which the votes are swapped. If the attacker is unable to distinguish these two instances then the system is deemed to satisfy *ballot privacy*. More formal definitions can be found, for example, in [10]. Note that even in extreme cases, for example when all voters vote for *X*, such a system will satisfy the above definition, even though in that case the attacker knows precisely how each voter voted.

It was later realised that simple notions of ballot privacy in the presence of a passive attacker are not enough. For E2E V schemes we have to worry about ways that the voter might be able to prove her vote to a third party. This motivates the requirement for *receipt-freeness* [4]: the voter cannot acquire evidence that would enable her to construct a proof to a third party as to how she voted.

In the face of a yet more active attacker who might interact with the voter before, during and after voting, potentially issuing detailed instructions and requiring the voter to reveal credentials, ephemeral random values etc, we need even stronger notions. This threat model motivates the property of *coercion resistance* for which many different definitions have been proposed [12], reflecting various subtle distinctions. We will adopt the following definition, informally stated:

A voting system *S* is coercion resistant if, for all $c \in \mathcal{C}$ there exists a voter strategy ψ such that for all attacker strategies ϕ , the voter can cast her intended vote *c* and the attacker cannot tell that she did not obey his instructions.

Such a style of definition appears to be the most powerful in that it captures the privacy failure in the case of unanimous votes, forced abstention and randomisation attacks.

2.3 Selene

We now give a sketch of how voter-verification is achieved in the Selene voting protocol. Full details can be found in [20]. In Selene, the verification is much more direct and intuitive than is the case for conventional E2E V systems: rather than checking for the presence of her encrypted vote on the *BB*, the voter checks her vote in cleartext in the tally on the *BB* identified by a secret, deniable tracker.

During the setup phase the set of distinct trackers are posted on the *BB*, verifiably encrypted and mixed and then assigned to the voters according the resulting secret permutation. This ensures that each voter is assigned a unique, secret tracker.

For each encrypted tracker, a trapdoor commitment is created for which the voter holds the secret trapdoor key. In essence this is the “ β ” term of an El Gamal encryption of the tracker, where the “ α ” term is kept secret for the moment.

Voting is as usual: an encryption of the vote is created, and sent to the server for posting to the *BB* against the voter (pseudo)Id. Once we are happy that we have the correct set of validly cast, encrypted votes, we can proceed to tabulation: the (encrypted vote, tracker) pairs are put through verifiable, parallel re-encryption mixes and decrypted, revealing the vote/tracker pairs in plaintext.

Later, the α terms are sent via an untappable channel to the voters to enable them to open the commitment using their secret, trapdoor key. If coerced, the voter can generate a fake α that will open her commitment to an alternative tracker pointing to the coercer’s choice. With the trapdoor, creating such a fake α is computationally straightforward. On the other hand, computing a fake α that will open the commitment to a given, valid tracker is intractable without the trapdoor. Thus, assuming that the voter’s trapdoor is not compromised, the α term is implicitly authenticated by the fact that it opens to a valid tracker.

3 Risk-Limiting Audits

A *risk-limiting audit* (RLA) [16] of a reported election outcome is any procedure that has a known minimum chance of correcting the reported outcome if the outcome is wrong (and that cannot render a correct outcome incorrect). In this case, the *outcome* means the winner or winners, not the precise tally. The reported outcome is *correct* if it is the outcome that an accurate manual tally of the underlying voter-verified records would show.⁴ The maximal chance that the procedure will fail to correct an outcome that is wrong is the *risk limit*.

⁴ The trustworthiness of the underlying records should be assessed by a *compliance audit* [25]. A RLA that relies on an untrustworthy record cannot reliably assess whether outcomes reflect how voters voted.

RLAs generally pose auditing as a sequential test of the hypothesis that the reported outcome is incorrect. The audit continues to examine more ballots until either the hypothesis is rejected or the audit has conducted a full manual tally. The use of sequential tests enables RLAs to stop as soon as there is convincing evidence that the reported outcome is correct, reducing the number of ballots the audit inspects.

RLAs check reported outcomes while RLTs determine what outcomes to report. However, similar sequential testing methods can allow RLTs to stop the tally (of a random permutation of the ballots) as soon as there is convincing statistical evidence of the electoral outcome, which the RLT then reports. A RLT declares “either this is the correct outcome, or an event occurred that had probability no larger than α ,” where $\alpha \in (0, 1)$ is any pre-specified risk limit. Minimizing the number of ballots that must be tallied maximizes the number of ballots kept shrouded, improving privacy and coercion-resistance.

There are two general strategies for RLAs: *ballot-polling* and *comparison*. Ballot-polling manually examines randomly selected ballots for evidence of who won. A comparison audit has three steps: first, the voting system must commit to its interpretation of physically identifiable individual ballots or groups of ballots comprising all ballots validly cast in the election. Second, auditors check that the exported data reproduces the reported results. Third, auditors compare the manual interpretation of a random sample of ballots or groups of ballots to the voting system’s interpretation. Further, “hybrid” methods combine ballot-polling for some groups of ballots and comparisons for other groups; see [18].

Comparison audits require auditors to know how the equipment interpreted the ballots, so they are not suitable for RLTs, where we seek evidence about who won just from a subset of the shrouded votes. Below, we show how a new procedure for ballot-polling RLAs can be adapted for RLTs.

4 Risk-Limiting Tallies

We propose a simple modification of the way that votes are tallied to address the issues outlined in the introduction. Rather than tallying all votes straight-off, the election authority reveals the votes for a random sequence of encrypted ballots, continuing until the sample gives the acceptable level of risk in the outcome (i.e., who won). If the true margin of victory is not too small, the outcome can be determined with high confidence (i.e., low risk) while leaving a substantial number of ballots unopened, thus allowing a voter to claim that they cast the ballot required by the coercer even if such a ballot was not revealed during the partial tally. The approach is thus inspired by the idea of Risk-Limiting Audits (RLAs), [24,16], but here we apply the approach to determining the correct outcome rather than checking whether a reported outcome is correct. That difference turns out to have surprising statistical implications; in particular, larger sample sizes are generally required to control the risk to the same level.

RLAs test the null hypothesis that the reported winner(s) did not actually win, rather than determine the correct outcome *ab initio*. Moreover, the operating

characteristics of existing RLAs depend on the reported results. For instance, *comparison audits* test whether the reported margin overstated the true margin by enough to cause the reported winners to be incorrect. Previous methods for *ballot-polling* audits, such as BRAVO [15] test the hypothesis that the reported outcome is wrong against the alternative that the reported vote shares are nearly correct.

For RLTs, we do not have reported results to leverage, so we need a new approach. Section 4.1 presents a probability inequality; Section 4.2 applies it to produce a new sequential ballot-polling test, the engine for the RLT scheme presented in Section 5 based on the Selene E2E V protocol.

4.1 Tests for the Mean of a Non-Negative Population

Extant methods for RLAs generally involve the reported results in some way. Here, we present a new sequential method to determine with high confidence who won, without specifying a particular alternative hypothesis. The method applies to plurality (including vote-for- k), majority, and super-majority social choice functions, but we present the method in detail only for plurality contests.

Our RLT method is based on tests about the mean of a non-negative population. Consider a population of N items, each labeled with a non-negative number.⁵ Let $x_i \geq 0$ be the label of item i , $i = 1, \dots, N$. Let $\mu \equiv \frac{1}{N} \sum_{i=1}^N x_i$ be the mean of the labels. Moreover, let t denote the hypothesized value of the population mean μ .

We sample items at random, sequentially, without replacement, such that the (conditional) probability that item k is selected in the j th draw is $\frac{1}{N-j+1}$, given that item k was not selected before the j th draw. X_j denotes the number on the label of the item selected on the j th draw. Define $S_j \equiv \sum_{k=1}^j X_k$, $\tilde{S}_j \equiv S_j/N$, and $\tilde{j} \equiv 1 - (j-1)/N$. Let

$$Y_n \equiv \int_0^1 \prod_{j=1}^n \left(\gamma \left[X_j \frac{\tilde{j}}{t - \tilde{S}_{j-1}} - 1 \right] + 1 \right) d\gamma. \quad (1)$$

It has been shown in [11] that if $\mu = t$ (i.e., if the null hypothesis is true), then $(Y_j)_{j=1}^N$ is a nonnegative closed martingale with expected value 1. Kolmogorov's inequality then implies that for any $J \in \{1, \dots, N\}$ and any $p \in (0, 1)$,

$$\Pr \left(\max_{1 \leq j \leq J} Y_j(t) > 1/p \right) \leq p.$$

This can be used as the basis of a ballot-polling RLA that does not require a reference tally, as we show below. The same result holds for sequential sampling *with* replacement, re-defining $\tilde{S}_j \equiv 0$ and $\tilde{j} \equiv 1$ (the limit of the finite-population result as $N \rightarrow \infty$). We also note that [11] provides a recursive algorithm for computing the integral (1).

⁵ In our case, the items will be ballots, and their labels will represent votes; see Section 4.2.

4.2 Risk-Limiting Tallies

Consider plurality contests that allow each voter to vote for $k \geq 1$ of C candidates. The winner(s) are the k candidates who receive the most votes. We ignore the possibility of ties; they are an easy extension. Majority and super-majority are straightforward generalizations; see [24].

Candidate w is one of the winners if w received more votes than at least $C - k$ other candidates. In general, some ballots will have invalid votes or votes for other candidates. Consider a single pair of candidates, w and ℓ . Let N_w denote the number of ballots in the population that show a vote for w but not for ℓ ; let N_ℓ denote the number of ballots in the population that show a vote for ℓ but not for w , and let $N_u \equiv N - N_w - N_\ell$ denote the number of ballots that show a vote for neither w nor ℓ or show votes for both w and ℓ .

Let W_j be the number of items labeled with w selected on or before draw j ; and define L_j analogously. The probability distributions of those variables depend on N_w , N_ℓ , and N_u , even though we only care about one parameter, $N_w - N_\ell$. Now $N_w \leq N_\ell$ if and only if $N_w + N_u/2 \leq N_\ell + N_u/2$. Since $N_\ell + N_u/2 = N - (N_w + N_u/2)$, we have $N_w + N_u/2 \leq N - (N_w + N_u/2)$. We can now divide by N to obtain $\frac{N_w + N_u/2}{N} \leq 1 - \frac{N_w + N_u/2}{N}$ from which we get

$$\frac{N_w + N_u/2}{N} \leq \frac{1}{2}. \quad (2)$$

Let

$$\mu_{w\ell} \equiv \frac{1 \times N_w + \frac{1}{2} \times N_u + 0 \times N_\ell}{N}.$$

This is the mean of a population derived from re-labeling each vote for w as 1, each vote for ℓ as 0, and the rest as 1/2. The mean of this population is greater than 1/2 iff w received more votes than ℓ . We can test the hypothesis $\mu_{w\ell} \leq 1/2$ (i.e., w did not beat ℓ) using the martingale-based test above by simply treating the sampled ballots that way: every ballot with a vote for w (but not ℓ) counts as 1, every ballot with a vote for ℓ (but not for w) counts as 0, and invalid ballots, ballots with votes for other candidates, and ballots with votes for both w and ℓ count as 1/2.

To determine the set of winners, we sequentially test the collection of $C(C-1)$ hypotheses

$$\{H_{w\ell} : \mu_{w\ell} \leq 1/2, w = 1, \dots, C; \ell = 1, \dots, C; w \neq \ell\}, \quad (3)$$

stopping when either

- there is a set \mathcal{W} of cardinality k such that we have rejected the hypothesis $\mu_{w\ell} \leq 1/2$ for every (w, ℓ) with $w \in \mathcal{W}$ and $\ell \notin \mathcal{W}$, or
- we have examined a too high percentage of votes from the privacy point of view, in which case the sampling strategy is abandoned and different means are used to determine with certainty who won, see Section 5.1 for details.

Proposition 1. *If every hypothesis is tested at level α , the probability that this algorithm misidentifies the set of winner(s) is at most $k(C - k)\alpha$.*

Proof. The approach misidentifies one or more winners iff it terminates in the first branch, but \mathcal{W} is not the set of winners: $\exists w \in \mathcal{W}, \ell \notin \mathcal{W}$ s.t. $\mu_{w\ell} \leq 1/2$. In a RLA, a wrong outcome can only be confirmed if *every* true null hypothesis is erroneously rejected. In contrast, in a RLT, a wrong outcome can be confirmed if just one particular true null hypothesis is rejected: the hypothesis that the candidate with the $k + 1$ st highest vote share got fewer votes than the candidate with the k th highest vote share.

There are $C(C - 1)$ hypotheses $\{H_{w,\ell}\}$ in all, of which $C(C - 1)/2$ are true. Of the true null hypotheses, those whose erroneous rejection would make the reported outcome wrong are the $k(C - k)$ that compare the vote share of a candidate in \mathcal{W} to the vote share of a candidate in \mathcal{W}^c : if none of those is erroneously rejected, the set of winners is correct. Observe that if we used the logical implications of the statistical rejections to entail rejections of other hypotheses—for instance, $H_{w\ell} \cap H_{\ell k} \rightarrow H_{wk}$ —this would not be true. Therefore, a Bonferroni multiplicity adjustment of $k(C - k)$ certainly suffices. Note that this may be conservative as an estimate, because there are logical dependencies among the hypotheses. \square

The aim of the sampling is to test the hypothesis “ $\mu_{w\ell} \leq 1/2$.” Rejecting $\mu_{w\ell} \leq 1/2$ means proving with risk at most α that w won the pairwise contest with ℓ .

Proposition 2. *If we reject $\mu_{w\ell} \leq 1/2$ at significance level α and reject $\mu_{\ell m} \leq 1/2$ at significance level α , then we reject $\mu_{wm} \leq 1/2$ at significance level α .*

Proof (sketch). This transitivity property follows from the monotonicity of the P -values in the number of votes for each candidate, at each sample size j . \square

4.3 Sample sizes

Because the underlying statistical test is sequential, the audit can start by looking at a single ballot selected at random, calculate the p -values for all not-yet-rejected null hypotheses, and continue to increase the sample one ballot at a time until the risk limit has been met. However, depending on the desired risk limit, the RLT will not be able to terminate until some minimum number of ballots has been tallied.

The minimum sample sizes required to identify the winner with a maximum error rate of α are given in Table 1, for sampling without replacement, for a plurality contest with 2 candidates and a plurality contest with 10 candidates. The sample sizes listed are exactly those that would be required if the votes were unanimously for one candidate; if more than one candidate receives votes, the sample size becomes random and becomes stochastically larger.

Similarly, if a fraction u of ballots do not have a valid vote for any candidate, the sample size will also be random, and the expected sample size will grow by a factor of $1/(1 - u)$. For instance, if 10% of ballots have no vote and 90% of ballots have a vote for candidate A in a 10-candidate plurality election, the expected sample size to identify the winner with risk limit 0.1% is $17/0.9 = 18.9$ ballots.

The closer the vote is to unanimous, the fewer ballots need to be revealed (the distribution is stochastically smaller the more nearly unanimous the vote). I.e., the protection a RLT offers is greatest when the risk is greatest.

For a two-candidate plurality election, only one of the two null hypotheses $\mu_{\ell_m} \leq 1/2$ can be true; thus, no multiplicity adjustment is needed. (This is consistent with the formula $k(C - k) = 1 \times (2 - 1) = 1$.) For a 10-candidate plurality election, the Bonferroni adjustment factor is $1 \times (10 - 1) = 9$. As the table shows, if the vote is (nearly) unanimous, the number of ballots required to identify the winner with negligible error probability is small: 35 suffices to have an error probability less than 10^{-9} for a two-candidate contest, and 38 suffices for a 10-candidate contest. Because the risk drops by an order of magnitude with an increase in sample size of about 4 ballots when the vote is (nearly) unanimous, the penalty for multiplicity is low in absolute terms. If the RLT sample is drawn *without* replacement, the expected sample sizes required to attain a given risk are smaller—but not by much unless the total number of ballots is small.

candidates	α									
	10^{-1}	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}	
2	5	9	13	17	21	24	28	31	35	
10	9	13	17	20	24	27	31	34	38	

Table 1. Minimum sample sizes to identify the winner of a two-candidate plurality contest and a 10-candidate plurality contest at risk limit α , for sampling with replacement. Actual sample sizes approach these minima (with high probability) as voter preferences approach unanimity.

5 Incorporating RLT in E2E V Voting Protocols

RLTs can be used in a straightforward way with any E2E V scheme in which the set of encrypted votes appears on a Bulletin Board (*BB*) and is applicable to either remote or in-person voting. The encryption should be homomorphic and probabilistic: for instance, ElGamal can be used. Helios [1], Prêt à Voter [8], Selene [20], etc., would all be amenable.

Conceptually, we can start with a random permutations of the encrypted votes and take samples from left to right, opening more ballots as required. The verifiable shuffles used in many schemes naturally give us a random permutation. However, we must be careful about simply taking the permutation output of the underlying scheme’s shuffles, as there may be opportunities to manipulate this and bias the sampling. The sampling must be truly random and demonstrably outwith the control of any entity. This brings us to the challenge of *certifiable randomness*, which arises in many contexts: lotteries, voting, auctions, public ledgers etc. A number of approaches have been proposed, for example using a seed derived from a hash of prices of previously agreed stock market options at

an agreed future time. Alternative approaches involve combining random values previously committed by a number of independent entities. Algorand [17] adopts such an approach combined with the use of verifiable random functions. Another possibility is to derive the seed from a cryptographic hash of suitable data posted to the *BB*. RLAs have employed seeds generated in a public ceremony of dice rolling. We might rely on a trusted third party such as the NIST random beacon service. For the purposes of this paper do not specify a particular approach but leave it for the stakeholders to select.

Sampling with replacement can be implemented straightforwardly by performing further mixes between samplings.

5.1 Guaranteeing Plausible Deniability

For most elections, the RLT approach will naturally leave a good proportion of unrevealed votes. However, there will be cases where the winning margins are narrow, and thus the RLT might result in all or almost all votes being revealed. It is not enough for a system to be (objectively) coercion resistant, it must also be seen as coercion resistant. Thus, for the RLT approach to be effective, we must ensure that the voters will never be, nor expect to be, in a situation in which plausible deniability fails. In this section we identify such situations, and describe some strategies to deal with the potential vulnerability.

Of course, a close run referendum will not be a problem, but a problematic scenario is a close margin between candidates X and Y , along with a low-support candidate Z . This could result in a full count where the low score Z opens up the possibility of coercion. We have already indicated that coercion for Z is possibly harmful for the outcome of the election, as it can be used to decrease the number of votes that either X or Y gets. Note also, again, that this kind of coercion is feasible not only when the voter *knows* (e.g., from polls) the a close run will occur. In many cases, it suffices that the voter thinks it *might* happen to get her worried and vulnerable to threats. We propose that, in such circumstances, the system should switch to a *fallback strategy* that works in all cases. Example fallback strategies are sketched below.

PET testing. In the event of a close race between X and Y , start Plaintext Equivalence Testing of randomly selected, unrevealed ballots against $\{X\}_{PK}$ and $\{Y\}_{PK}$, until we reach the required confidence for the winner.

Tally hiding. One can also fall back to computationally heavy methods e.g. MPC for only disclosing the winner, see e.g. [9,26,6,27]. Note that the revealed votes and reduced number of possible winners will make these methods more efficient than if used from the onset.

A possibility is to have the tellers perform a secret computation of the tally, and announce the winner(s), but not the numerical tally, on which to base a null hypothesis. This allows the RLT to be computed much more efficiently, and the secretly computed tallies can guide the appropriate strategy to adopt in the event of narrow margins.

6 Security Assumptions

In this section we briefly state the security guarantees and give some arguments for their validity. For the exposition below, we introduce the following three authorities besides the voters: The Tally Tellers TT holding the secret election key in a threshold manner, the Mixnet Tellers MT mixing the encrypted votes before doing the risk-limiting tally, and a random sampling authority RSA organising the random sampling of votes for the tally.

For simplicity let us also assume that the underlying voting scheme that we build on is mixnet-based, i.e., the main difference between the RLT version and the original version is that not all ciphertexts output from the mixnet are decrypted, but only a proportion of them.

In general, if RSA is acting honestly, or bound to do so e.g. via a verifiable proof based on a computational assumption, then the security reduces to that of the underlying scheme. For privacy, we normally have to trust that a threshold set of TT is not colluding and at least one server in MT is honest. For verifiability most schemes will not impose verifiable trust in TT or MT but might rely on computational assumptions and the RO-model or a CRS setup.

Verifiability. When random sampling procedure is corrupted, the adversary could possibly adjust the outcome in his favour. However, note that in this situation we can still achieve verifiability by having RSA committing to the sampling order before mixing, and assuming the last mix node is honest (or assuming one arbitrary mix node is honest and no threshold set of TT is corrupted). This will ensure that the final sampling is random.

Ballot-Privacy. Obviously, a necessary assumption for ballot-privacy is that a threshold set of TT are not colluding, and at least one mix node is honest. If the random sampling is also honest, we get strictly less information from the tally than in the original scheme, and we thus achieve better privacy in an information theoretic sense. When using standard ballot-privacy definitions on the scheme it should also be possible to reduce the ballot-privacy to that of the underlying scheme, the only subtlety being that tally functions differ in the two schemes.

It might seem that a corruption of the random sampling procedure should not influence ballot-privacy. However, there is one assumption to make: the random sampling should, in the computational view of the adversary, be uncorrelated with the cast votes. Having input from the voters to the random sampling could indeed make sense from a verifiability viewpoint, like in Demos [13], but should not depend on the vote choice unless this is computationally hidden.

Coercion-Resistance and Vote-Buying Resistance. As we have discussed above, the RLT protocol in general improves the coercion-resistance especially when candidates are expected to have a low vote count. It would be interesting to relate this to the coercion-resistance level δ in Kusters et al. [14]. On the other hand, the security against vote-buying is not increased in the same way since the voter here has an intent to obtain a receipt. The vote buyer could indeed follow the Italian attack method and the marked ballot would often appear.

The above is reminiscent of a distinguishing example between vote-buying and coercion resistance due to Rivest:⁶ the system chooses at random whether or not to provide the voter with a plaintext receipt. Such a system is, arguably, coercion resistant (the voter can claim to have received no receipt) but is vulnerable to vote buying (the voter might comply in the hope of getting the pay-off).

7 Risk-Limiting Verification

The idea of using risk-limiting techniques to improve coercion resistance can also be applied to verification of votes. Here, we apply the idea to Selene, allowing us to ensure that a proportion of the trackers remain unrevealed. In consequence, the coerced voter can always claim that her tracker was amongst those that remained shrouded. Some subtleties have to be handled in the case of an obnoxious coercer who demands the voter divulge their tracker; we describe those below. Indeed, these considerations require some modifications of the way Selene works.

7.1 Risk-Limiting Verification in Selene

A drawback of Selene, as noted in the original paper, is that when a coerced voter claims a fake tracker, the coercer (who is also a voter) could maintain that this is in fact his tracker. By construction, the coercer cannot prove this to the voter, but the voter is now in a difficult position: she knows that the claim might be true. Elaborations of the basic scheme are proposed, but they complicate things and render the verification less transparent: the final tally contains dummy votes that must be subtracted out to get the true result.

The RLT idea can be extended to avoiding revealing all the trackers in a run of a Selene election. The natural step is to apply the RLT mechanisms described above to reveal as many votes as necessary and then reveal the corresponding trackers. There would seem to be little point in revealing trackers for which the corresponding vote has not been revealed. There may, however, be some merit in revealing a subset of the trackers for which the votes have been revealed, as we discuss below.

Risk Limiting Verification (RLV), as applied to Selene, can ensure that not all trackers are revealed, thus allowing a coerced voter to simply claim that their tracker did not appear. There is still a problem, however, if we use Selene in its original form: the full set of trackers is published, so the coercer could require the voter to reveal her tracker anyway, and still claim that it is his.

We can fix this fairly easily: one purpose of revealing all the trackers in the setup phase is to demonstrate that they are all distinct, so the EA could publish a list of encrypted trackers for which the trustees run pairwise PETs to show that all the plaintexts differ. This would be computationally heavy and does not scale well, but is no worse than, e.g., JCJ [12]. Moreover, we can use some of the approaches to linearising the JCJ-style checks, for example by raising the tracker ciphertexts to the same, secret exponent and then verifiably decrypting.

⁶ private communication

We note that we get a form of partial random checking anyway when we reveal a random sample of the trackers: if all the revealed trackers in say a 90% random sample are distinct then we have very high confidence that they all are distinct. The drawback of this approach is that if the EA has cheated and included collisions then we will not discover this until rather late. Note, however, that we could reveal the trackers first, before revealing the votes. Now, if we find collisions, we can abort the election before any tally results have been revealed.

Still, one problem remains: another reason to publish the set of trackers is to allow voters to confirm that the α term sent to them is authentic: it opens the commitment to a valid tracker, i.e., a member the published set. If we do not publish the set of trackers then we need another mechanism for voters to confirm that their notified tracker is “valid.” We can achieve this by requiring that valid trackers are drawn from a negligible subset of the full space, e.g., numbers with say six digits. Now it is still intractable to produce fake α terms that will open a given commitment to a member of this set, but, by adjusting the number of digits we can ensure that the chance that a fake tracker will collide with the coercer’s is greatly reduced, so improving the plausible deniability.

If this reduced probability of tracker collision is deemed unacceptable, then we could allow the voter to request a fake tracker from the Notification Authority. This authority knows which valid trackers have not been assigned and so can provide an unassigned tracker to the coerced voter. This requires a level of trust in this entity, to keep tracker-related information secret but such trust is needed anyway.

There remains the question of whether all the voters should be notified of their tracker, even when their tracker has not been revealed on the BB. The immediate thought is not to notify unrevealed trackers, but this introduces possibilities of the authorities exploiting this: leading many voters to think that their tracker was not revealed and so denying them the possibility to verify their vote. It is not clear how we could verify that all the voters whose trackers are revealed are notified, so it seems wiser to notify each voter of their tracker.

8 Related Work

A number of papers [9,26,6,27] try to achieve tally hiding, either by only calculating the winner(s), or via multi-party computation and other cryptographic means. An idea closer to RLTs is that of Random Sample Voting (RSV) by Chaum [7]. A scheme that seeks to implement RSV in a fully verifiable fashion is Alethea [3]. RSV typically samples a small and predetermined number of voters, regardless of the margins. In contrast, RLTs adjust the sample size to obtain the desired level of confidence in the reported outcome.

The idea of Risk-Limiting Verification is somewhat analogous to Rivest’s ThreeBallot protocol [19]. Recall that, in ThreeBallot, each voter can verify a random 1/3 of her cast ballot. Thus, RLV gives “vote handles” to a fraction of voters, whereas in ThreeBallot each voter gets a handle to a fraction of her vote.

9 Conclusions

This paper presents two simple methods, RLT and RLV, for reducing the amount of information provided in the tally and verification stages. In consequence, we enhance the coercion-resistance by giving coerced voters plausible deniability, while achieving whatever confidence level in the outcome is required. An important future step will be to understand how well this method protects coerced voters in practice. It would be good also to understand better the trade-off between confidence in the outcome and plausible deniability levels.

There exist other methods that leak less information in the tally process, e.g., by using multi-party computation to only reveal the winner of the election. Such methods might be better suited to avoid strategic voting in runoff elections, and may provide somewhat better deniability. However, those methods require more elaborate and computationally expensive cryptography; arguably, our methods are more efficient and transparent.

The novel Risk-Limiting techniques introduced here should be of independent interest and have applications beyond the RLTs and RLVs described here.

Acknowledgements. WJ and PYAR acknowledge the support of the Luxembourg National Research Fund (FNR) and the National Centre for Research and Development (NCBiR Poland) under the INTER/PolLux project VoteVerif (POLLUX-IV/1/2016). PBR was supported by the EU Horizon 2020 research and innovation programme under grant agreement No. 779391 (FutureTPM).

References

1. Ben Adida, Olivier de Marneffe, Olivier Pereira, and Jean-Jacques Quisquater. Electing a University President using Open-Audit Voting: Analysis of real-world use of Helios. In *Proceedings of EVT/WOTE*, 2009.
2. Ben Adida and C. Andrew Neff. Ballot casting assurance. In *Proceedings of the USENIX/Accurate Electronic Voting Technology Workshop 2006 on Electronic Voting Technology Workshop*, EVT'06, pages 7–7, 2006.
3. David A. Basin, Sasa Radomirovic, and Lara Schmid. Alethea: A provably secure random sample voting protocol. In *31st IEEE Computer Security Foundations Symposium, CSF 2018*, pages 283–297, 2018.
4. J. Benaloh and D. Tuinstra. Receipt-free secret-ballot elections. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of Computing*, pages 544–553. ACM, 1994.
5. Josh Benaloh. Simple verifiable elections. In *Proceedings of the USENIX/Accurate Electronic Voting Technology Workshop 2006 on Electronic Voting Technology Workshop*, EVT'06, pages 5–5, 2006.
6. Sébastien Canard, David Pointcheval, Quentin Santos, and Jacques Traoré. Practical strategy-resistant privacy-preserving elections. In *Computer Security*, pages 331–349, Cham, 2018. Springer International Publishing.
7. David Chaum. Random-sample voting. http://rsvoting.org/whitepaper/white_paper.pdf.
8. David Chaum, Peter Y. A. Ryan, and Steve A. Schneider. A Practical Voter-Verifiable Election Scheme. In *Proceedings of ESORICS*, LNCS, volume 3679, pages 118–139. Springer-Verlag, 2005.

9. Josh Cohen. Improving privacy in cryptographic elections. Technical report, 1986.
10. Stéphanie Delaune, Steve Kremer, and Mark Ryan. Towards trustworthy elections. chapter Verifying Privacy-type Properties of Electronic Voting Protocols: A Taster, pages 289–309. Springer-Verlag, 2010.
11. S.N. Evans and P.B. Stark. Confidence bounds for the mean of a non-negative population, 2019. in prep.
12. Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-resistant electronic elections. In *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, pages 61–70. ACM, 2005.
13. Aggelos Kiayias, Thomas Zacharias, and Bingsheng Zhang. DEMOS-2: scalable E2E verifiable elections without random oracles. In *Proceedings of CCS*, pages 352–363, 2015.
14. Ralf Küsters, Tomasz Truderung, and Andreas Vogt. A game-based definition of coercion-resistance and its applications. In *Proceedings of IEEE Computer Security Foundations Symposium (CSF)*, pages 122–136, 2010.
15. M. Lindeman, P.B. Stark, and V. Yates. BRAVO: Ballot-polling risk-limiting audits to verify outcomes. *Proceedings of EVT/WOTE '11*, 2012.
16. Mark Lindeman and Philip B. Stark. A gentle introduction to risk-limiting audits. *IEEE Security and Privacy*, 10:42–49, 2012.
17. Silvio Micali. ALGORAND: the efficient and democratic ledger. *CoRR*, abs/1607.01341, 2016.
18. K. Ottoboni, P.B. Stark, M. Lindeman, and N. McBurnett. Risk-limiting audits by stratified union-intersection tests of elections (suite). In *Electronic Voting. E-Vote-ID 2018. Lecture Notes in Computer Science*, 2018.
19. Ronald L. Rivest. The ThreeBallot Voting System. <https://people.csail.mit.edu/rivest/Rivest-TheThreeBallotVotingSystem.pdf>.
20. Peter Y. A. Ryan, Peter B. Rønne, and Vincenzo Iovino. Selene: Voting with transparent verifiability and coercion-mitigation. In *Financial Cryptography and Data Security: Workshops*, pages 176–192, 2016.
21. Peter Y. A. Ryan, Steve A. Schneider, and Vanessa Teague. End-to-end verifiability in voting systems, from theory to practice. *IEEE Security & Privacy*, 13(3):59–62, 2015.
22. Peter Y A Ryan and Vanessa Teague. Pretty good democracy. In *WORKSHOP ON SECURITY PROTOCOLS*, 2009.
23. Steve Schneider and Abraham Sidiropoulos. CSP and anonymity. In *Proceedings of ESORICS*, pages 198–218, 1996.
24. P.B. Stark. Conservative statistical post-election audits. *Ann. Appl. Stat.*, 2:550–581, 2008.
25. Philip B. Stark and David A. Wagner. Evidence-based elections. *IEEE Security and Privacy*, 10:33–41, 2012.
26. Alan Szepieniec and Bart Preneel. New techniques for electronic voting. *USENIX J. of Election Technology and Systems (JETTS)*, 3(2):46–69, 2015.
27. Vanessa Teague, Kim Ramchen, and Lee Naish. Coercion-resistant tallying for STV voting. In *2008 USENIX/ACCURATE Electronic Voting Workshop, EVT 2008, Proceedings*, 2008.

VAULT: Verifiable Audits Using Limited Transparency

Josh Benaloh^{*1}, Philip B. Stark²[0000-0002-3771-9604], and Vanessa Teague³[0000-0003-2648-2565]

¹ Microsoft Research benaloh@microsoft.com

² University of California, Berkeley stark@stat.berkeley.edu

³ University of Melbourne vjteague@unimelb.edu.au

Abstract. Risk-limiting audits (RLAs) can provide strong evidence that reported election outcomes are correct, on the assumption that the paper trail of voter-verified ballots is trustworthy. Ballot-comparison RLAs involve comparing a human reading of voter intent from the paper ballot to the voting system’s electronic representation of voter intent for that ballot, the cast-vote record (CVR). Ballot-comparison RLAs first check that the full list of CVRs reproduces the reported results, then compare manual readings to CVRs for randomly selected ballots. For a ballot-comparison RLA to deserve public trust, the public must be able to validate those two steps. The easiest way to do that is to publish the entire list of CVRs. However, if every CVR is published, “Italian attacks” via pattern voting can be used to coerce voters or to facilitate selling votes.

Keywords: risk-limiting audit · homomorphic encryption · elections

1 Introduction

Over the last decade, *risk-limiting audits* (RLAs) [19,12] have gained traction as a method for verifying whether reported election outcomes⁴ accurately reflect the underlying paper trail. A recent report of the National Academies of Science, Engineering, and Medicine [16] advocates RLAs. They are performed routinely in Colorado, and are mandated by law now in Colorado, Rhode Island, Virginia, and Texas. There have been about 40 pilot audits in California, Colorado, Indiana, Michigan, New Jersey, Ohio, Rhode Island, Virginia and in Denmark.

RLAs involve manually examining random samples of paper ballots. If and when the sample provides adequately strong evidence that the reported outcome is correct, the audit stops; otherwise, it progresses to a full manual tally to set the record straight.

* Authors listed alphabetically.

⁴ *Outcome* means the political outcome—the candidate(s) or position(s) that won—not the exact vote counts.

However, auditing using rigorous statistical criteria is not enough to justify public confidence in election outcomes. An audit should not only allow insiders or approved auditors to check the results, it should also provide the public with enough information to verify that the audit was conducted properly and did not stop prematurely. At the same time, the public information should not compromise voter privacy. When RLAs are considered as a public verification process, their requirements closely resemble the *public verifiability* property of end-to-end verifiable elections.

The most efficient kinds of RLAs require a commitment to the interpretation of each ballot in advance of the audit. Traditionally, this commitment is made by producing a complete, plaintext statement of the contents of each ballot. Unfortunately, this can introduce a privacy problem for some election types. In this paper we show how a cryptographic commitment can be used as the basis of an RLA with essentially the same public verifiability as a traditional plaintext statement, but much better protection of individual vote privacy.

The methods are immediately useful in California, Colorado, Rhode Island (USA) and New South Wales (Australia).

We first describe how ballot-comparison risk-limiting audits work, then explain the privacy problem we are solving (Section 1.2) and the cryptographic tools we can use instead of plaintext commitments. Section 2 outlines the main advantages and shortcomings of VAULT compared with prior art. Section 3 then gives an overview of current audit law and practice in some example jurisdictions. The technical details of our approach are explained in Section 4, with some detailed examples in Section 5 and an informal argument for its main security properties in Section 6.

1.1 Ballot-comparison risk-limiting audits

Unlike traditional post-election audits, RLAs adjust the sample size to attain a desired level of confidence that electoral outcomes are correct, given what the audit finds as it progresses. There are many methods for conducting RLAs. The most efficient, measured by the number of ballots that need to be inspected when reported outcomes are correct, is a *ballot-comparison audit*. Ballot-comparison audits are possible only if the vote tabulation system creates an electronic interpretation of voter’s preferences for each ballot—a *cast vote record* (CVR)—in such a way that the corresponding paper ballot is uniquely identified and can be retrieved for manual inspection by auditors, so that their interpretation of the ballot can be compared to the CVR.

Existing protocols for ballot-comparison RLAs start with:

1. A *ballot manifest*, which describes in detail how the physical ballots are stored, so that ballots can be selected randomly and retrieved.
2. A *commitment* by the voting system to the full set of CVRs.⁵

⁵ Here, *commitment* is a term of art. It means that something about the CVRs must be published in such a way that observers can tell whether the CVRs that the audits

To conduct the audit, auditors first confirm that applying the social choice function to those CVRs yields the reported results, and that there are not more CVRs than ballots.⁶ (The social choice function is the rule for figuring out who won, such as plurality, multi-winner plurality, majority, IRV, or D’Hondt.) The audit proceeds by randomly selecting ballots and checking whether the corresponding CVRs match a human reading of the paper.

Ballot-comparison audits are like checking an itemized expense report using paper receipts. The first step is to check whether the itemized expenses add up to the total requested, and whether there is a receipt for every item. The second step is to spot-check the amounts of the reported expenses against the amounts listed on the receipts. Requiring the traveler to itemize expenses keeps the traveller from being able to fudge the numbers after the fact. Checking whether the itemized expenses add up to the requested reimbursement prevents a traveler from reporting every receipt accurately, but adding the expenses incorrectly.

Analogously, requiring a commitment to the CVRs before the audit starts keeps the system from simply generating CVRs that match whatever ballots the audit selects; and verifying that the collection of commitments imply the reported electoral outcomes ensures that if the commitments accurately reflect their corresponding ballots, the reported electoral outcomes must be correct.

A public auditing algorithm would therefore consist of:

1. Checking that the social choice function, applied to the CVRs, does indeed produce the announced election result.
2. Checking that the Risk Limiting Audit has been properly applied to the CVRs and paper ballots. This includes verifying that the random ballot selections are properly computed, checking that the correct paper ballot is retrieved according to the ballot manifest, applying the RLA risk computation to the ballot’s true value, and checking that the audit stops only when the RLA instructs it to (or falls back to a full manual recount).

In this work, we assume that VAULT takes as input a valid ballot-comparison RLA algorithm and concentrate only on the use of cryptographic rather than plaintext commitments. Important details such as how to verify that the ballots are properly selected at random, are out of scope.

1.2 Public evidence and voter privacy

Ballot-comparison RLAs provide strong public evidence that reported outcomes are correct if the commitment to the CVRs is public, if the ballot selection process is publicly verifiable, and the public can observe whether the selected ballots match the commitments about the corresponding CVRs.

check against the ballots are altered during the audit. One way to commit to the CVRs is simply to publish them all.

⁶ There are conservative methods for dealing with a mismatch between the number of CVRs and the number of ballots in the ballot manifest; see [2].

However, committing to the full set of CVRs by publishing them all may compromise the anonymity of the vote and enable an attacker to coerce voters through “pattern voting.” For instance, suppose an employer is running for mayor and wants to ensure getting the vote of all employees. The employer can select a lesser office on the ballot (e.g., “dog-catcher”) and threaten that each employee who wants to remain employed should cast a vote with the employer selected for mayor and the employee’s own name written in for dog-catcher. When the CVRs are released, the employer can check which employees complied with the demand. Even if write-in votes are not possible, the employer could select, for each employee, a unique pattern of votes on “downballot” contests and then check whether the patterns show up in the published CVRs. Complex voting systems such as Range Voting and Instant Runoff Voting (IRV) are susceptible even when there is only one race on the ballot.

1.3 Cryptographic commitments and homomorphic tallying

Here we show that cryptography provides an alternative way to commit publicly to the CVRs. This *cryptographic commitment* still lets the public check whether—on the assumption that the cryptographic commitments accurately reflect the votes on the underlying ballots—the reported results are correct, and also lets audit observers check (statistically) whether the commitments were accurate enough that the reported outcome is correct. Using appropriately designed cryptographic commitments protects voter privacy while still allowing the public to verify the audit.

Effective verification, of course, depends upon the protocol being sound. A verification mechanism may seem to be secure but actually leave gaps that make an election’s results unverifiable. For instance, part of the protocol may require the system to prove that the commitment for a CVR does not hide a negative vote, or more than one vote for a particular candidate. If the system could fake a proof that the committed value was valid, it could fake election results and evade detection with probability much higher than the RLA’s risk limit. This is not merely hypothetical: the protocol for the Scytl/SwissPost Internet voting system⁷ contains just such a flaw (Lewis, Pereira, and Teague 2019).

The remainder of this paper describes how techniques that for decades have been used in end-to-end verifiable (E2E-V) systems can be re-purposed to enable publicly verifiable ballot-comparison RLAs without revealing the contents of ballots other than those selected at random in the audit.

We use a cryptographic commitment scheme and denote by $c = E(m, r)$ the commitment to message m with randomness r . The commitment is *opened* when the committer produces (m, r) , thus allowing anyone to check that $E(m, r) = c$. The scheme must be both *hiding* and *binding*, meaning that the commitment does not reveal the message, and that it is infeasible to open a commitment in

⁷ The flaw also affects the iVote Internet voting system deployed in New South Wales, Australia.

two different ways, *i.e.* to find (m, r) and (m', r') s.t. $m \neq m'$ but $E(m', r') = E(m, r)$. Precise definitions can be found in any cryptography textbook [5,10].

E must also satisfy the homomorphic addition property:

$$E(m_1, r_1) \cdot E(m_2, r_2) = E(m_1 + m_2, r_1 \oplus r_2)$$

where \cdot and \oplus are easily-computed functions, usually modular multiplication and addition. For example, E might be an El Gamal encryption putting the message in the exponent. This is perfectly binding (because there is only one possible decryption of any ciphertext), but only computationally hiding (because an attacker who guesses the key can compute the plaintext). Alternatively, we could use Pedersen commitments [18], which are perfectly hiding but only computationally binding.

Using the homomorphic property, committed values can be combined by any observer to form a commitment to the sum of those values. The committed value can then be publicly opened, so anyone can verify that the claimed total is correct. With homomorphic tallying, individual votes are never decrypted or revealed.

Homomorphic tallying has been used in numerous cryptographic voting protocols to enable independent verification that a set of encrypted votes corresponds to an announced tally, without revealing the contents of individual ballots.

Some systems use perfectly hiding cryptographic commitments to achieve *everlasting privacy* [14,15,9,1], meaning that the published data does not expose information about the individual ballot even to an attacker with unlimited computational power. VAULT can be implemented with perfectly hiding cryptographic commitments and hence provide everlasting privacy for those ballots that are not audited.

The key contribution of this paper is the observation that homomorphic tallying also makes it possible to conduct ballot-comparison audits without revealing the contents of any ballots other than those selected at random in the audit, which is generally a small fraction of the ballots that were cast.

2 Related Work and VAULT’s advantages and limitations

2.1 SOBA

SOBA (*secrecy-preserving observable ballot-level audits* [3]) addresses the coercion problem by splitting ballots into their constituent votes and then creating a complex web of hash commitments that can be used to verify the required ballot properties without publishing full ballots. While SOBA is effective, it is complicated and unintuitive. No jurisdiction has used it to alleviate the real and practical problem of enabling public verification of ballot-comparison audits without putting voter privacy at risk.

SOBA and VAULT both rely heavily on cryptographic commitments, but in different ways. Public perception of the methods might be quite different, as a

result. In SOBA, people need to trust the cryptographic commitments to ensure that the plaintext votes for different contests really do correspond to “slicing” ballots into separate contests. The commitments prevent a cheating authority from reassembling sliced ballots any way they like—but the public can tally the plaintext votes themselves.

For VAULT, the public must rely on homomorphic encryption to check whether the commitments imply the reported outcomes. In SOBA it is less obvious that integrity relies on the cryptography, so SOBA may engender more public trust even though it relies just as essentially on cryptographic commitments.

Also, SOBA works by splitting up a ballot, which solves the problem for some social choice functions (such as Borda Count or Condorcet methods), and for US-style ballots with multiple questions on one ballot paper. It does not extend in an obvious way to Instant Runoff Voting (IRV), in which one vote may contain enough information for coercion, but it can not be divided into smaller parts while still allowing the social choice function to be computed.

Here we show how privacy-preserving ballot-comparison audits can be conducted far more simply and convincingly.

2.2 End-to-End Verifiable Elections

E2E-verifiability is generally achieved by publishing an encryption of all votes recorded in an election. An election is then *end-to-end (E2E) verifiable* if two properties are satisfied.

- Voters can confirm that their own votes have been correctly recorded.
- Voters and observers can confirm that all recorded votes have been correctly tallied.

The first of these properties is often referred to as *individual verifiability* while the second is typically known as *universal verifiability*. It is the universal verifiability property that is of interest for RLAs, because it closely matches the properties required of CVRs in a publicly observable ballot-comparison audit. However, there are some important differences.

The primary difference between VAULT and E2E-V is the level of protection that needs to be afforded to the raw data. In E2E-verifiability, releasing even a single raw ballot can directly compromise a voter’s privacy, because each voter in an E2E-verifiable election receives a receipt tied to the voter’s encrypted ballot. To prevent rogue individuals from decrypting the CVRs, decryption keys used for E2E-verifiability are typically shared amongst multiple independent parties in a way that some subset must cooperate to decrypt anything.

In contrast, for a ballot-comparison audit, the electoral process is assumed to have already done something to disassociate ballots from the identities of the voters who cast them. Thus the threat is lower: releasing an individual CVR does not immediately compromise privacy because ballots and CVRs are not linked to individual voters. Ballot-comparison audits require unsealing individual CVRs

as the audit progresses. It is inefficient to require a quorum of keyholders to convene and execute a decryption protocol every time an RLA selects additional ballots. It is therefore both desirable and sufficient for the encrypted CVRs to have a single decryption key—presumably held by election administrators.

2.3 Effectiveness of VAULT’s coercion-resistance

While it might appear as though the release of the complete set of votes on even a single ballot creates a privacy risk, the true risk comes from the release of the contents of *most or all* ballots. In order for effective coercion to take place, there needs to be a means by which a coercer can determine that a coerced voter *did not* vote as prescribed. However, if the contents of only a minority of ballots are revealed (at random), a coerced voter can simply assert that the voter’s ballot was not among those that were revealed.

While the new approach thwarts coercion, it is less effective against voluntary vote buying and vote selling. Even when the contents of only a small fraction of ballots are released, a lottery bounty might effectively purchase votes. A voter who might sell a vote for, say, \$10 might be just as willing to sell a vote for, say, a 1% chance of getting \$1,000. A vote buyer could therefore assign patterns to individual voters and pay a large bounty to any voter whose assigned pattern appears in a released CVR. A vote buyer’s potential payout could even be protected by tying the size of the bounty to the number of ballots whose contents are revealed.

3 Current audit law and practice

3.1 Colorado

Colorado counties that perform ballot-comparison audits upload ballot manifests and CVRs to state-provided, open-source software called RLATool. The Secretary of State publishes a cryptographic hash of the entire CVR file,⁸ but not individual plain text CVRs. The officials who audit the paper ballots manually and enter their reading of voter intent into RLATool generally do not have access to the CVRs, and do not calculate whether there is a discrepancy between the CVR and their interpretation: that is calculated by RLATool. Members of the public do not have access to the CVRs, before, during, or after the audit. After each round of the audit, the state generates a report that lists each ballot inspected and whether or not the CVR had a discrepancy, contest by contest.⁹ The public currently has no way to check whether the comparison was done correctly.

⁸ See, e.g., https://www.sos.state.co.us/pubs/elections/RLA/files/2018G/round_1/cvr_hash.csv (last visited 15 May 2019).

⁹ See, e.g., <https://www.sos.state.co.us/pubs/elections/auditCenter.html> (last visited 15 May 2019).

3.2 California

California AB2125, signed into law in 2018,¹⁰, authorizes pilot RLAs in 2020. Section 15366(b) defines *ballot-level comparison audits* (i.e., ballot-comparison audits):

- (1) The elections official uses an independent system to verify that the cast vote records created by the voting system or ballots created independent from the tally or ballot marking system yield the same election results as those reported by the voting system.
- (2) The elections official compares some or all of those cast vote records to a hand-to-eye, human interpretation of voter markings from the corresponding ballot marked by the voter or the voter verified paper audit trail, as defined by Section 19271.

Section 15367(b)(2)(G) requires the Secretary of State to establish regulations so that “the audit process is observable and verifiable by the public.” We interpret 15366(b)(1) in conjunction with 15367(b)(2)(G) to mean that the regulations must allow the public to verify that the CVRs used in ballot-comparison audits yield the reported results, and that the correct CVR is compared to each ballot selected for audit. That could be accomplished by publishing the entire set of CVRs in plain text—which could compromise voter privacy and facilitate vote-selling and coercion, as discussed above. Hence, it would be preferable to provide the public a way to ensure that the CVRs used in the audit yield the reported results without revealing every CVR. The approach we develop here solves the problem.

Constraints of Existing California Voting Systems

Conversations with California elections officials lead us to expect that the counties most likely to participate in the pilot have voting systems that produce CVRs that can be matched to ballots by relying on the order in which ballots are scanned, and that those counties are more likely to pilot ballot-comparison RLAs than ballot-polling RLAs.

Ballot imprinters were recently certified by the California Secretary of State; at least one jurisdiction likely to conduct audits under AB2125 plans to purchase imprinters. However, voting systems in most California counties cannot imprint

¹⁰ https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180AB2125, last visited 18 April 2019

identifiers or salts on ballots.¹¹ To our knowledge, no current voting system in California can print salts on ballots.

Thus, to comply with AB2125 and still protect voter privacy to the maximum extent possible, a method that does not rely on imprinting salts on ballots is needed.

3.3 Australia

Australian federal elections, and some state elections, use an automated scanning process to digitize paper ballots before counting, but currently no law requires any auditing of any paper records at all. As far as we know there is no public auditing in practice either.

4 Technical Details

For a simple plurality election, the process of proving that a set of encrypted ballots corresponds to an announced tally is identical to what is done for E2E-verifiability. However, by generalizing this approach, we can accommodate RLAs for a broader class of elections including instant-runoff voting.

Assume a ballot manifest, known to the electoral authority but not publicly released, providing a unique ID number for each ballot and attesting to its contents.

Suppose we have a set of asserted tallies $\mathcal{A} = \{A_1, A_2, \dots, A_J\}$ for the election. An assertion claims something about numbers that can be derived from the ballots, for example that a certain candidate’s tally has some particular value. These each contribute to some (perhaps several) null hypotheses $\mathcal{N} = \{N_1, N_2, \dots, N_I\}$ to be examined by RLA. Each paper ballot contributes some numerical value (most often 1, -1 or 0).

For example, if the election consists of a simple plurality election, then each assertion A_j might be the announced total of candidate j , and a_{ij} might be 1 if the ballot i is a vote for candidate j , zero otherwise. N_1 might be the hypothesis that a certain losing candidate actually got a higher tally (according to the paper ballots) than the announced winner. See Section 5 for detailed examples.

Note that some asserted tallies might be wrong though their dependent null hypotheses might still be demonstrably false—a small number of misrecorded votes, and hence some small errors in the announced tallies, don’t usually alter the election result.

¹¹ Experience in Colorado shows that printing sequential identifiers on ballots substantially increases the speed and accuracy of retrieving ballots. The Humboldt County Elections Transparency Project does imprint the ballots before re-scanning the ballots using an independent, unofficial system. <https://electionstransparencyproject.org/> While imprinting and rescanning could be the basis of a ballot-comparison *transitive* RLA of the kind conducted in pilots in California and Colorado, we do not anticipate that any California jurisdiction will attempt such an audit in 2020.

Let n be the total number of ballots. The audit proceeds as follows.

For each ballot b_i , for each assertion A_j , the EA posts a commitment to a_{ij} , which is a number representing b_i 's contribution to A_j . This commitment is denoted by C_{ij} .

$$C_{ij} = E(a_{ij}, r_{ij}) \text{ where } r_{ij} \text{ is randomly chosen.}$$

Then for each assertion A_j , the sum of the contributions of all b_i 's are computed (which is a public operation) and opened (which the authority has sufficient information to do, having produced the summands). That is,

$$C_j = \prod_{i=1}^n C_{ij}$$

and the EA publishes $\sum_{i=1}^n a_{ij}$ and $\bigoplus_{i=1}^n r_{ij}$. This opening can be immediately checked.

The audit consists of randomly selecting a paper ballot b_i , locating its electronic record, and then for each assertion A_j ($j = 1 \dots J$), opening the commitments C_{ij} , by publishing the pair (a_{ij}, r_{ij}) . This allows observers to check the commitment opening and verify that the committed values a_{ij} ($j = 1, \dots, J$) correctly describe ballot b_i 's contributions to each assertion.

Each committed value a_{ij} is expected to fall within some set S_j of valid entries, defined at the beginning of the election. For example, in a standard first-past-the post election the set of valid contributions to a candidate's tally is $\{0, 1\}$; in a Borda election it is $\{0, 1, \dots, n - 1\}$. The RLA defines the assumed sets of expected values for each assertion, and the EA proves that each committed value is within its corresponding set. *It is critically important that the proven ranges match the RLA's assumptions.* We will denote the proof that a commitment c contains a value in set S as $ZKP_S(c)$. Depending on the set, these could be instantiated as witness-indistinguishable disjunctive proofs (Cramer, Damgård, and Schoenmakers 1994), range proofs (Mao 1998), (Camenisch, Chaabouni, and others 2008), (Bünz et al. 2017), etc.

The first step is for the EA to define the assertions, RLA algorithm and corresponding sets of valid committed values. This is shown in Algorithm 1. The idea is that the assertions form the set of facts to be audited—it is up to the public to verify that their conjunction implies the announced election outcome. More precisely, the set \mathcal{N} of null hypotheses should obviously, when eliminated, imply that the announced election outcome is true, and the list of asserted tallies \mathcal{A} should (if true) imply that all the null hypotheses are false.

The commitment process is shown in Algorithm 2. There, $r_{ij} \leftarrow R$ means that r_{ij} is chosen randomly and uniformly from set R .

Verification is Algorithm 3. If there are some committed votes that do not have corresponding paper ballots, this can be dealt with using the phantom/zombie approach of [2].

Algorithm 1 Election outcome statement–EA

Input: Election outcome; social choice function; Risk Limiting Audit algorithm \mathcal{RLA} .

- 1: Announce the election outcome
 - 2: Define the set \mathcal{N} of null hypotheses to be examined by RLA.
 - 3: Define each assertion A_j for $j = 1..J$
 - 4: **for** $j=1..J$ **do**
 - 5: Define the set S_j of valid single-ballot contributions to A_j .
-

Algorithm 2 Commitment and opening algorithm–EA

Input: Ballot manifest; election outcome statement; commitment algorithm E with randomness range R ; set inclusion proof NIZKP ZKP .

- 1: **for** each ballot b_i **do** ▷ Make Commitments
 - 2: **for** each assertion A_j **do**
 - 3: $a_{ij} = b_i$'s contribution to assertion A_j
 - 4: $r_{ij} \leftarrow R$
 - 5: publish $C_{ij} = E(a_{ij}, r_{ij})$
 - 6: publish $ZKP_{S_j}(C_{ij})$
 - 7: **for** each assertion A_j **do**
 - 8: compute $C_j = \Pi_i C_{ij}$ ▷ Aggregate commitments
 - 9: publish $\Sigma_i a_{ij}$ and $\bigoplus_i r_{ij}$ ▷ Open the aggregate commitment
 - 10: When ballot b_i is audited ▷ Auditing
 - 11: Publish a_{ij}, r_{ij} for $j = 1, \dots, J$
 - 12: *Note: actually it is necessary to open the commitments only for those assertions for which the audit has not terminated.*
-

Algorithm 3 Commitment and opening verification algorithm—public

Input: EA’s election outcome statement; audited paper ballots; Risk Limiting Audit algorithm \mathcal{RLA} ; Commitment algorithm E ; set inclusion proof verification algorithm.

- 1: Check that the conjunction of $\{A_j\}$ over all j
 - 2: implies all the null hypotheses \mathcal{N} are false.
 - 3: Check that if all \mathcal{N} are false, this implies that the announced election outcome is true.
 - 4: If either of these checks fail, STOP and perform a full manual recount.
 - 5: **for** each assertion A_j and each ballot b_i **do**
 - 6: checking that S_j matches the assumed set in \mathcal{RLA} .
 - 7: verify $ZKP_{S_j}(C_{ij})$.
 - 8: **for** each assertion A_j **do** ▷ COMMITMENT VERIFICATION
 - 9:
 - 10: If the EA does not open C_j , STOP and conduct a full manual recount.
 - 11: Recompute C_j and check that $\Sigma_i a_{ij}, \bigoplus_i r_{ij}$ is a valid opening
 - 12: Check that $A_j = \Sigma_i a_{ij}$
 - 13: **for** each ballot b_i that is audited **do** ▷ AUDITING VERIFICATION
 - 14: **for** $j = 1, \dots, J$ **do**,
 - 15: verify that
 - 16: a_{ij}, r_{ij} is a valid opening of C_{ij} and
 - 17: a_{ij} accurately describes the paper ballot
 - 18: **if** the commitment opening is invalid or absent **then**
 - 19: **if** r_{ij} makes a valid opening of C_{ij} for some other value $a'_{ij} \in S_j$ **then**
 - 20: follow \mathcal{RLA} , with a'_{ij} as the apparent vote and the physical ballot as the true one.
 - 21: **else**
 - 22: follow \mathcal{RLA} , making the worst-case assumption about a_{ij} .
 - 23: **if** a_{ij} differs from the paper ballot **then**
 - 24: follow \mathcal{R} , with a_{ij} as the apparent vote and the paper ballot as the true one.
-

4.1 Defining the worst-case assumption

If the EA refuses (or is unable) to open a commitment, C_{ij} , or if a commitment opening doesn't verify, we must make the worst-case assumption about the message that was committed to. The worst-case assumption about a_{ij} is defined by the audit method and the valid set S_j . It might be different for each null hypothesis being tested.

Suppose for example that A_j declares a tally for some announced loser c_j , and that c_1 is the announced winner, in a single-winner plurality contest, with a tally announced by A_1 . Then $S_j = \{0, 1\}$. Suppose we have retrieved some particular ballot b_i and observed its contents, but the EA refuses (or is unable) to open the commitment C_{ij} . The commitment must have contributed to A_j 's homomorphic tally some value in the set S_j . Consider the implications for a particular RLA testing a particular null hypotheses N_k , which states that A_1 is a tally lower than or equal to A_j . The worst case assumption about a_{ij} is the maximum, over all values in S_j , of the discrepancy in favour of the announced winner compared with the true value on ballot b_i . This is one if b_i contains a vote for c_j , and zero otherwise. (The worst case is that a true vote for a loser was instead tallied as zero.) If the EA also refuses to open the commitment to a_{i1} , then a similar analysis shows that the worst case interpretation is another 1 if b_i shows a vote for the announced loser—if both of these happen, the RLA treats it as a two-vote overstatement.

The case in which the EA refuses to open the commitment might be ameliorated by using encryption (rather than other kinds of commitments) because then there is some set of authorities who hold the decryption keys, and may therefore open the commitment without having generated it. These authorities may still refuse to decrypt the message, however, so there still needs to be a way of incorporating this refusal into the audit.

For more expressive voting schemes such as Range Voting, if we let $A_j(b_i)$ be b_i 's numerical contribution to assertion A_j , then the worst-case assumption for the discrepancy is $d_{worst_k} = \operatorname{argmax}_{s \in S_j} \{s - A_j(b_i)\}$.

Note that the worst-case assumptions are chosen independently across different assertions. We never prove or check that the commitments about a single ballot are consistent—a cheating authority could have made various assertions about a ballot that are not consistent with any real ballot.

The above is sufficient data to conduct a Risk Limiting Audit, which must be parameterised s.t. the set of possible committed values corresponds to S_j for each assertion A_j .

4.2 Putting it together with an RLA

We now have all the ingredients necessary to conduct a Risk Limiting Audit of the announced outcome, by testing the null hypotheses associated with each assertion.

The basis for RLAs described by [19] is simply to test a hypothesis about the mean of a finite non-negative population—in our case, we are testing the

hypothesis that the discrepancies between the paper ballot data and the committed values are large enough to alter the outcome. The set membership proofs guarantee that each individual discrepancy is bounded by a known value (it might be negative, but it is bounded below). Hence the statistics of the RLA work exactly as they would do in a traditional open-CVR-based audit with the same parameters.

4.3 Locating ballots and keeping track of salts

There are several different ways of doing the bookkeeping necessary to implement the algorithm.

1. The random openings r_{ij} could be printed directly on the paper ballots – either in plaintext or encrypted.¹²
2. In the case of multiple commitments per ballot, the random openings could be generated from a cryptographic PRNG, for which the seed was printed directly on the paper ballot.
3. The random openings could be posted, encrypted, on the WBB.
4. The index i could be printed on ballot b_i .
5. There could be no printing on the ballots, but they could be stored in a way that made the index associated with each ballot obvious to an observer.

When the only option available is 5, it is important to ensure a publicly-verifiable correspondence between the ballot IDs and their paper ballots. This protects against substitution of ballots during the audit. Accidental errors of this kind have caused problems during audits (Ottoboni 2019)—deliberate substitution could render the audit meaningless. Printing either ballot IDs (Option 4) or random commitment openings (Options 1 and 2) conveniently prevents this substitution, assuming that observers can see that all the ballots have been printed in advance.

Whether the IDs or the random openings are printed on the ballots seems to matter for convenience but not for security: if only a few random values are used, printing them on the ballot obviates the need for secure storage elsewhere.

5 Specific examples

5.1 California: multiple-winner first-past-the-post

Consider an election with multiple winners elected by first-past-the-post. The process here is identical to that which is currently performed for E2E-verifiability. Each assertion A_j can simply be the tally of candidate c_j .

For ballot b_i ,

$$a_{ij} = \begin{cases} 1, & \text{if } b_i \text{ contains a vote for candidate } c_j, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

¹² This was suggested by Marc Rosen of Galois, Inc.

So commitment C_{ij} should be a commitment to 1 or 0, with a proof that the committed value is 1 or 0.

Then $A_j = \sum_i a_{ij}$, so assertion j can be checked by homomorphically summing the commitments and accepting the outcome if, for all $j = 1 \dots, J$, the opened value of commitment C_j matches the announced tally A_j .

This check includes the proof of commitment range.

The null hypotheses \mathcal{N} correspond to each case in which an announced loser's tally is higher than or equal to the winner's.

In the audit step, when a paper ballot has been retrieved, observers simply have to check whether a_{ij} has the right value as required above. If b_i is a vote for an announced loser but the commitment C_{ij} is not validly opened, the worst-case assumption is that a_{ij} is 0 if c_j is an announced loser, or 1 if c_j is an announced winner.

5.2 Instant runoff voting

Instant runoff voting (IRV) is used in numerous Commonwealth countries and some US state and local government elections. Each vote is a list of candidates in preference order. The social choice function first tests whether there is anyone with a strict majority of first-preference votes. If not, the candidate with the lowest tally is eliminated and their votes redistributed according to the next-listed preference on each ballot. This proceeds iteratively until one candidate has a strict majority.

To apply VAULT, the assertions \mathcal{A} could be a description of each elimination in sequence, but a much more efficient audit could be conducted by using a set of assertions derived using the techniques of (Blom, Stuckey, and Teague 2019). In this case, \mathcal{A} is a set of assertions about ballot preferences which, in conjunction, are sufficient to prove the election outcome (though not necessarily the exact elimination sequence that is claimed). For example, it would suffice to prove that one candidate received more first-preference votes than any other candidate received mentions (if it were true).

Using the notation of (Blom, Stuckey, and Teague 2019), define:

$$\begin{aligned} \tilde{f}(c) &= \text{the number of first preference votes for } c, \\ \tilde{t}_S(c) &= \text{the tally of candidate } c \text{ assuming the uneliminated candidates are those in set } S \end{aligned}$$

Note that $\tilde{f}(c)$ is the minimum tally c can possibly have, while $\tilde{t}_{\{c_1, c_2\}}(c_2)$ is the maximum tally that c_2 can possibly have in any election in which c_1 has not been eliminated. If $\tilde{f}(c_1) > \tilde{t}_{\{c_1, c_2\}}(c_2)$, then c_2 cannot possibly be eliminated before c_1 .

The algorithm of (Blom, Stuckey, and Teague 2019) can produce various kinds of assertions that suffice, together, to prove that the reported winner truly won, and could therefore be immediately used for the set \mathcal{A} .

To take a simple example, suppose that in some particular IRV election with $n + 1$ candidates, it happened to be the case that for all $j \neq n + 1$, $\tilde{f}(c_{n+1}) >$

$\tilde{t}_{\{c_{n+1}, c_j\}}(c_j)$. So define $a_j = \tilde{f}(c_{n+1}) - \tilde{t}_{\{c_{n+1}, c_j\}}(c_j)$ for $j = 1 \dots, n$. Then c_{n+1} won the election if, for all $j = 1 \dots, n$, $a_j > 0$.

Although this is not always true, it turns out to be true surprisingly often in real IRV elections, in which case it provides a simple and efficient test of the announced election outcome.

The audit can proceed by testing the set of n assertions $\mathcal{A} \equiv \{a_j > 0\}_{j=1}^n$. More specifically, for ballot b_i ,

$$a_{ij} = \begin{cases} 1, & \text{if } b_i \text{ has candidate } c_{n+1} \text{ as its first preference,} \\ -1, & \text{if } b_i \text{ has candidate } c_j \text{ preferred over } c_{n+1}, \\ 0 & \text{otherwise.} \end{cases}$$

So commitment C_{ij} should be a commitment to one of these values, with a proof that the committed value lies in the set $\{-1, 0, 1\}$.

Then $a_j = \sum_i a_{ij}$, so the public can check whether the CVRs satisfy the assertion $A_j \equiv \{a_j > 0\}$ by homomorphically summing the commitments and accepting the outcome if, for all $j = 1 \dots, J$, $a_j > 0$. This check includes the proof of proper range.

In the audit step, when a paper ballot has been retrieved, observers simply have to check whether a_{ij} has the right value as required above. If the commitment is not validly opened, the worst-case assumption is that a_{ij} is 1.

6 Overall risk-limit argument

Here we state our main security claim and sketch an argument to support it. The adversary controls the EA but not the verification algorithm. The security is based on the risk limit of a traditional RLA with plaintext CVRs—we assume correct functioning of the cryptographic aspects of that, including public verification that the random choices are correctly made and that the correct physical ballot is retrieved.

Recapping our setup:

- Let \mathcal{A} be a set of assertions which, in conjunction, suffice to prove the accuracy of the announced election outcome.
- For each assertion $A_j \in \mathcal{A}$, S_j is the set of possible contributions to a_j for any valid ballot. (Note, it will usually be a range of integer values, but this is not necessary.)
- For each commitment C_{ij} , the authority proves and the verifier checks that C_{ij} is a commitment to a value in S_j .

We want to argue that the overall probability of mistakenly accepting a wrong election outcome (as defined by the physical ballots) is the (negligible) probability of breaking the cryptography, plus the risk limit of the RLA. We *don't* need to prove consistency across different commitments for the one ballot.

Claim. Let VAULT be parameterised with an RLA with Risk Limit α for plaintext CVR commitments. Then the risk limit obtained by substituting VAULT

for the traditional RLA procedure is at most $\alpha + \epsilon$, where ϵ is the combined probability of the attacker undermining the soundness property of either the ZKPs or the commitments, *i.e.*

- being able to open a commitment in two different ways, or
- producing a set-inclusion proof that passes verification for a value that is out of range.

Proof. (Sketch)

If the election outcome is wrong, then at least one of the null hypotheses is true. Wlog call it hypothesis N_1 , and suppose it is negated by assertions A_1 and A_2 . Then we have a series of commitments C_{i1}, C_{i2} for $i = 1, \dots, n$ s.t. the homomorphically-added commitments

$$C_1 = \prod_{i=1}^n C_{i1} \text{ and } C_2 = \prod_{i=1}^n C_{i2}$$

can be opened as commitments to A_1 and A_2 (resp), and ZKPs ZKP_{i1}, ZKP_{i2} for $i = 1, \dots, n$ s.t. z_{i1} (resp z_{i2}) passes verification for the statement that C_{i1} (resp. C_{i2}) commits to a value in S_1 , (resp S_2) though in fact the claimed comparison between A_1 and A_2 is false according to the physical ballots.

If the authority can produce either a commitment opening to two different values, or a set-inclusion proof $ZKP_{S_i}(m_i)$ that passes verification though $m_i \notin S_1$, then cheating may succeed with probability greater than α .

We assume this happens with probability at most ϵ and, for the rest of this proof, assume that it has not happened.

Then the authority knows, for each C_{i1} , at most one tuple (m_{i1}, r_{i1}) that constitutes a valid opening (and likewise for C_{i2}). (Note that perfectly binding commitments, like El Gamal encryptions, get uniqueness automatically. *i.e.* there exists a unique valid opening, we don't have to assume that the authority knows only one.) Similarly, for the product commitments C_1 and C_2 , the authority knows at most one valid opening (M_1, R_1) (resp (M_2, R_2)).

We have taken a random selection \mathcal{I} of paper ballots (leaving aside for now the question of cheating on the predictability of those selections) and, for each of them, either had the corresponding commitment opened as (m_i, r_i) and checked whether it is a proper opening of C_{i1} , or had no commitment opened and made the worst-case assumption.

All commitments have been proven to come from some set, which has been checked to match the assumptions of the RLA. Some have been opened; others not, for which we made the worst-case assumption. Thus the process is equivalent to an RLA in which every ballot's contributing value was in S_1 , with the CVRs being equivalent to the openable values for everything in \mathcal{I} , and the worst-case values for everything else.

So apart from the ϵ probability of cryptographic failure, everything about the audit is identical to an RLA (whichever RLA is being conducted) with m_i as the apparent/claimed CVR. Thus the overall probability of accepting a wrong election result is ϵ plus the risk limit of the RLA.

7 Privacy guarantees

VAULT exposes the exact contents of those ballots that are audited. This still allows for some coercion, because a randomly selected fraction of voters can prove that their ballots were part of the tally. There is also, always, the possibility of a full manual recount, which exposes all individual ballots. Hence the privacy guarantees of VAULT are usually better than an RLA that publishes plaintext CVRs, but they are not always strictly better and not better for all voters. However, if we consider an attacker who observes the WBB but not the full manual recount, then VAULT does not reveal extra information about those ballots that are not audited.

We assume that the election authority is trusted for privacy, and that the identity of the voter is separated from the CVR before it is committed on the WBB.

Claim. Against an attacker who observes the WBB but not the (possible) full manual recount, and assuming that the election authority is trusted for privacy, VAULT does not reveal information about votes that were not audited except what can be derived from the election outcome statement (Algorithm 1). The guarantee depends on the form of commitment: it is perfect if perfectly-hiding commitments are used, or computational if computationally-hiding commitments are used.

8 Conclusions

Risk-limiting audits are an important tool to ensure election integrity and to provide trustworthy public evidence that reported outcomes are correct: that tabulation errors did not result in reporting the wrong winner(s). However, the most efficient approach to RLAs—ballot-comparison audits—are publicly verifiable only if three conditions hold:

- i. There is a public commitment to the full list of CVRs before the audit starts.
- ii. The public can verify that applying the appropriate social choice function to the committed list of CVRs yields the reported election results.
- iii. The public can verify how the contents of each CVR selected for audit compares to its corresponding paper ballots.

While these can be accomplished by publishing the entire list of CVRs as plain text, that would enable voter coercion. The only published approach to mitigate the risk of coercion while meeting (i)–(iii), SOBA [3], has never been used in a real election, possibly because of its complexity. We have shown that existing homomorphic tallying techniques used for end-to-end verifiability can make publicly verifiable, privacy-preserving ballot-comparison audits simpler, for instance, by publishing a complete list of homomorphically encrypted CVRs before the audit starts.

The minimal set of required cryptographic elements do not entail any change to voting systems, only post-processing the CVRs to create a set of cryptographic commitments for each ballot, and posting the results.

References

- [1] Arapinis, M., V. Cortier, S. Kremer, and M. Ryan. 2013. “Practical Everlasting Privacy.” In *International Conference on Principles of Security and Trust*, 21–40. Springer.
- [2] Bañuelos, J.H., and P.B. Stark. 2012. “Limiting Risk by Turning Manifest Phantoms into Evil Zombies.” arXiv.org. <http://arxiv.org/abs/1207.3413>.
- [3] Benaloh, J., D. Jones, E. Lazarus, M. Lindeman, and P.B. Stark. 2011. “SOBA: Secrecy-Preserving Observable Ballot-Level Audits.” In *Proceedings of the 2011 Electronic Voting Technology Workshop / Workshop on Trustworthy Elections (EVT/WOTE '11)*. USENIX. <http://statistics.berkeley.edu/~stark/Preprints/soba11.pdf>.
- [4] Blom, M., P.J. Stuckey, and V. Teague. 2019. “Risk-Limiting Audits for IRV Elections.” *arXiv Preprint arXiv:1903.08804*.
- [5] Boneh, D., and V. Shoup. 2016. “A Graduate Course in Applied Cryptography.” *Draft of a Book, Version 0.4*.
- [6] Bünz, B., J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. 2017. “Bulletproofs: Efficient Range Proofs for Confidential Transactions.” *IEEE SP*.
- [7] Camenisch, J., R. Chaabouni, and A. Shelat. 2008. “Efficient Protocols for Set Membership and Range Proofs.” In *International Conference on the Theory and Application of Cryptology and Information Security*, 234–252. Springer.
- [8] Cramer, R., I. Damgård, and B. Schoenmakers. 1994. “Proofs of Partial Knowledge and Simplified Design of Witness Hiding Protocols.” In *Annual International Cryptology Conference*, 174–187. Springer.
- [9] Demirel, D., J. Van De Graaf, and R. Araújo. 2012. “Improving Helios with Everlasting Privacy Towards the Public.” *EVT/WOTE 12*.
- [10] Katz, J., and Y. Lindell. 2014. *Introduction to Modern Cryptography*. CRC Press.
- [11] Lewis, S. J., O. Pereira, and V. Teague. 2019. “Ceci N’est Pas Une Preuve: The Use of Trapdoor Commitments in Bayer-Groth Proofs and the Implications for the Verifiability of the Scytl-SwissPost Internet Voting System.”
- [12] Lindeman, M., and P.B. Stark. 2012. “A Gentle Introduction to Risk-Limiting Audits.” *IEEE Security and Privacy* 10: 42–49.
- [13] Mao, W. 1998. “Guaranteed Correct Sharing of Integer Factorization with Off-Line Shareholders.” In *International Workshop on Public Key Cryptography*, 60–71. Springer.
- [14] Moran, T., and M. Naor. 2006. “Receipt-Free Universally-Verifiable Voting with Everlasting Privacy.” In *Annual International Cryptology Conference*, 373–92. Springer.

- [15] Moran, T., and M. Naor. 2010. “Split-Ballot Voting: Everlasting Privacy with Distributed Trust.” *ACM Transactions on Information and System Security (TISSEC)* 13 (2). ACM: 16.
- [16] National Academies of Sciences, Engineering, and Medicine. 2018. *Securing the Vote: Protecting American Democracy*. Washington, DC: The National Academies Press. <https://doi.org/10.17226/25120>.
- [17] Ottoboni, K. 2019. “Classical Nonparametric Hypothesis Tests with Applications in Social Good.” PhD thesis, University of California, Berkeley.
- [18] Pedersen, T.P. 1991. “Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing.” In *Annual International Cryptology Conference*, 129–140. Springer.
- [19] Stark, P.B. 2008. “Conservative Statistical Post-Election Audits.” *Ann. Appl. Stat.* 2: 550–581.
- [20] Stark, P.B. 2009. Risk-limiting post-election audits: P -values from common probability inequalities, *IEEE Transactions on Information Forensics and Security*, 4: 1005–1014.

The Swiss Voting Experience

The Swiss Postal Voting Process and its System and Security Analysis

Christian Killer and Burkhard Stiller

Communication Systems Group CSG, Department of Informatics IfI,
Universität Zürich UZH, Binzmühlestrasse 14, CH-8050 Zürich
{killer, stiller}@ifi.uzh.ch

Abstract. The Swiss postal voting system builds on trust in governmental authorities and external suppliers. The federal structure of Switzerland of cantons and municipalities leads to a distributed architecture. Detailed information on the current postal voting procedure are manifested as implicit knowledge within fragmented institutions and are not easily accessible. This work serves (i) as an overview of the Swiss remote postal voting system, (ii) a detailed insight into the process flow, and (iii) a respective risk assessment.

Keywords: Remote postal voting, risk assessment.

1 Introduction

Around the globe, government services are becoming increasingly digitized [1]. Naturally, these efforts include electoral processes. In Switzerland, the federal government defined strategies enabling digitization for public authorities and processes, including Electronic Voting (EV) [32,11]. Private companies collaborate with Swiss authorities to actively define standards across e-Government processes [35]. The Swiss EV typically refers to *Remote* EV (REV) carried out over the internet, which is also often referred to as Internet Voting (I-Voting) [19].

According to recent studies [32], 47% of Swiss citizens would be more likely to vote if EV were available, and almost 70% of Swiss citizens welcome an EV system [21]. Despite the positive sentiment surrounding EV, a current political position proposes a moratorium on EV in Switzerland [15]. According to their initiative [15], a REV system has to be “*at least as secure as the current remote postal voting (RPV) system*”. Thus, the key question is: what exactly does such a minimal level of security involve? Which security metrics and mechanisms are mandatory? In the general public perception, EV often provokes a fear of change, presuming the current RPV system to be mostly analog and tamper-proof. However, it can be argued that the current Swiss RPV system is already partially EV, since many steps already involve distributed electronic systems. Thus, defining and comparing the security properties of a REV also requires an analysis of the current RPV system in Switzerland.

Reducing cost and increasing the voter turnout by providing a convenient way to vote are important considerations for Swiss authorities [20]. By 1994,

all cantons accepted votes by postal mail. As of today, RPV is the dominant voting channel, used by approximately 90% of the voters in Switzerland [16]. Most eligible voters in Switzerland show trust in the authorities on the federal, cantonal, and municipal level to handle electoral processes and protect voter privacy [32]. The trust placed in authorities encompasses state-owned companies, which are important stakeholders in the current RPV system.

Due to the federal and decentralized structure of Switzerland, each canton and municipality autonomously manages their respective jurisdictional electoral procedures. Cantons and municipalities execute a degree of independence in decisions on how to handle certain parts of the voting process. Therefore, the current RPV system in Switzerland is neither universally documented or specified, nor homogeneous across entities.

This paper, therefore, summarizes major related work and terminology to formalize the Postal Voting Process Flow (PVPF) in Switzerland. The approach taken formalizes the PVPF in a step-based model, for which major assumptions made, such as trust, people involved, and technology applied, are made explicit, if known. The dedicated interpretation of social trust assumptions is discussed within Sec. 3, along with the risk analysis, weaknesses and strengths of a person-based RPV approach. Finally, the paper performs an overall risk assessment in Sec. 4, providing the basis for discussions of security-relevant comparisons to REV or I-Voting, while Sec. 5 draws main conclusions.

2 Legal Background and Related Work

Switzerland is organized as a decentralized system of municipal and cantonal entities, working together under the umbrella of the Federal Government. The federal structure is also mirrored in the legal framework (*cf.* Figure 1). At the root rests the Federal Constitution of the Swiss Confederation, wherein Art. 39 [6] forms the basis for the Federal Act on Political Rights (BPR) [8]. In turn, Art. 91 BPR [8] is the foundation for the Federal Decree for Political Rights (VPR) [9]. On a cantonal level, the VPR builds the foundation for the Cantonal Decrees (*e.g.*, for the canton of Aargau [7]). Every canton is an independent legal entity and defines its own constitution on the basis of the Federal Constitution.

The political system is under the authority of the cantons, *i.e.*, cantonal laws and ordinances regarding political rights define elements for these processes. Various aspects of those elements are relevant for the RPV system in Switzerland, and each canton has its own decrees regarding political rights. The federal structure is mirrored down to the municipal level: each municipality decides on certain processes, again, aligned to cantonal laws and decrees. For instance, keeping record of the electoral register is under the authority of municipalities, leading to different approaches.

A direct comparison of the Swiss RPV system to REV was performed in [29]. Other countries discuss the usage of RPV critically because the secrecy of the ballot cannot be fully ensured [29]. From a practical standpoint, thorough documentation is the easiest way to achieve verifiability for RPV. Supervisory bodies

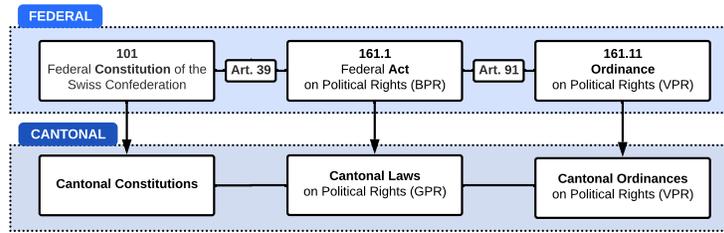


Fig. 1. Swiss Legal Framework

and authorities should check the documentation and verify it [29]. Also, trust is crucial for all voting methods. And the relationship between verifiability and trust is neither linear nor one-dimensional. Technical measures are not sufficient to create trust, sociopsychological aspects also have to be considered carefully. An extended literature review is provided in [18] with a focus on Switzerland, but also outlining the work done in Canada [22], Estonia [19], and Australia [30].

In order to analyze the RPV system with a focus on security aspects, the US National Institute for Standards Technology (NIST) serves as reference, outlining and standardizing terminology on the “Effort, Detection and Impact levels of Threat Events” [26,33,34]. Past work applied such principles to a RPV system used in the United States [25]. To consistently apply terminology, Table 1 defines the corresponding terminology used in Swiss legislation and their English translations.

Tab. 1. Official German Terminology with Corresponding English Translations

GERMAN	ENGLISH
Zwei-Weg Abstimmungskuvert	Two-Way Voting Envelope (VE)
Abstimmungsresultat	Voting Result (VR)
Erwahrung des Abstimmungsresultates	Legally valid determination of VR
Die Schweizerische Post	The Swiss Post (SP)
Stimmkuverts	Paper Ballot Envelope
Stimmrechtsausweis	Voting Signature Card (VSC)
Stimmregister	Electoral Register (ER)
Stimmzettel	Paper Ballot (PB)
Vertrauenswürdiger Dritter	Trusted Third Party (TTP)

3 Postal Voting Process Flow (PVPF) in Switzerland

This section details the illustrated Postal Voting Process Flow (PVPF) in-depth (*cf.* Figure 3), containing an end-to-end process as it is currently implemented in Switzerland. The detailed sub-steps are formalized and vary between cantons and municipalities. However, the general process adheres to the federal laws and

ordinances. The PVPF is divided into six main phases from *A* to *F*, each phase containing one or multiple sub-stages from 1 to *N* (*cf.* Figure 2). The ensuing subsections are structured according to the PVPF within Figure 3 and describe all the different steps in detail.



Fig. 2. Paper Voting Process Phases

3.1 Setup Phase

The Setup phase *A* contains four sub-stages 1-4 describing the production and assembly for dispatch of all necessary ballots and envelopes (*cf.* Figure 4). The two-way Voting Envelopes (VE), the Voting Signature Cards (VSC), the Paper Ballot Envelopes (PBE), and the Paper Ballots (PB) are the physical artifacts produced in the Setup phase. The secure execution of the Setup phase is crucial, since all following phases rely on the sound production and assembly of those artifacts. The main stakeholders of this phase are the municipal and cantonal authorities supervising the process. Due to cost, time, and capability constraints, Trusted Third Parties (TTP) support the authorities during the Setup phase as External Suppliers (ES).

Production of voting envelopes: In Step 1, the certified two-way VEs are produced by an ES. In the canton of Aargau, the municipality secretaries place a centralized buying order [28,17] for the two-way VEs at least a year in advance. After production, the VEs are distributed among the municipalities. In municipalities where Step 3 is outsourced, the VEs are directly delivered to the corresponding ES. The exact process steps are under municipal authority and can differ accordingly. Some cantons contract a single ES to handle the complete Setup phase *A*, mainly due to the special requirements of EV systems [10].

Production of Paper Ballots and Voting Signature Card: Step 2 consists of the production of the PBs and the VSCs. The printing of VSCs and PBs is predominantly commissioned to an ES. Each political layer in Switzerland (Federal, Cantonal and Municipal) commissions the PBs within their legal responsibility, *i.e.*, the production of federal referendum PBs are commissioned by the federal government, cantonal PBs are commissioned by the cantonal authorities, and municipal PBs are commissioned by the municipal authorities.

A VSC contains the name and address of the eligible voter, embedded within a template customized by the municipality. It is essential that the printed credentials are valid, since the assembled VE is delivered to the credentials printed on the VSC. The voter has to sign the VSC for the ballot to be valid. A substantial amount of ballots are not counted because many VSCs remain unsigned.

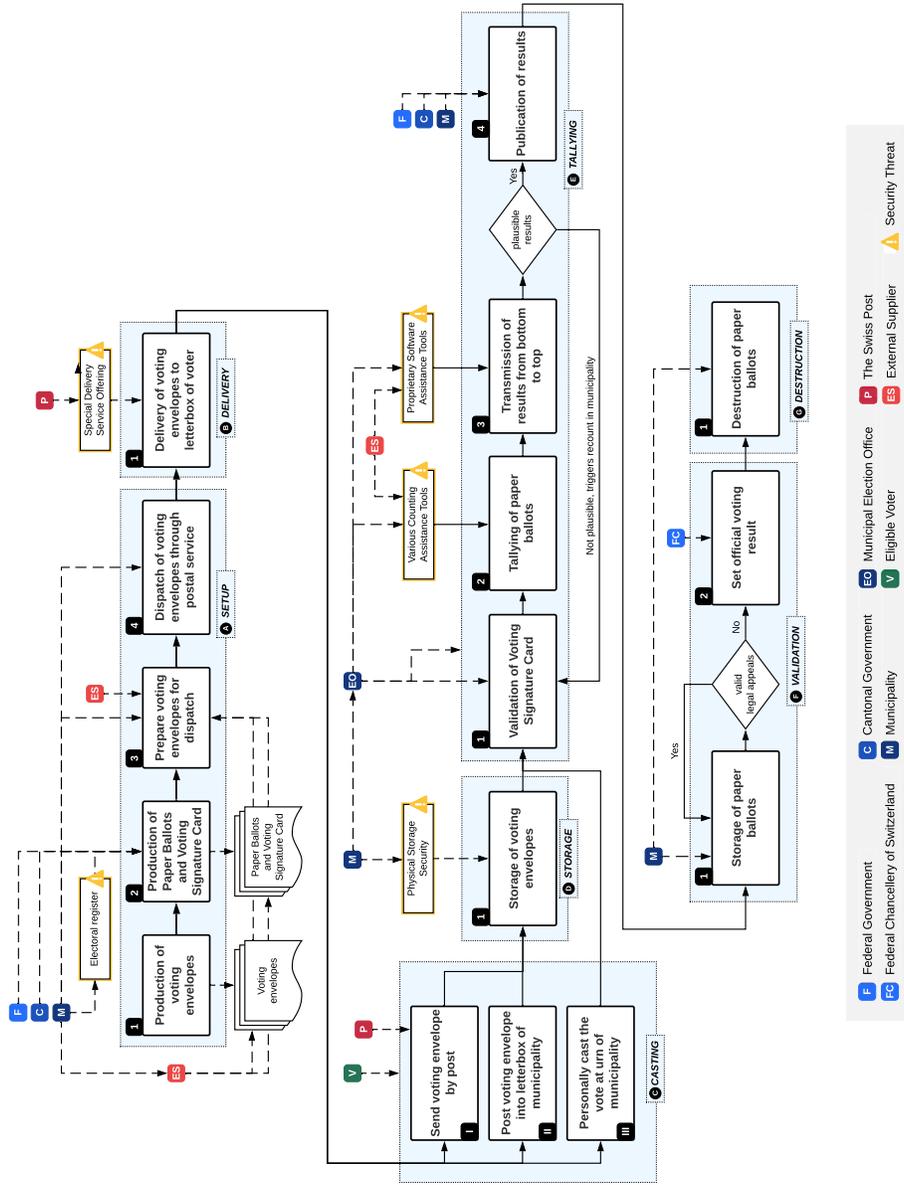


Fig. 3. Paper Voting Process Flow (PVPF)

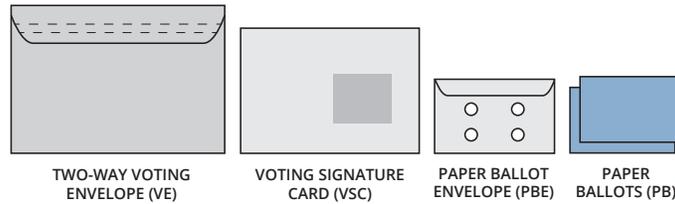


Fig. 4. Abstract representation of the necessary paper artifacts

The individual VSCs are printed according to an electoral register (ER). The ER is a centralized register containing all eligible voters. The ER is under the authority of each municipality. Neither the cantonal, nor the federal authorities have copies of the ER [28].

Since most municipalities contract an ES for the printing of the VSCs, a dataset containing the eligible voters needs to be transferred to the corresponding ES. Most municipalities export a file from the ER (*e.g.*, XLS, CSV) and send the snapshot to the ES directly via email [31,27]. The transmission of an unencrypted, unsigned dataset via standard email is critical, since the dataset could be tampered with (*e.g.*, the creation of fake identities, entries being removed), either after the export, during transmission, or when the export file reached the ES. Most ERs are administered by proprietary software systems provided by companies targeting Swiss public administrations. Some cantons also provide process checklists to municipalities. For instance, the election office (EO) of the Canton of Aargau provides such checklists [17]. These checklists state that the number of VSCs should be identical to the number of eligible voters present in the ER. If issues arise, an *in-depth* manual control should take place. Whether or not to adhere to these checklists is under the authority of the municipality. Also, since printing and assembly of the VE is mostly done by ESs, the ES should verify the integrity of the printed artifacts (*e.g.*, content and amount).

Prepare voting envelope for dispatch: Step 3 concerns the final assembly of the VEs. For each eligible voter (*i.e.*, each VSC), a VE containing the VSC, PBs, and the PBE is assembled (*cf.* Figure 4). The assembly is a monotonous task, often outsourced to ESs or social institutions and foundations [18]. Receiving an incomplete VE increases the possibility of the voter’s abstention. According to cantonal checklists [17], handing out new PBs is only allowed if the voter can *make the loss credible*. Then, the voter’s credentials should be recorded to check for attempted dual voting in the Tallying phase *E* [17].

Dispatch of voting envelopes through postal service: The final step involves the dispatch of the assembled VEs. In some municipalities, the VEs are directly dispatched by the ES commissioned with the assembly of the VEs. The Swiss Post (SP) offers a special service [14] for the dispatch and delivery of VEs.

3.2 Delivery

In Switzerland, the postal market was partially deregulated in 2009 [13]. Still, the SP maintains a monopoly on postal letters below 50 gr. Therefore, the SP is a crucial TTP, since the secure delivery to the municipality falls under the responsibility of the SP. When using the special service provided by the SP, VEs can be dispatched on a work day in the week prior to the specified delivery week [14]. Then, the SP guarantees the delivery of the VEs will take place during the specified delivery week [14].

3.3 Casting

Phase *C* outlines the three different options to cast a vote. The vast majority of ballots are not cast at the urn [16]. Statistics do not indicate whether VEs arrive through postal services (*I*) or were delivered to the letterbox by the voter (*II*).

I: The most popular way to cast the vote is to *send the VE by postal mail*. Some cantons pre-stamp the VSCs, which can then be used to return VEs free of charge [3]. For the voter, it is impossible to verify whether the ballot was successfully delivered to the municipal office. The SP offers the ability to track deliveries for an additional cost.

II: Thus, a favoured alternative is to *deliver the VE into the letterbox of the municipality*, which is then emptied by municipal officials and safely stored. According to [12], this option is still a favoured option by many voters.

III: The third option is to *personally cast the ballot at the urn*, which guarantees ballot secrecy. Casting ballots at the urn remains the most secure option to cast a vote, since the PBE (containing the ballot) is directly cast into the urn and separated from the VSC (containing the voter's credentials and signature).

3.4 Storage

Phase *D* deals with the storage of VEs that were delivered via postal service (*I*), or directly cast to the municipal letterbox (*II*). Often, an employee is tasked to fetch the postal mail addressed to the municipal office. During votes, the VEs are collected from the SP and municipal letterbox, and then carried to the safe storage location. Past incidents describe where municipal employees misused that trust [4]. The storage safety varies heavily, depending on the municipality. The Federal Act for Political Rights (BPR) [8] does not specify any security requirements. Additional considerations include the exact definition of an access control for the VE storage, (*e.g.*, Who should have physical access to the VEs?). Also, the definition of a process for incoming VEs can increase process security (*e.g.*, How many ballots arrived at which date and time? Who got the ballots from the letterbox or postal office and transported them where?).

Thus, stricter access control and a secured ballot arrival process can maximize the physical storage security. In practice, physical storage security is not prioritized, since the municipal infrastructure is often not sufficiently equipped [27,31].

3.5 Tallying

Phase *E* specifies the process of tallying. The main stakeholders of phase *E* are the municipality and the local EO. The tallying is not regulated on a federal level and is heterogeneous among cantons and municipalities [29].

Art. 14 No. 1 BPR [8] states that every polling station should create a report containing the total number of eligible voters, the total eligible voters living abroad, the total of blank, invalid, and valid ballot papers, and the number of votes in favour and against the proposal [8]. Thus, the BPR serves as a federal guideline, without specific requirements regarding the tallying process.

Approximately 10% [23] of the ballots cast are counted with the help of Electronic Counting (e-Counting) tools, provided by ESs. The parliamentary control of the administration investigated e-Counting and concluded that the federal requirements are neither functional, nor practical, and the control mechanisms of the federal government are not sufficient [23].

Tallying of all ballots: The local EO usually hires paid and elected helpers to assist with the manual counting. In large cities, thousands of helpers are engaged to count the paper ballots [24]. The EO defines the details of the tallying process. Some municipalities use e-Counting solutions or deploy high-precision scales to weigh PBs and derive the tally from averaging the weight of (sometimes pre-counted) batches.

Transmission of results from bottom to top: According to Art. 14 No. 2 BPR [8], the cantonal government is responsible for compiling provisional results from the entire canton and notifying the Federal Chancellery (FC) of the results, and publishing the same result in the Cantonal Gazette (or a special issue thereof) within 13 days of the polling day. As soon as the EO concludes tallying, the result is transmitted from the municipal EO to the cantonal EO, and from there to the FC. Some cantonal EOs deploy dedicated software systems to verify results using statistical methods. Also, most cantons make use of software provided by ESs to transmit the results. Thus, this phase also includes the use of web-based assistance tools [18].

Publication of results: The tallying phase is finalized with the publication of all results on the municipal, cantonal, and federal level. Generally, the FC publishes the collected results from the Cantonal Chancelleries in the Federal Gazette. Cantons publish the results and protocols in their Cantonal Gazettes. Each municipality publishes a final tally and tallying protocol with respect to the cantonal law. Mostly, the publishing process is performed by uploading documents to a public web-server and displaying print-outs outside the municipal offices.

Validation: Art. 15 BPR [8] defines the validation and publication of the results. The official results can only be declared when no valid appeals are in process at the Swiss Federal Court. After that, the official result is published by the FC in the Federal Gazette and can not be appealed anymore.

Storage of paper ballots: Before the results are ascertained, the counted PBs and VSCs have to be safely stored in the municipalities. It is important that the ballots remain unaltered because a recount could be triggered before the

official result is determined. Past cases have shown that premature destruction of the PBs and VSCs made a full re-count impossible [2].

Set official voting result: As soon as no more valid appeals are with the Federal Supreme Court, or as soon as a decision has been made on such an appeal, the legally valid voting result can be determined. According to Art. 15 No. 2 of the BPR [8], the validation decree shall be published in the Federal Gazette. The officiation by the Swiss Federal chancellery finalizes the Validation phase. Since a recount is no longer possible and the result is *untouchable*, the final phase can be started.

3.6 Destruction

The final phase, *G*, involves the destruction of the stored VSCs and PBs. According to Art. 14 No. 3 BPR [8], “*following validation of the result of the vote, the ballot papers shall be destroyed.*” In practise, the destruction is usually done by physically shredding all PBs and VSCs [27].

4 PVPF Risk Assessment

A risk signifies the level of impact on the operation of an information system’s task, given the potential impact of a threat and the likelihood of that threat occurring [5,34]. Therefore, a risk assessment (RA) serves as the identification and determination of the impact of vulnerabilities that an adversary can exploit. A threat covers any event with a potentially adverse impact on the assessed process [26]. With respect to the RPV in Switzerland, threat sources are groups or individuals who could feasibly attack the RPV system. Threat sources can stem from insiders or external adversaries. All Threat Events (TE) in the following RA are general in nature and require multiple co-conspiring hostile individuals or groups to achieve a large-scale effect. The effort for each threat defines the relative level of difficulty of performing a successful attack based on a threat [25]. Three relative levels of effort are defined:

- **Low** (–): An attack requires little / no resources or detailed knowledge of the system.
- **Moderate** (◦): An attack requires significant resources (or the ability to obtain these resources) or knowledge of the system. Inside attacks involving a small number of co-conspirators fall into this category.
- **High** (+): An attack requires excessive resources, in-depth knowledge of the system, or even access to the systems. It also requires specific tactics, techniques, and procedures [26]. Insider attacks involving a large number of co-conspirators fall into this category [25]

Detection describes the relative level of difficulty to notice whether a particular threat has been executed in an attack [25]. Thus, attacks are more severe when they remain undetected. Three estimated levels of likelihood of detection exist [25]:

- **Low** (–): An attack is unlikely to be detected without extraordinary resources.
- **Moderate** (◦): An attack may be detectable, but could require a large amount of resources and time. Such attacks are unlikely to be detected during the election.
- **High** (+): An attack would most likely be detected, given proper monitoring.

The impact on PVPF was analyzed according to [26] with the focus on confidentiality (*C*), integrity (*I*), and availability (*A*) as defined in [33]. Some TEs - all are shown in Tab. 2 - are interdependent or can be combined as depicted in Figure 5. The following RA serves as a major discussion of potential TEs that can lead to a loss of voter confidence. The mitigation of identified TEs concerns the actions of establishing trust and confidence in a system.

Tab. 2. Threat Events on the Swiss Postal Voting Process Flow (PVPF)

Phase	TE	Description	Effort	Detection	Impact
A	TE1	Delay production of physical artifacts	–	◦	A+
A	TE2	ER master records	+	◦	I+
A	TE3	ER data snapshot	◦	◦	I+
A	TE4	Forge physical artifacts	+	◦	I◦
A	TE5	Steal assembled VEs before dispatch	◦	+	I+
B	TE6	Re-route VEs	<i>unknown</i>	<i>unknown</i>	A+
B	TE7	Steal VEs from voter letterboxes	◦	+	A◦
C	TE8	Steal VEs from municipal letterbox	◦	◦	I◦
C	TE9	Re-route VEs	+	◦	C+
C	TE10	Cast stolen or forged VEs	◦	◦	I◦
D	TE11	Access stored VEs	–	–	I+
E	TE12	Manipulate tallying	+	◦	I+
E	TE13	Manipulate final tally	+	◦	I+
F	TE14	Initiate premature destruction	–	+	I+
G	<i>no major threat events identified</i>				

4.1 Risk Assessment Phase A

The Setup Phase *A* produces all the necessary artifacts for the secure execution of the whole PVPF.

TE1 describes the *malicious delay* of *A1* and *A2* in the PVPF. For instance, delaying the production can be achieved by targeting contracted ESs or directly attacking the municipal information systems.

TE2 describes the *tampering of the ER master records*. Often, the ER is provided and deployed by an ES. A targeted attack of an ES provider or municipal information systems with access to the ERs creates the ability to tamper with ER master records. The modification of master records can damage the integrity of ER data and the exported subset of eligible voters.

- TE3** describes the *tampering of ER snapshot data*. Instead of modifying the ER master records, the snapshot used to print the VSCs can be modified. When the snapshot is neither digitally signed nor encrypted, an adversary could modify the data before, during, or after transmission to the ES.
- TE4** describes the *forgery of physical artifacts* with (stolen) digital templates. If an adversary gains access to digital templates used to produce the physical artifacts, the adversary can forge VSCs and PBs. Additional information may be necessary to obtain (*e.g.*, weight and type of paper used). PBEs and VEs may also be forged, stolen, or even ordered from an ES. Since most municipalities do not perform a validation of incoming VSCs (by comparing the list of eligible voters with incoming VSCs), the attack can remain undetected. Practically executing TE4 requires a high effort and specific knowledge of the PVPF down to a municipal level.
- TE5** describes the *physical theft of the assembled VEs*. By stealing assembled VEs, the adversary can either destroy or cast ballots. The detection of this threat event relies on individual voters noticing that they did not receive their VEs, *i.e.*, the detection probability increases with every voter notifying municipal authorities.

The integrity and availability of ER is crucial for the Swiss RPV. By targeting ERs, substantial damage can be inflicted on data integrity, but also on trust in local authorities and can undermine voters' confidence. Requirements for EV systems can serve as a reference for process improvements [10].

4.2 Risk Assessment Phase B

The Delivery phase *B* is a black-box. The internal processes of the Swiss Post (SP) are not publicly available. When using the dedicated SP service to dispatch and deliver VEs, the VSC design must adhere to special layout rules to facilitate automatic batch processing [14]. The special layout of VEs could simplify identification of VEs, but requires an adversary to achieve partial control of the SP routing system. To achieve such control, a hostile individual or group can create an Advanced Persistent Threat (APT) within the SP and from there, *e.g.*, identify VEs according to specific attributes and re-route identified VEs, or attempt to delay the delivery deliberately.

- TE6** describes the *re-routing of VEs*. This TE requires adversarial access to internal SP systems and the capability to covertly manipulate the postal routing. A re-routing may require a co-conspiring postal employee because re-routing a large number of VEs could raise suspicion. Assuming a successful re-routing of VEs, the adversary is offered multiple options: Either to destroy the VEs, or open, modify, and re-cast them (*cf.* Figure 5).
- TE7** describes the *theft of VEs from voter letterboxes* before successful retrieval by the recipient voter. In contrast to TE5, TE7 requires the adversary to steal from individual letterboxes, not only at a single location. Similar to TE5, detection increases with every voter noticing the absence of VEs.

Phase *B* is characterized by the trust placed in one large entity, the SP. Thus, the effort and detection probability of TE6 can only be analyzed with additional information or access to internal SP systems, operations, and processes. Generally, however, an insider can achieve a low detection with moderate effort.

4.3 Risk Assessment Casting Phase C

TE8 describes the *theft of VEs from the municipal letterbox*. As shown in Figure 5, stolen VEs can either be destroyed or opened and modified.

TE9 describes the *re-routing of VEs (before delivery to the municipality)*. Similar to TE8, re-routing offers two different options: either the adversary can decide to destroy the VEs, or open, modify, and re-cast. Similar to TE6, a co-conspiring postal employee is crucial, since delivering a large amount of VEs to a different location than the authorities may alarm an honest employee.

TE10 describes the *casting of stolen or forged VEs*. An adversary can attempt to cast stolen and modified or forged artifacts to influence the voting result. The interdependence among TEs is visualized in Figure 5.

In official logs provided by the municipalities, there is no differentiation between channels *I* and *II*, both count as delivered by the SP. Even though 90% of votes are cast through *I* and *II*, keeping *III* remains crucial: Multiple channels strengthen confidence in results because it enables cross-channel comparison with statistical methods.

4.4 Risk Assessment Storage Phase D

TE11 describes TEs originating from *physical storage security*. Depending on the municipality, one or *N* employees have access to the cast VEs. The access to VEs offers similar options as presented in Figure 5. Since most municipalities do not log the amount of incoming VEs, the destruction of VEs can remain undetected.

The physical access to the ballots stored allows an adversary to either destroy VEs, modify them, or open VEs and break ballot secrecy. As past incidents show [4], access control to the stored VEs is again a question of trust.

4.5 Risk Assessment Tallying Phase E

TE12 describes the *risk of manipulation during tallying*. According to [23], over 10% of ballots cast in Switzerland are electronically counted. In 2014, sample checks identified errors in these counting mechanisms and concluded that e-Counting is neither more exact nor more secure than manual counting [23]. The manipulation of e-Counting requires an adversary to write targeted malware to influence the counting mechanism in his favor.

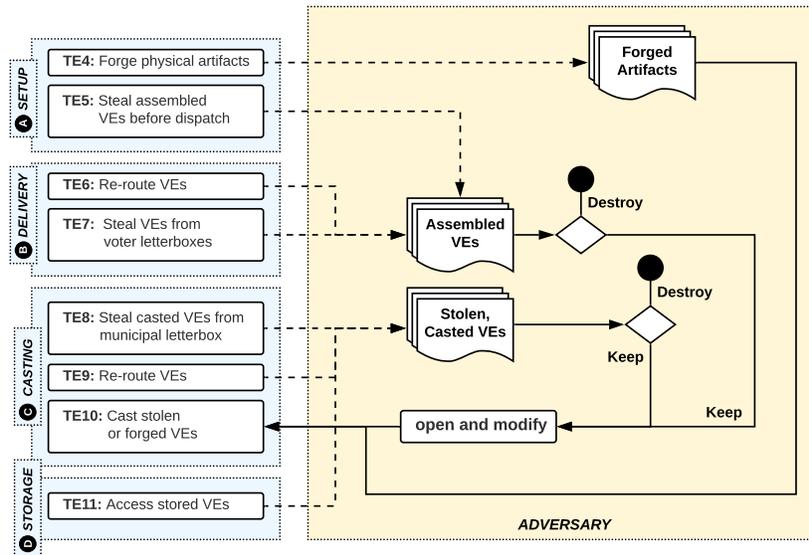


Fig. 5. Threat Event Interdependencies

TE13 describes the possible *manipulation of the final tally*. Some cantons use proprietary software to handle vote transmission from municipalities to the cantonal EO [18]. An adversary with access to these tools can tamper with the final tally. Since the manual tallying process produces logs published on a municipal level, large discrepancies can be detected by attentive observers. However, a sophisticated adversary can anticipate that and tamper with all digital traces to further obfuscate detection. Hence, the risk increases when PBs were exclusively counted electronically, without any redundancy from manual counting.

The tallying phase *E* builds on the integrity of each and every individual member of the municipal Election Offices (EO). The distribution of trust builds the cornerstone of the Swiss RPV system.

4.6 Risk Assessment Validation Phase F and Phase G

TE14 describes the *prematurely initiated destruction* of the PBs and VSCs. The destruction of PBs and VSCs before validation by the FC makes full recounts impossible, which already occurred in 2011 [2].

Since the validation finalizes and validates the official result within Phase *G*, a recount is no longer an option. Also, PBs are now irrelevant, since legal appeals are impossible at this point.

5 Conclusions

The Swiss postal voting system is highly successful, because substantial trust is placed in third parties, which includes a wide range of governmental authorities, state-owned companies, and various private companies and suppliers, and the individual voter. The current Remote Postal Voting (RPV) system is inherently built on external suppliers and trusted relationships among all parties involved. For a regular citizen, the current process is hard to decipher. Thus, this paper provides a coherent insight into the Postal Voting Process Flow (PVPF) and identifies its weaknesses as well as strengths with practical examples.

The main advantage of the current RPV system is its physical decentralization, which is undercut by using centralized information systems to administer or transfer crucial data (*e.g.*, Electoral Registers (ER) or Web-based assistance tools to transmit votes). Many aspects regarding the ER, assistance tools, the Voter Signature Card (VSC), or the physical storage of voting envelopes offer room for improvements from a security perspective.

The deployment of a Remote Electronic Voting (REV) system potentially decreases the necessary amount of trust placed in institutions and people, shifting trust to verifiable processes instead [29]. As this work showed, assessing the risks of the Swiss RPV system is reliant on the specific process across governmental entities. This work identified crucial Threat Events (TE) and showed that the system cannot serve as a suitable reference for electoral processes [15].

Furthermore, the Swiss federal structure leads to fragmented processes across jurisdictional barriers, from federal to cantonal, down to municipal authorities. The real-world deployment of the threat events identified requires a group of hostile individuals with specific knowledge. In small municipalities, authorities and citizens are intertwined and manipulations would either be not widely effective or detected rather swiftly. In large municipalities or large cities, processes are secured. Releasing an attack would require substantial effort from an attacker. Hence, an attack on the RPV is most likely to be successful in medium-sized municipalities, *e.g.*, where processes have not yet adapted to the larger size of the formerly smaller municipality.

Apparently, federal laws are not complete yet in guiding the deployment of secure e-Counting tools [23]. Thus, the compilation of an open and transparent list of all the electronic tools in use in the current PV flow can help to identify further threat events and enable the design of mitigation measures to handle risks better. Further, the Risk Assessment (RA) needs to be extended and ultimately applied to full real-world processes of cantons. In turn, TEs identified can be assessed in more detail and improvements can be provided to act as a comparison to EV systems.

Acknowledgements

The authors would like to thank Anina Sax, Annina Zimmerli, Dr. Christian Folini, Marco Sandmeier, and Dr. Benedikt van Spyk for their valuable input.

This paper was supported partially by (a) the University of Zurich UZH, Switzerland and (b) the European Union's Horizon 2020 Research and Innovation Program under Grant Agreement No. 830927, the Concordia project.

References

1. Andersen, K., Medaglia, R., Vatrapu, R., Zinner Henriksen, H., Gauld, R.: *The Forgotten Promise of E-Government Maturity: Assessing Responsiveness in the Digital Public Sector*. Government Information Quarterly, Vol. 28, Issue 4, October 2011, pp. 439–445
2. Berner Zeitung: *Stimmzettel fehlen, Nachzählung über Motorfahrzeugsteuern ist gefährdet*, August 2011, [Online] <http://pvpf.ch/bz-pb>, last visit July 9, 2019
3. Bühlmann, M.: Schweiz am Sonntag, Aargauer Zeitung: *Das Stimmcouvert per Post verschicken - ein Gratisangebot, das viele Aargauer ausschlagen*. [Online] <http://pvpf.ch/az-ve>, February 2016, last visit July 9, 2019
4. Bumbacher, B.: Neue Zürcher Zeitung: *Hauswart fälscht aus Frust Stimmzettel bei Gemeindewahl*, October 2005, [Online] <http://pvpf.ch/nzzfraud>, last visit July 9, 2019
5. Computer Security Division, Information Technology Laboratory: *Minimum Security Requirements for Federal Information and Information Systems*. (FIPS PUB 200), US Department of Commerce, NIST, March 2006
6. Der Regierungsrat des Kantons Aargau: *101 -Bundesverfassung der Schweizerischen Eidgenossenschaft vom 18. April 1999 (Stand am 23. September 2018)*. [Online] <http://pvpf.ch/bv>, last visit July 9, 2019
7. Der Regierungsrat des Kantons Aargau: *131.111 - Verordnung zum Gesetz über die politischen Rechte (VGPR) in Kraft seit 01.01.2013, Beschlussdatum: 30.05.2012*. [Online] <http://pvpf.ch/vgpr>, last visit July 9, 2019
8. Der Schweizerische Bundesrat: *161.1 Bundesgesetz über die politischen Rechte (BPR) (Stand am 1. November 2015)*. [Online] <http://pvpf.ch/bpr>, last visit July 9, 2019
9. Der Schweizerische Bundesrat: *161.11 Verordnung über die politischen Rechte (VPR) (Stand am 15. Januar 2014)*, [Online] <http://pvpf.ch/vpr>, last visit July 9, 2019
10. Die Schweizerische Bundeskanzlei: *Anforderungskatalog Druckereien für Vote électronique*, [Online] <http://pvpf.ch/bkreq>, last visit July 9, 2019
11. Die Schweizerische Bundeskanzlei: *Vote électronique*, [Online] <http://pvpf.ch/ve>, last visit July 9, 2019
12. Die Schweizerische Bundeskanzlei: *Änderung des Bundesgesetzes über die politischen Rechte*. Erläuternder Bericht zur Vernehmlassung, December 2018
13. Die Schweizerische Post AG: *Das Briefestromopol - Finanzierungspfeiler für die Grundversorgung*, [Online] <http://pvpf.ch/spmon>, last visit July 9, 2019
14. Die Schweizerische Post AG: *Factsheet, Briefe Wahl- und Abstimmungssendung*, [Online] <http://pvpf.ch/spfs>, last visit July 9, 2019
15. E-Voting-Moratorium: *Initiativtext*, [Online] <http://pvpf.ch/evmor>, last visit July 9, 2019
16. Grünhelfelder, P.: Neue Zürcher Zeitung: *Digitale Demokratie verlangt Pioniergeist*, September 2015, [Online] <http://pvpf.ch/nzzpio>, last visit July 9, 2019
17. Kantonales Wahlbüro Aargau: *Wahlen und Abstimmungen, Checkliste Allgemeine Arbeiten (Rahmenorganisation)*, November 2017

18. Killer, C., Stiller, B.: A Flow Analysis of Today's Swiss Postal Voting Process and a Respective Security Scrutiny. IfI Technical Report No. 2019-02, Department of Informatics IfI, University of Zurich, April 2019
19. Krimmer, R., Triessnig, S., Volkamer, M.: *The Development of Remote E-Voting Around the World: A Review of Roads and Directions*. First International Joint Conference on Electronic Voting and Identity (E-VOTE ID 2007). Bochum, Germany, October 2008 , pp. 1–15
20. Luechinger, S., Rosinger, M., Stutzer, A.: *The Impact of Postal Voting on Participation: Evidence for Switzerland*. Swiss Political Science Review , January 2007, pp. 167–202
21. Milic, T., McArdle, M., Serdült, U.: *Haltungen und Bedürfnisse der Schweizer Bevölkerung zu E-Voting = Attitudes of Swiss Citizens Towards the Generalisation of E-Voting*. Studienbericht, Aarau, September 2016
22. Pammet, H. Jon and Goodman, Nicole: *Consultation and Evaluation Practices in the Implementation in the Implementation of Internet Voting in Canada and Europe*. Research Study, November, 2013
23. Parlamentarische Verwaltungskontrolle (PVK): *Elektronische Auszählung von Stimmen (E-Counting) Bericht der PVK zuhanden der Geschäftsprüfungskommission des Nationalrates*, Februar 2017
24. Pauchard O. - Swissinfo: *Tausende beim Zählen der Wahlzettel*. [Online] <http://pvpf.ch/swi>, October 2003, last visit July 9, 2019
25. Regenscheid, A., Hastings, N.: *A Threat Analysis on VOCAVA Voting Systems*. Threat Analysis, US Department of Commerce, NIST, December 2008
26. Regenscheid, A., Hastings, N.: *Guide for Conducting Risk Assessments*. NIST Special Publication 800-30, US Department of Commerce, NIST, September 2012
27. Sandmeier, M.: Stadtschreiber und Leiter Stadtkanzlei Baden, February 22, 2019. Personal Conversation, Stadtkanzlei, Baden
28. Sax, A.: Leiterin Wahlen und Abstimmungen, February 20, 2019. Personal Conversation, Staatskanzlei, Generalsekretariat, Kanton Aargau, Regierungsgebäude, Aarau
29. Serdült, U., Dubuis, E., Glaser, A.: *Elektronischer versus brieflicher Stimmkanal im Vergleich. Überprüfbarkeit, Sicherheit und Qualität der Stimmabgabe*. Jusletter IT, September, 2017
30. Smith, R.: *Implications of Changes to Voting Channels in Australia*. Research Report comm. by the Electoral Regulation Research Network, December, 2018
31. van Spyk, B.: Vizestaatssekretär Kanton St. Gallen, February 27, 2019. Personal Conversation, Staatskanzlei, Recht und Legistik, Regierungsgebäude St. Gallen
32. Staatssekretariat für Wirtschaft SECO, Schweiz: *Nationale E-Government Studie 2019*. [Online] <http://pvpf.ch/egov19>, last visit July 9, 2019, March 2019
33. Stine, K., Kissel, R., Barker, W.C., Fahlsing, J., Gulick, J.: *Volume I: Guide for Mapping Types of Information and Information Systems to Security Categories*. NIST Special Publication 800-60, Vol. I, Rev. 1, US Department of Commerce, NIST, August 2008
34. Stine, K., Kissel, R., Barker, W.C., Lee, A., Fahlsing, J.: *Volume II: Appendices to Guide for Mapping Types of Information and Information Systems to Security Categories*. NIST Special Publication 800-60, Vol. II, Rev. 1, US Department of Commerce, NIST, August 2008
35. Verein eCH: *eCH-Standards*. [Online] <http://pvpf.ch/ech>, last visit July 9, 2019

How Do the Swiss Perceive Electronic Voting?

Social Insights from an Exploratory Qualitative Research

Emmanuel Fragnière^{1,2}, Sandra Grèzes¹ and Randolph Ramseyer¹

¹ University of Applied Sciences and Arts Western Switzerland (HES-SO), Sierre, Switzerland

² University of Bath, School of Management, Bath, UK

{emmanuel.fragniere, sandra.grezes, randolf.ramseyer}@hevs.ch

Abstract. Electronic voting is enjoying growing interest within the scientific community. However, the focus is on systems (algorithms, mathematical cryptographic models, user experience, reliability, traceability, security, etc.). Consequently, the purpose of this exploratory research on e-voting is not to address aspects that have already been well-studied by scientists, but rather to understand, through a qualitative research, bottlenecks and sociological obstacles. This understanding will help to explain the reasons that might prevent its adoption by Swiss citizens and also the dissemination of e-voting in the digital age. Based on 25 semi-directed interviews (in German, French and Italian) that we have analyzed, we are able to provide new insights that are more sociological than technological. These insights are essentially related to the social acceptance of e-voting. We observe in particular that the vote in Switzerland has an almost sacred dimension and that the trust that surrounds the voting “ritual” is of supreme importance.

Keywords: Democratic Values, Field Studies, Self-Determination, Social Acceptance, Perception, E-Voting Operations.

1 Introduction

1.1 Context

Electronic voting is enjoying growing interest within the scientific community. However, the focus is on systems (algorithms, mathematical cryptographic models, user experience, reliability, traceability, security, etc.). Estonia has become a well studied and known case since it has been systematically using it for many years. Switzerland could also become a reference in this field since it already represents a life-size laboratory because of its internationally recognized status as a semi-direct democracy. However, even if a law allows it, the Federal Chancellery remains extremely cautious about these developments. Last June 2019, the e-voting project has been postponed until the end of 2020. As a matter of fact arguments are often used to undermine the credibility of electronic voting, such as its unreliability (amplified hacking in the case of the last presidential elections in the United States), or people's tradition and attachment to the voting “physical process that are known and understood for a long time. The intergenerational digital divide is also an argument frequently used by some political parties. Thus, our

intention regarding this exploratory research on e-voting is not to address aspects that have already been well-studied by scientists, but rather to understand, through a qualitative field research, bottlenecks and sociological obstacles. This understanding will help to explain the reasons that prevent this type of development, as well as its adoption by citizens, in Switzerland. While postal voting has not caused any such resistance, and its use is well accepted and widespread, this is not the case with electronic voting.

1.2 Research Purpose and Contribution

The research project's purpose consists in providing the Swiss parliament and government with sociological rather than technical elements to develop a relevant public policy on electronic voting in Switzerland. Indeed, for the time being, within the framework of the Swiss e-voting project, the Confederation has only surrounded itself with experts (academics, the Confederation and the Cantons), with the aim to implement a fully functioning e-voting system. However, it is indispensable to integrate the opinions, perceptions and the rationale of the Swiss population, as the final users of e-voting is the population who is entitled to vote. Hence, this is also a matter of social acceptance of the new voting system by the population. Our project therefore has the advantage of questioning the population in depth (practically this has been done through 25 semi-structured interviews administered in the three national languages, respectively, German, French and Italian, during the period from January to June 2019). Moreover, it is important to take into account the cultural and socio-demographic characteristics of the different profiles interviewed, to consider all the stages of the different operational processes leading to the vote (ballot box, postal voting, electronic voting), to also integrate the major changes in all sectors of society caused by the global phenomenon of digitalization. Through the analysis of the interview transcripts, we have produced a synthesis of the main findings of this field study. This will allow us to develop, in a next step of this research, a theoretical model explaining the population's perceptions about the development of electronic voting in Switzerland. In a third step, everything will be ready to set up a national quantitative survey to validate the assumptions of our theoretical model in order to make statistical inferences at the national level.

From a scientific point of view, our study would be, to our knowledge, the first of its kind corresponding to a qualitative research taking into consideration the different Swiss cultures (French-speaking, German-speaking and Italian-speaking) that would allow the generation of "meanings" explaining the public's perceptions about the development of electronic voting. Indeed, the scientific literature on this subject is rather sparse compared to that with a more technical orientation. Even if we focus on a small country, Switzerland nevertheless represents a highly relevant and credible "democratic laboratory" on a global scale. It is indeed the country where the most votes are cast in the world and since 2000 more than 200 electronic voting trials have already been carried out (<https://www.admin.ch/gov/fr/accueil/documentation/dossiers/E-Voting.html>).

1.3 An Approach Anchored on the Notion of Voting Operations

The usual way to form a political opinion in Switzerland is to read official documentation and follow traditional media such as television or the press and recently the Internet. We are also largely influenced by poster campaigns and all-household distributions. However, new media as well as digitized democratic processes are completely changing the situation. Switzerland is the country with the highest number of popular votes in the world. The traditional process consists of a preparation period that can be quite long (about 6 months) in which parties and the media contribute to the formation of voters' opinions. They send their ballot papers by post in advance or go to a voting room, where booths will be made available during the weekend of the vote, as well as voting materials and the ballot box to cast their ballot. Over the years, there has been a decline in participation in these elections, especially among the younger generation. At the same time, our economy is currently undergoing very intense digitization. We are increasingly talking about the "ubiquitous" nature of the economy, which changes the roles assigned to each person and redefines what an expert, a provider or an agent is. In terms of information and opinion forming, roles are also changing with a transformation of the role of the journalist, editorialist and expert. The scandals in the last US presidential elections and the Brexit vote highlighted the fragility of our democratic processes. We intend to address the topic of electronic voting from the perspective of a business process (i.e. voting operations), while focusing on the human aspects. In Switzerland, at the present time, two voting options are available: either the voter goes to the polling station or the voter votes by mail. A third possibility, electronic voting, has recently been accepted in Switzerland, after years of testing. Due to recent developments, the e-voting process is on hold [1]. In the coming years, it could (or not) become a fully-fledged voting option on a par with the two options already in place. Electronic voting is based, like ballot box voting and postal voting, on an operational process, except that most of the steps are dematerialized, since they are digitized. This is called a digital process.

1.4 Organization of the Text

This paper is organized as follows, In Section 2, we present a brief literature review about e-voting. In Section 3, we describe the methodology that has been employed in this research. In section 4, we present an overview of the main elements of electronic voting perception. In Section 5, we present an overview of the main elements of electronic voting security. In Section 6, we present an overview of the main elements of electronic voting operations. In Section 7, we provide a discussion about the notion of trust and confidence in the voting process. Finally, we conclude and provide directions for further research.

2 Literature review

2.1 The digitalization of service operations

Electronic voting refers to any process that benefits from the use of electronic technology by electoral authorities for the conduct of elections [2]. As part of our research, it is important to place electronic voting in a broader field of investigation that generally concerns the digitalization of society. Digitization can be defined as the integration of multiple digital technologies into all aspects of daily life that can be digitized by the conversion of analogue information into digital form so that information can be processed, stored and transmitted through circuits, equipment and networks digital [3]. Thus, referring to the concept of digitization (or digitalization), rather than talking about digital processes, implies that it is an emerging transformation process, in progress, still in the development phase rather than a completed and clearly defined process [4]. These authors believe that that digitization corresponds to the characteristic of the information society, as defined by [5], i.e. it is not simply something that is imposed on individuals and organizations, but something that individuals and organizations "do" and produce themselves through daily practice and social interaction. Therefore, if we consider e-voting as a phenomenon of digitalization of our democratic society, we must also analyze it in its integration into the life of voters and in the context of the social interactions in which it is integrated. We live in a digital age, because digital technologies are used today in almost every aspect of life [6] and they play a key role in shaping and regulating societies, communities, organizations and individuals [7].

2.2 Research on the Topic of E-Voting

Let us now return to scientific research that focuses more specifically on the field of electronic voting. We note that the production of scientific articles on this subject over the past decade is very abundant [8], with pioneering scientific works already published at the beginning of the new millennium [9]. One country, Estonia, is a precursor in this field, having made it possible to use electronic voting in its elections since 2005 [10] [11]. However, it should be noted that Estonia has developed strongly from the point of view of digitalization in all sectors of society. This Estonian practice attracts the attention of all regions of the world. This is particularly the case in Europe. The Council of Europe has also issued a specific recommendation on this subject. The Recommendation "Rec(2004)11" on legal, operational and technical standards for electronic voting is a unique reference source in this field. Europe, through the Organization for Security and Cooperation in Europe (OSCE), has sent experts to Switzerland to observe the various test phases that have been planned since 2005. Switzerland can thus be considered as a kind of laboratory for the process of digitizing votes, which has made it possible to develop original scientific research [12] [13].

Most of the scientific contributions in the field of electronic voting deal mainly with technical aspects. There are many articles on the following themes: design and evaluation of electronic voting systems, identification and authentication of voters, reliability, security and safety issues, end-to-end traceability, etc. A recent article [14] provides a

comprehensive literature review on all these technical and usability aspects. Research on e-voting also incorporates the "Blockchain" to address supposed vulnerabilities inherent in most existing systems [15]. Some researchers even argue that traditional paper-based voting is subject to the same security problems. However, since we have been using paper for a long time, people are no longer even aware of security problems related to paper-based voting [16].

2.3 Taking into Account the Human Factor in E-Voting

What is not much studied in academic research, however, is everything that touches on the human aspects of electronic voting, and in particular the public's perception of this new possibility of voting. However, some rare studies of this type exist. This is the case with a survey conducted in Australia that shows that there is a correlation between perceived ease of use and perceived utility of e-voting technologies to determine their acceptance and use [17]. A Malaysian study on a campus shows that when students use e-voting to express themselves on university activities, there is a need to have confidence in the electronic system to ensure a real commitment to voting [18]. As part of this research project, we employ qualitative research to identify and explain the perceptions of e-voting by Swiss citizens about the widespread adoption of electronic voting. To our knowledge, no scientific study has yet been carried out on this subject specifically.

3 Methodology

3.1 Methods

The aim of this study is first of all, through documentary analysis, to better understand the innovations linked to the mode of democratic process does not lead to digitalization within our societies. It is therefore a question of taking stock of the e-voting initiatives carried out at the global level. Through this documentary research, it was also necessary to understand the situation in Switzerland (legal bases, parliamentary debates, motions, party politics, test phases carried out in the various cantons, etc.). Secondly, we have conducted field research based on an ethnographic approach. Ethnography represents the descriptive study of the activities of a specific human group. More specifically, we use ethnomethodology, which is not based on an a priori theoretical framework. Ethnomethodology makes it possible to identify the latent needs of the target population, to detect social trends for the design and improvement of given public service processes and finally to write scenarios to highlight intangible elements that bring added value to users. In practice, we have launched a field research based on 25 semi-structured interviews administered in German, French and Italian to directly question citizens about their perceptions regarding electronic voting. This research approach is therefore essentially based on the notion of constructivism (or interpretivism). Its main objective is to understand how and why electronic voting, as a new method of voting (on a par

with voting by depositing it in the ballot box and voting by post), is at the root of skepticism and opportunities. We therefore believe that this inductive approach is the most appropriate for our research. It is indeed well adapted to the understanding of the perception we have of the environment under study. The data collected through semi-structured interviews have been analyzed on a content analysis basis (with the help of RQDA and NVivo) according to the codes or code categories used for the analysis of the transcripts. On the basis of the synthesis of the results, we will develop in a subsequent research a new theoretical model to explain the public's major concerns about the adoption of electronic voting. In a third phase of research the generated model will be validated through a quantitative survey.

3.2 Purposeful sampling strategy and interview guide

For this qualitative research, we used the purposeful sampling technique [19]. The aim is not to be representative of the population studied in order to draw statistical inferences, but rather to "go around the issue". Indeed, a qualitative field research has an exploratory purpose and not the validation of research hypotheses. For inductive and exploratory research, qualitative methods are most suitable, as they can lead to hypothesis building and explanations. Qualitative research also delivers better understanding of motivations, values and attitudes on a given context. Data collected from a small number of carefully selected samples on relevant issues can be sufficient in this case, as it demands a limited number of observations to explain different aspects of the problem area. Low numbers are justified to do an in-depth study [20].

Nevertheless, in our case, we have tried to take into account Swiss cultural diversity (14 interviews in German, 8 in French and 3 in Italian). To have broad social insights, we chose people from different cantons. Moreover, to understand the perceptions of different generations, we chose people covering an age spectrum from 29 to 75 years. Apart from that, the other socioeconomic parameters are not representative of the Swiss population since the sampling strategy was purposeful. All of the interviewees are highly educated [21], most of them have a University diploma and they appear to vote on a regular basis. In our sample, there are solely a few people who are Swiss living abroad and who have already experience with e-voting, to enlarge our insights concerning e-voting. The main topics, guided by authors' discussions (translated into 14 questions) addressed during our semi-structured interviews were: e-voting "customer journey/mode of operation", personal views about electronic voting, trust and security aspects, meaning of e-voting in relation to democracy, habits related to digitization, e-voting for Swiss abroad/disabled people, obstacles and barriers to e-voting. Choices of the different dimensions retained in our interview guide have been based on discussions taking place before the fieldwork and on the literature review of section 2.

4 Findings Regarding E-Voting Perception

4.1 Different Attitudes Towards E-Voting

The qualitative analysis of the interviews revealed different attitudes towards electronic voting. In the following, we tried to identify some patterns of attitudes towards e-voting. Type 1 is a strong supporter of the e-voting system. He or she is accustomed to the use of digital devices for many daily activities and e-administration, such as e-banking. He or she rather or strongly trusts the political system and/or e-voting. One respondent e.g. has been eager for years to vote electronically: “I would love it. It makes everything more flexible and easier for me” (German-speaker from Valais, female). Respondents of this type of attitude relativize the risks associated to e-voting, as an interviewee underlines “If someone wants to manipulate, he or she will succeed in doing it, no matter the voting type” (German-speaker from Bern, male).

On the other side of the spectrum, type 4 is totally against e-voting and would never vote like that. However, there was only one person who has this attitude among those questioned. This person argues that “a voice cannot be reduced to an action on Internet” (German-speaker from Valais, male). Another argument of this person is that via e-voting “not only Swiss can manipulate the vote, but theoretically the whole world” (German-speaker from Valais, male).

In between this spectrum of supporter and opponent of e-voting, there are different nuances of attitudes towards e-voting. Such a profile (type 3) is rather skeptical towards e-voting. However, this type acknowledges the convenience aspects of e-voting and trusts the system for better or for worse by being aware of the various risks associated with this system. One respondent says: “I have always been very skeptical, because I asked myself: is it really enough secure?” (German-speaker from Solothurn, male). Then, we could identify another type of attitude (type 2) that is characterized by a limited interest in e-voting and do not have a strong feeling of support of or opposition of e-voting systems. The reasons for this attitude can be different. Some of them say that “the introduction of e-voting corresponds to modern times” (German-speaker from Schaffhausen, female / German-speaking from Valais, female). People with this attitude think that e-voting is a logical consequence of current technological developments, which also influence the way people vote. Another reason can be that “they have not yet had the time to deal with that topic” (German-speaker from Basel, female / German-speaker from Fribourg, male). One respondent mentioned that he has not yet become accustomed to this “e-government logic” (German-speaking from Fribourg, male). Yet, these latter persons are also aware of the risks associated with e-voting.

Several interviewees mentioned that they are more or less obligated to trust the system, as they have not the time, nor the interest or competence to understand in detail the technological process of e-voting. Therefore, interviewees trust the system in general, but one of the respondents says that a lack of security would indeed be a no-go. Another interviewee says “if there is somewhere a chance, to forge the vote, then it is a no-go for me.” (German-speaker from Schaffhausen, female).

4.2 Pros and Cons of E-Voting in general

Pros	Cons
Better for young people	Could be a problem for elderly people
Flexibility, efficiency (process and concerning resources), simplicity	Is it really easier? Logins, passwords could be more complicated than by postal vote
Less people and resources needed	The current system functions very well, very easy
	Lack of social contact
	Less work for postman, at the beginning a lot of technical and financial effort
Security	Manipulation (also a problem for other vote systems), hacking, data abuse
Additional technological elements possible (like interactive tools, mistake detection, e.g. concerning elections)	You need internet (but it is a matter of course today)
More ecological	

Table 1. Pros and Cons of E-Voting

When asked about the advantages of e-voting (see Table 1), respondents stressed in particular the flexibility (independence in terms of time and place) and the efficiency and simplification of the process (German-speaker from Aargau, male). One respondent says “you do no longer have an excuse for not going to vote” (Italian-speaker from Ticino, male). Some respondents also mentioned the environmental aspect, which means that less paperwork is needed. However, some interviewees wondered whether it really would be a simplification of the process. They are therefore a little skeptical on this point. An atypical answer was that one advantage would be that “the envelope does not have to be licked like during the postal vote” (German-speaker from Valais, male). Even if many respondents mention the convenience aspect as one of the main advantages of electronic voting, the two interviewees having already used e-voting systems (one abroad, one in Switzerland), underline that they perceived the process itself as being more complicated than the postal voting process because of additional security barriers. However, most interviewees not having used e-voting so far, think that e-voting has the potential to be easier than the other types of voting. They however underline that login and password issues, loading problems or a missing confirmation, that the vote has been validated, could complicate things. Many respondents also think that the current system works very well and they wonder why it should be changed. This general lack of interest is reflected in the responses of the interviewees concerning the question, if they have already informed themselves about e-voting; a question that almost all of the interviewees denied.

Cons (see Table 1) are particularly related to the possibility of manipulation and hacking. The respondents mentioned that people might be afraid of not having guaranteed the anonymity of the data. However, they mentioned this point when talking about

possible reasons why some Swiss citizens are not satisfied with the introduction of e-voting, not concerning themselves. Except for one person, most respondents express some concerns about the security aspects. Many of them also mention that today's systems already work very well and that there is no immediate need to change a well-functioning system to which the Swiss are accustomed. Another reason one respondent mentioned about possible concerns of the Swiss population is that they are worried that other people will not vote as seriously as they need to when voting online. Another argument is the loss of social contact, which already affects postal voting. Respondents also mention that some time must pass before changes are implemented. Respondents also note that older people may have some difficulty voting online.

Therefore, it would be important to retain the other voting options such as postal voting and go to the ballot box. Some respondents argue that it would certainly be more interesting for younger people. One respondent stressed that e-voting is simply not trustworthy in terms of technical aspects and data abuse. Another respondent mentions the Internet's associations with electoral influence in the US and Europe. Another disadvantage of e-voting could be that paper seems to be more "binding" than voting on the Internet.

4.3 Opportunity or Risk for Democracy?

Several respondents think that e-voting is rather an opportunity for democracy than a risk, although one of them underlines the importance of security being guaranteed. Two interviewees think that e-voting will quickly become the norm. Another respondent thinks that it makes voting easier for everyone and more comfortable for Swiss people living abroad.

We also asked the respondents, if they perceive e-voting being a threat for a potential "sacred" meaning ritual dimension of current Swiss voting practices. Most of them do not think that e-voting would mean a "desacralization" of the voting process. Some of them could imagine that this could be the case for other people, who celebrate the social and ritual element of going to the ballot box. One interviewee indicated that this kind of "desacralization" has already happened by introducing the postal vote.

5 Findings Regarding E-Voting Security

5.1 Trust in Relation to Security Aspects

Except for one respondent, the majority trusts the Swiss political system and the e-voting technology. In general, interviewees tend also to trust the government and the professionals dealing with the technology. Two of the interviewees say that it is important to have confidence, as technology is far too complicated for normal citizens to be able to understand everything.

People also generally are accustomed to e-banking, which includes the management of incidents and insurance protection. Most of the interviewees are accustomed to e-banking and thus trust the system. The respondent who is totally against e-voting also

uses e-banking. This respondent adds that e-voting differs from e-banking insofar as the banking system has not been hacked so far, which is however the case for the e-voting system managed by the Swiss post enterprise. This same respondent underlines that “a problem with e-voting leads to a collective damage, whereas a problem with the e-banking system only concerns the money of an individual” (German-speaker from Valais, male). Another interviewee however says that for her the personal damage would be much more important than the collective one. Even another respondent says that, normally it should be possible to realize when a vote is manipulated, as there exist surveys in the run-up to the votes that give an approximate picture of the voting results. Even if the security aspect is very important for all the respondents, some of them relativize the problem by highlighting that the risks of manipulation and loss of anonymity also exist with regard to postal vote and when going to the ballot boxes. However, some of them underline that the damage extent could be much higher with an e-voting system, especially at a national scale, whereas with the conventional types of voting (postal vote and ballot box) the risk seems to be higher at the local level. One respondent even says that e-voting seems more secure to her than the other types of voting. Moreover, some of the interviewees mention that e-voting is especially secure regarding the counting procedure.

5.2 Information Needed for Transparency

Respondents said that it would be quite important for them to be informed about the concrete procedure of the e-voting process in terms of an instruction, as well as in terms of security measures that have been taken. One interviewee however adds that the government’s main aim was to reassure the citizens concerning security measures, which relativized the value of the information. Another interviewee says that an info-button dealing with security aspects could be helpful (German-speaker from Valais, female). As to the concrete procedure to vote online, some of the respondents would like to have an easy operating manual, or explanation in a video.

5.3 Influence on Voting Participation, Voting Results and Voting Decision

While most of the interviewees think that the introduction of e-voting would only slightly influence the voting participation or even not at all, one interviewee guesses that the participation from Swiss people living abroad could increase. This respondent also believes that the possibility to vote directly with the smartphone could be a reason for people to vote more often.

In general, however, the interviewees think that voting participation is a matter of interest and education rather than of the means of voting. Nevertheless, two interviewees say that they personally would vote more often, if they could do it online. In this context, it is important to bear in mind that all the interviewees have a very high participation rate in voting.

As to a potential influence on voting results, most interviewees do not think that the possibility to vote online changes the voting results in a significant way. Some of them

think that it could maybe lead to a slightly higher participation of young people, but not a significant one.

We also asked interviewees what they think about e-voting influencing the voting decision of people. A majority of the interviewees guess that this could be possible in some rare cases or not at all. An example of such a rare case would be that the voter changes his or her mind regarding the voting decision later in time. Hence, when having voted electronically, he or she will not be able any more to correct his or her opinion. When doing a postal vote however, this would still be possible. However, this aspect does not seem to have a big importance for the interviewees. Moreover, one interviewee even said that the e-voting process strengthens the seriousness of voting, as one is conscious of the “seriousness of the moment and the consequences of your choice” (Italian-speaker from Ticino, male). It implies a shift of responsibility, in a digital platform all decisions are made autonomously and at one’s own risk.

6 Findings Regarding E-Voting Operations.

6.1 Important Aspects of the Voting process and Suggestions for improvement

As to the e-voting process, many respondents say that it is especially important to have an easy process. The user should receive a voting confirmation, and even if there are some loading problems, he or she should know if his or her vote was sent. Some of the interviewees also mention that e-voting via smartphone should be possible. The respondent having experience with e-voting abroad mentions that he “really appreciated to be able to choose between postal vote and e-voting” (German-speaker from Aargau, male), meaning that he did not have to choose once and for all for between the two possibilities. Another respondent suggests to use a password that can be scanned with a smartphone and that leads you directly to the right website.

However, when talking about the current voting process, most interviewees say that the process already works very well. One interviewee says that “it is important to ensure in the future that the e-voting will not be the only way but one way to vote” (Italian-speaker from Ticino, male). The most important aspects needing improvement rather concern the steps in the voting process than the voting means.

Two interviewees mention that the voting questions are sometimes asked in a confusing way. Some of the interviewees criticize the content of the brochure the federal government adds to the voting documents. This brochure summarizes the arguments of the advocates and opponents and informs the reader of the recommendations of the federal government. One interviewee thinks that there is not enough space for the people having started an initiative to express their arguments. The opinion of another interviewee goes in the same direction. He thinks that the brochure is too propagandistic and that the neutral aspect is somehow missing. One respondent thinks that there should be more options to vote about, not only yes or no.

When talking about other possibilities how technology could improve the voting process, some of the respondents said that it would be interesting to create an application that informs about the opinions of the parties, and provide access to neutral information sources. Another interviewee mentions the creation of a reward system for participation in votes to increase voting participation. Another idea is to show the party and voting financing in real time. This could make the population more aware of the differences in budget of the different parties. Another respondent says that it is not so much about how and where to vote, but about the collection of largely neutral, correct and professional information that is independent from political power struggles or economic interests. Another person says that it would be interesting if newspapers would provide online information files dealing with the voting topic. One respondent suggests an online “one-stop-shop” like an application, where you can inform yourself about different points of view, vote and see the results.

To sum up, among the respondents there was only one person who is totally against e-voting mainly out of trivialization of the process and security reasons. In general, the interviewees have not or not much informed themselves about e-voting and they are rather happy with and accustomed to the current system, being a matter of habit. Therefore, people need some time to become accustomed to e-voting. Some of them think that it could be more convenient and comfortable to vote online, but estimate the risk of manipulation of voting results as higher than with traditional voting means. In general, e-voting must provide a personal additional benefit for voters, otherwise, they are not so much interested in changing their habits. This confirms the findings of the Australian study [17].

E-voting seems to be an advantage especially for Swiss people living abroad. However, in general, most interviewees think that participation issues do not depend on the voting means, but rather on general political interest and education or on the complexity of the voting topic. Most of them guess that e-voting would not significantly change the participation or even voting results. They are also skeptical of the voting means having an influence on the decision. Some of them think that technology could be useful to improve the information collection stage, as they think that it is not an easy task to inform oneself in a neutral way about quite complex voting topics.

Moreover, the majority of the interviewees trust in the government and technology. They however agree that voting results can always be manipulated, and that the extent of manipulation increases with e-voting. Table 2 summarizes the findings of Section 6.

7 Discussion: Trust and Confidence in the Voting Process

The results of the interviews generally confirm that the perception of electronic voting is strongly associated and correlated with the professional demographics of voters. It is understood that the professional background of electors influences their attitude towards electronic voting. Those working in a highly digitalized world and directly witnessing the digital transformation of their businesses are likely to adopt the electronic voting system as another information and communication technology tool for democracy.

1. Information collection and voting documents	
<i>Difficulties</i>	<i>Proposed solutions</i>
Too much, scattered and biased information, not enough time available for information	<ol style="list-style-type: none"> 1) Application regrouping neutral information 2) Linking e-voting platform to neutral information 3) Newspaper provide online information files 4) Interactive platforms with customized information based on one's political profile 5) State-organized sensitization with regard to the potential of being influenced by social media e.g. (to promote critical opinion forming skills) 6) Inform citizens about budget of the parties for financing their campaign 7) One stop shop for information on different points of view, reminder for votes information, results 8) A personalized digital communication, e.g. save the personal decisions of past votes
Voting questions are not clear enough	Easier formulation of questions
Complex voting topics are reduced to yes or no questions	More choice
Content of the official voting brochure not neutral enough, too easy / propagandistic	Two parts: one easy part and one more intellectual part
Forgetting about vote	Reminder via e-mail or application e.g.
2. Voting process depending on voting type	
<i>Difficulties</i>	<i>Proposed solutions</i>
e-voting: <ul style="list-style-type: none"> - Too complicated - Lack of security 	It has to be as easy as possible (time-saving, comprehensible, intuitive), accessible for everyone <ol style="list-style-type: none"> 1) Password scanning and direct arrival on the website 2) Voting confirmation 3) Possibility to do everything orally
Postal vote: <ul style="list-style-type: none"> - The process is ok as it is - You have to lick the envelope, which is not comfortable - You have to put a stamp 	
3. Information about results	
<i>Difficulties</i>	<i>Proposed solutions</i>
No particular difficulties	Inclusion in application that integrates the whole process, from information collection, to voting, to results

Table 2. Current difficulties linked to the voting process and potential of improvement along the different voting stages

On the other hand, the perception of citizens, who work in security activities that consist in ensuring the security of people and property, with regard to electronic voting is cautious and sometimes very negative. For the latter, there is no tangible evidence of the correct recording of votes because a computer screen can display one thing and record another.

In this field research, trust is highlighted as a particularly important condition for the use of the electronic voting system. Beyond the technical malfunctions that can affect the accuracy and validity of votes, there seems to be a major concern regarding the vulnerability of electronic voting systems to manipulation by external hackers. Although manipulations can occur by internal local agents, these intrusions can only affect local and isolated structures. There are also risks with traditional voting systems with recent examples in some democracies where it was necessary to recount votes. Electronic voting opens up new possibilities for a single hacker, who may be anywhere in the world, to affect the central system disrupting the entire nation. The perception of risk from an electronic perspective is more global, where a group of people of an individual can intervene and hack from anywhere on the planet.

The need to combine secrecy and individual verifiability is one of the most important attributes identified in this research. Indeed, unlike the e-banking system, where citizens can check their account statements at any time to ensure that all transactions are recorded correctly. One can check the accuracy of a bank transaction afterwards, by checking account statements online or by printing on official banking documents. Electronic voting operations concerning secrecy and verifiability (universal and individual) are still not clear to voters. In addition, in the event of embezzlement, the bank has insurance, through its after-sales service, to compensate its customers. All information necessary for data integrity can be stored and tracked, eliminating any secrecy between the customer and the bank. In the event of technical malfunction or external interference, these failures or manipulations may go unnoticed.

Swiss citizens will probably trust the electronic voting system and their officials. In the case of electronic voting, their trust does not need to be earned at this time, but it can only be easily lost. The technical challenges of electronic voting open doors for private companies and investors to sell their technologies. There is concern that personal data and individual voting preferences may be in their possession. It is imperative to maintain the confidence that only the government should control all systems. The government must be the sole guarantor of the entire system. With the trust established with the government, the voter is not very interested in learning more about verifiability or verification.

The results also show that the electronic voting system is accepted as a complementary means, but should not replace the postal voting system and the ballot box. This new possibility of voting will not significantly increase the number of people who participate in the vote or change their voting habits, as people vote because they want to, and not because of the different possibilities offered.

In a traditional voting system, voters have learned to trust their fellow assessors not to open the ballot box before the counting, which is a public operation understood by

all voters. In the case of wrongdoing, citizens are able to accept human fault because they too can make mistakes in their own work. However, it is still difficult to accept and understand the mistakes made by software. The government official could then keep some form of electronic voting ceremony such as broadcasting live the count happening in the back-stage process.

8 Conclusion

The objective of this exploratory research was to focus on the social aspects of electronic voting. Indeed, an abundant scientific literature already covers the technical aspects of e-voting (algorithms, security, etc.). The justification for this research objective, which focuses on rather social aspects, stems first of all from the fact that electronic voting, for various reasons, is barely fully established in Switzerland, whereas many test phases have already demonstrated the feasibility of this operational process and it has recently been accepted as an official way of voting. In this research, we intended to study the perceptions of Swiss citizens regarding electronic voting, and to understand their meanings through a qualitative field research (based on semi-structured interviews). The most important findings are presented in the three following categories: electronic voting perception, electronic voting security and electronic voting operations. These insights are essentially related to the social acceptance of electronic voting. We observe in particular that the vote in Switzerland has an almost sacred dimension, as opposed to other cultural contexts, such as [5] [13] [17] [18] and that the trust that surrounds the voting “ritual” (i.e. voting operations and processes) is of supreme importance.

Through this research project on electronic voting, which focuses on the human factor dimension, we believe that this study will help to develop relevant and practical managerial precepts to better implement these digitized operations processes in the future.

Nevertheless, this study is not without limitations. As the data collection was just completed when writing the paper, more in-depth analysis iteration of database must still be completed to improve the relevancy of our conclusion. The size of the sample (25) is still too limited to draw more generic conclusions. It was however our aim to investigate the “how” and “why” of perceptions of Swiss citizens rather than percentages as it is the case in a quantitative survey. In a further research, we intend to conduct a quantitative survey based on the qualitative findings presented in this paper.

References

1. Fenazzi S., Jaberg S.: Le vote électronique ne devient pas un canal de vote ordinaire en Suisse. Available via Swissinfo.ch (2019)
https://www.swissinfo.ch/fre/marche-arri%C3%A8re-gouvernementale_le-vote-%C3%A9lectronique-ne-devient-pas-un-canal-de-vote-ordinaire-en-suisse/45060752
Accessed 02 July 2019

2. Budurushi J., Neumann S., Renaud K., Volkamer K.: Introduction to special issue on e-voting. *Journal of Information Security and Applications*. 38, 122–123 (2018)
3. Gray, J., Rump, B. (2015). Models for digitalization. *Software and Systems Modeling*. 14, 1319–1320 (2015)
4. Hagberg, J., Sundstrom, M., Egels-Zandén, N.: The digitalization of retailing: an exploratory framework. *International Journal of Retail and Distribution Management*, 44, 694–712 (2016)
5. Moisander, J., Eriksson, P.: Corporate narratives of information society: Making up the mobile consumer subject. *Consumption, Markets and Culture*. 9, 257–275 (2006)
6. Eshet-Alkalai, Y.: Real-time thinking in the digital era, in: *Encyclopedia of Information Science and Technology*, Second Edition. IGI Global. 3219–3223 (2009)
7. Liyanage, J.P.: *Hybrid Intelligence Through Business Socialization and Networking: Managing Complexities in the Digital Era*. IGI Global, Hershey, USA (2012)
8. Gibson J. P., Krimmer R., Teague V., Pomares J.: A review of e-voting: the past, present and future. *Annals of Telecommunications*. 7, 279–286 (2016)
9. Mahrer H., R. Krimmer, R.: Towards the enhancement of e-democracy: identifying the notion of the ‘middleman paradox’. *Information systems journal*. 15, 27–42 (2005)
10. Alvarez, R., Hall, T., Trechsel, A.: Internet Voting in Comparative Perspective: The Case of Estonia. *Political Science and Politics*. 42, 497–505 (2009)
11. Vassil K., Solvak M., Vinkel P., Trechsel A.H. Alvarez R.M. The diffusion of internet voting. Usage patterns of internet voting in Estonia between 2005 and 2015. *Government Information Quarterly*. 33, 453–459 (2016)
12. Germann, M., Serdült, U.: Internet voting and turnout: Evidence from Switzerland. *Electoral Studies*. 47, 1–12 (2017)
13. Germann M., Serdült U.: Internet voting for expatriates: The Swiss case. *JeDEM-eJournal of eDemocracy and Open Government*. 6, 197–215 (2014)
14. Wang, K.-H., Mondal, S. K., Chan, K., Xie, X.: A review of contemporary e-voting: Requirements, technology, systems and usability. *Data Science and Pattern Recognition*. 1, 31–47 (2017)
15. Tarasov P., Tewari H.: The future of e-voting. *IADIS International Journal on Computer Science and Information Systems*. 12, 148–165 (2018)
16. Willemsen J.: Bits or paper: Which should get to carry your vote? *Journal of Information Security and Applications*. 38,124–131 (2018)
17. Zada, P., Falzon G., Kwan P.: Perceptions of the Australian public towards mobile internet e-voting: risks, choice and trust. *Electronic Journal of e-Government*. 14, 117–34 (2016)
18. Suki N.M., Suki N.M.: Decision-making and satisfaction in campus e-voting: moderating effect of trust in the system. *Journal of Enterprise Information Management*. 30, 944–963 (2017)
19. Ghauri, P., Gronhaug, K.: *Research Methods In Business Studies* (3rd ed.). England: Pearson Education (2005)
20. Palinkas L.A., Horwitz S. M., Green C. A., Wisdom J. P., Duan, N., Hoagwood, K: Purposeful sampling for qualitative data collection and analysis in mixed method implementation research. *Administration and Policy in Mental Health and Mental Health Services Research*. 42, 533-544 (2015)
21. OECD Education at a glance 2014. Available via Country note. <http://www.rsc.org/dose/title> of subordinate document. Accessed 02 July 2019 (2014)

The Swiss Post/Scytl transparency exercise and its possible impact on internet voting regulation

Ardita Driza Maurer¹

¹ Centre for Democracy Studies Aarau / University of Zurich
ardita.drizamaurer@uzh.ch

Abstract. In Switzerland, internet voting has been in the experimental phase for over fifteen years. With a view to putting an end to trials and normalizing its use alongside the paper-based channels (polling station and postal voting), a thoroughly updated federal regulation entered into force in January 2014. Only systems that are formally certified and offer complete verifiability can be authorized to propose internet voting in an unrestricted manner, i.e. to all the electorate. Furthermore, since July 2018, the publication of the source code of fully verifiable systems is mandatory. A major transparency exercise took place in February – March 2019. The first system to introduce complete verifiability – the Swiss Post/Scytl system – was submitted to a public intrusion test (PIT), open to anyone interested. In a parallel development, the source code of the same system was published on the internet. Researchers found critical errors in the source code of both individual and universal verifiability. The PIT revealed other, less critical issues. This experience has fuelled the already heated debate over the future development of internet voting in Switzerland. It questions the procedures for controlling verifiability solutions and, ultimately, the consensus to develop such solutions. Lessons learned will most probably be reflected in the future update of the regulation.

Keywords: Switzerland, internet voting, regulation, security, transparency, public intrusion test (PIT), source code publication.

1 Introduction

Debate on internet voting in Switzerland focuses on security and transparency. After initial experiences with “black-box”¹ internet voting systems in political elections in several countries, including Switzerland, at the beginning of 2000, consensus emerged within the research community that end-to-end verifiable voting systems are a necessary condition for internet voting [1].² Systems started to be developed that may allow the voter and anyone else to verify important aspects of the election, namely his/her

¹ We use this term to characterise first generation internet voting systems introduced in the beginning of 2000 which did not provide for independent, transparent verifications.

² See also the 2007 Dagstuhl Accord, <http://drops.dagstuhl.de/portals/index.php?semnr=07311>.

All links were last checked on 28 June 2019.

own vote and the final tally, while protecting the secrecy of the vote, without introducing any additional danger of improper influence of the voter as compared to postal voting and without relying on trust in persons, processes, devices or software. According to this consensus, the challenge for government and civil society should be to find ways to foster development and testing of new election paradigms in general and to allow them to be assessed and expeditiously rise to meet their potential to improve elections, the goal being to develop systems that increase transparency regarding the correctness of the election results and yet maintain secrecy of individual votes. Improved voter confidence may follow.³ Proper implementation of such systems as well as voter education are considered important to avoid misuse. Recent developments in Switzerland show that control of end-to-end verifiability solutions and requirements thereof are crucial.

Complete verifiability is required by federal regulation if a system is to be authorized to cover more than 50 per cent of the cantonal electorate [2].⁴ It is the sum of extended individual verifiability and universal verifiability. Extended individual verifiability allows the voter to ascertain whether their vote has been manipulated or intercepted on the user platform or during transmission. Voters must receive proof that the server system has registered the vote as it was entered by the voter on the user platform as being in conformity with the system. The proof must also confirm to the voters that the data relevant to universal verification has reached the trustworthy part of the system. Voters (rather “electors” in this case, i.e. persons with voting rights but who did not vote) must be able to request proof after the electronic voting system is closed that the trustworthy part of the system has not already registered a vote cast using their client-sided authentication. For universal verification, the auditors receive proof that the result has been ascertained correctly. The proof must confirm that the result ascertained: a. takes account of all votes cast in conformity with the system that were registered by the trustworthy part of the system; b. takes account only of votes cast in conformity with the system; c. takes account of all partial votes in accordance with the proof generated in the course of the individual verification.⁵ Verifiability relies on several trust assumptions.⁶

The development of end-to-end verifiable systems provides valuable real-world experience. One of the two Swiss internet voting systems, the Swiss Post/Scytl system, became the first to allegedly introduce complete verifiability⁷ after it had been certified to offer individual verifiability.⁸

³ *Ibid.*

⁴ The definition of complete verifiability is to be found in article 5 read in combination with article 4 of the federal Chancellery Ordinance on Electronic Voting (VEleS), RS 161.116.

⁵ *Ibid.*

⁶ See e.g., art. 4 para. 4 and 5 as well as art. 5 para. 3 let. c and para. 5 and 6 VELeS

⁷ <https://www.post.ch/-/media/post/evoting/dokumente/complete-verifiability-security-proof-report.pdf?la=fr&vs=1>

⁸ Individual verifiability is required for authorization for more than 30 per cent of the cantonal electorate, whereas complete verifiability is required for more than 50 per cent (art. 27f PRO and articles 4 and 5 VELeS). The Swiss Post system was the first and eventually only system

The Swiss Post set the objective to present a system offering complete verifiability by the end of 2018.⁹ In this context, it underwent the most complete transparency exercise organized so far on a Swiss internet voting system and, to our knowledge, the most complete on an internet voting system for political elections. The system was submitted to a public intrusion test (PIT) decided by the federal Chancellery and the cantons¹⁰ which took place from 25 February to 24 March 2019.¹¹ In a parallel development, the Swiss Post and its partner, the Spanish firm Scytl, published the source code of their software on 7 February 2019,¹² in accordance with the federal requirement to do so which came into force in July 2018.¹³ The publication of the source code should take place when the system has the property of complete verifiability in terms of article 5 VELeS and after successfully passing the examinations foreseen in article 7 VELeS.¹⁴

(as Geneva decided to stop developing its system) to be certified for more than 30% of the cantonal electorate. See fn. 14.

⁹ <https://www.post.ch/en/business/a-z-of-subjects/industry-solutions/swiss-post-e-voting?shortcut=evoting>

¹⁰ <https://www.bk.admin.ch/bk/en/home/dokumentation/medienmitteilungen.msg-id-73898.html>. See also <https://www.evoting-blog.ch/en/pages/2019/public-hacker-test-on-swiss-post-s-e-voting-system>

¹¹ See <https://onlinevote-pit.ch> and <https://pit.post.ch/en>

¹² <https://www.post.ch/en/business/a-z-of-subjects/industry-solutions/swiss-post-e-voting/e-voting-source-code?shortcut=evoting-sourcecode>

¹³ Article 7a and 7b of the federal Chancellery Ordinance on Electronic Voting (VELeS).

¹⁴ Two types of examinations are foreseen: for less than 30% of the electorate (paragraph 3) and for more than 30% (paragraph 2). The Swiss Post system had successfully passed a number of examinations required by paragraph 2 of art. 7 VELeS for more than 30 per cent of the electorate, in May and June 2017. The certificates issued end June 2017 are valid till end June 2020. The information is published on <https://www.post.ch/en/business/a-z-of-subjects/industry-solutions/swiss-post-e-voting>. The examinations/certificates published are the following:

- Verification of the cryptographic protocol <https://www.post.ch/-/media/post/evoting/dokumente/zertifikat-pruefung-des-kryptographischen-protokolls.pdf?la=en&vs=1>
- Verification of functionality <https://www.post.ch/-/media/post/evoting/dokumente/zertifikat-pruefung-der-funktionalitaet.pdf?la=en&vs=1>
- Verification of infrastructure and operation <https://www.post.ch/-/media/post/evoting/dokumente/zertifikat-pruefung-infrastruktur-und-betrieb.pdf?la=en&vs=1>
- Verification of protection against attempts to infiltrate the infrastructure <https://www.post.ch/-/media/post/evoting/dokumente/zertifikat-pruefung-des-schutzes-gegen-versuche-in-die-infrastruktur-einzudringen.pdf?la=en&vs=1>

We could not find information on the internet about the examination required by art. 7 paragraph 2 let. e (printing offices) and f (control components) VELeS on the internet. We take for granted that “the disclosed source code relates to the implementation of the cryptographic protocol for complete verifiability at application level” (see <https://www.post.ch/en/business/a-z-of-subjects/industry-solutions/swiss-post-e-voting/e-voting-source-code>) and that all preconditions for doing so (art. 7a VELeS) were respected.

A group of researchers discovered significant flaws in the source code [3].¹⁵ As for the PIT, a total of 16 responses were classified as breaches of best practice. According to the federal Chancellery, they do not constitute major risks.¹⁶

The federal Chancellery declared itself satisfied that these measures (PIT and publication of source code) led to the discovery of weaknesses and allowed important findings to be made. It also declared that it would conduct a review namely of the licensing and certification procedures for e-voting systems. The Swiss Post decided to suspend internet voting until the source code and other identified errors are addressed and not to offer e-voting at 19 May 2019 federal vote. The federal Chancellery considered the decision on the part of Swiss Post not to make its system available for the vote on 19 May to be logical under the circumstances.¹⁷

The next federal vote is the federal (National Council) election of 20 October 2019. Requirements for authorizing use of e-voting at federal elections are stringent [4]. The correction of the source code most probably classifies as “substantive change” which should be followed by tests and a new certification [5].¹⁸ The certification requirements are currently under review by the federal Chancellery.¹⁹ Given this, it is questionable whether the Swiss Post system can be ready in time for the 2019 federal election. The federal Government will decide on authorizing the Swiss Post system to use electronic voting in the federal election of 20 October foreseeably on 14 August 2019, provided cantons working with the Post will apply for such an authorization.²⁰

The second system belongs to the canton of Geneva and is operated by its administration. It offers individual verifiability but not the universal one. Geneva system was used for the 19 May 2019 vote. It has not been formally certified so far and is authorized for less than 30% of the cantonal electorate. End November 2018, the Geneva Government announced it would cease operating its e-voting system in 2020 for lack of financial support to upgrade it to a fully compliant system, namely, to set up a new system that offers complete verifiability and have it certified.²¹ On 19 June 2019 the Geneva Government decided to stop e-voting with immediate effect because of uncertainties around the possible authorization by the federal Government to use e-voting at the October 2019 federal election. The canton of Geneva and the other cantons working with it on internet voting estimated that the expected moment for the federal Government decision (14 August 2019) did not leave enough time to adapt the procedures in case the decision would be negative.²²

¹⁵ <https://www.bk.admin.ch/bk/en/home/dokumentation/medienmitteilungen.msg-id-74307.html>

¹⁶ <https://www.bk.admin.ch/bk/en/home/dokumentation/medienmitteilungen.msg-id-74508.html>

¹⁷ *Ibid.*

¹⁸ See article 27/ 2 PRO and article 7 paragraph 1 VEleS.

¹⁹ <https://www.bk.admin.ch/bk/en/home/dokumentation/medienmitteilungen.msg-id-74508.html>

²⁰ This was still pending on 28.06.2019 when this paper was last reviewed. See <https://www.swissinfo.ch/ger/schweiz-demokratie-volksabstimmungen-evoting/45061040>

²¹ <https://www.ge.ch/document/point-presse-du-conseil-etat-du-28-novembre-2018>

²² <https://www.ge.ch/document/point-presse-du-conseil-etat-du-19-juin-2019>

Another potentially disruptive development started in March 2019: the collection of signatures in support of a popular initiative – the so-called e-voting moratorium initiative – to stop any form of e-voting for at least five years.²³ The initiative aims at changing the federal Constitution to prohibit e-voting. It foresees a possible ban lift by the federal Parliament, through a law, which can be introduced at the earliest five years after the introduction of the ban. Several cumulative conditions should be fulfilled to lift the ban, namely: e-voting should offer at least the same level of security against manipulations than paper voting; it should allow voters without specialized knowledge to verify the main steps of the e-voting procedure and enable count-as-cast of cast-as-intended votes while also respecting vote secrecy; the system should exclude external influences and should make sure that results are unequivocal and un falsified; results can be verified in a sure manner and without special knowledge through new counting; it should be possible to exclude results that do not respect the before-mentioned requirements. One of the conditions, namely the possibility for the layman to understand and control every important step without specialized knowledge, seems, at first reading, impossible to achieve.²⁴ The 18 months signature collection period ends on 12 September 2020. If the initiative committee gathers the required one-hundred-thousand valid signatures, the fate of internet voting will be decided in a popular vote by the majority of the people and cantons.

These developments unfold in the context of the implementation of a federal Government's decision of April 2017 to introduce internet voting into regular operation alongside the postal and polling-station voting.²⁵ The federal Council submitted in December 2018²⁶ a proposal to modify the federal Act on political rights (PRA) [6] in this sense. The proposed modification upheld the current requirements for internet voting and proposed to improve the structure of the regulation by bringing core principles of complete verifiability, transparency, certification, risk assessment framework and accessibility to the level of the law instead of having them at the ordinances' level, as is currently the case. The normalized use of e-voting would have put an end to the experimental phase that lasts since 2004. The proposed revision of the PRA was submitted to a consultation procedure from 19 December 2018 to 30 April 2019. Cantons and interested organizations were invited to comment on the proposal. The results of the consultation were published end June 2019.²⁷ They show that developments around the Swiss Post transparency exercise influenced the debate. The consultation revealed that most respondents, including a clear majority of the cantons and political parties, support the introduction of e-voting in principle. However, most respondents, including political parties which support e-voting in principle, also considered its introduction into regular operation to be premature. On 26 June 2019, the

²³ Initiative populaire fédérale « Pour une démocratie sûre et fiable (moratoire sur le vote électronique) », FF 2019 2081 (“FF” is an abbreviation of the Swiss Federal Gazette).

²⁴ A few months earlier the federal Parliament had refused such a “layman control” on e-voting https://www.swissinfo.ch/eng/boost-for-expat-swiss-group_opponents-of-e-voting-suffer-setback-in-parliament/44395904

²⁵ <https://www.bk.admin.ch/bk/fr/home/documentation/communiqués.msg-id-66273.html>

²⁶ <https://www.bk.admin.ch/bk/fr/home/documentation/communiqués.msg-id-73491.html>

²⁷ <https://www.newsd.admin.ch/newsd/message/attachments/57568.pdf>

federal Council took the decision “to provisionally forgo introducing electronic voting into regular operation” and “not to proceed with the partial revision of the Political Rights Act at the present time”.²⁸ Internet voting’s introduction as a regular voting channel is thus technically delayed. The federal Council also commissioned the federal Chancellery “to amend the general conditions for future trials” namely “to redesign the way in which the trials are operated, and to present the results in a report by the end of 2020. The aim is to establish stable trial operations using the latest generation of systems. Other measures include extending independent audits, increasing transparency and trust, and greater involvement of scientific specialists”.²⁹

The following sections focus on lessons learned from the PIT and the publication of source code, from a regulatory point of view. After an overview of the federal legal requirements on security and transparency (section 2), we present the PIT and the publication of the source code and related events (section 3). The results call into question the current regulation, particularly the control requirements for end-to-end verifiable systems and, ultimately, the consensus on the role and adequacy of such solutions. The experience has already had an impact on regulation as it prevented the amendment of PRA and the introduction of e-voting into regular operation. It will continue to have an impact as the federal Chancellery is expected to amend the conditions for future trials and decide later on its transformation into a regular voting channel (section 4).

2 Internet voting development in Switzerland

2.1 Federal regulation of internet voting

Switzerland has adopted a cautious approach to internet voting which is reflected in the long experimental phase. E-voting has been tested with binding effect in political votes and elections for more than 15 years. The motto is “security before speed”. At the same time, Switzerland has a unique situation: its direct democracy system imposes frequent votes at all levels of government. Electors, i.e. the persons with voting rights, are invited to vote on issues or elect representatives at local, cantonal (state) and federal levels an average four times a year. It is thus important to find ways and means to offer effective voting channels to a maximum of electors, including those living abroad and those with special needs.

At the beginning of the years 2000 Swiss authorities concluded that any use of internet voting in the political field required a legal basis [7]. Federal regulation of internet voting, including a dedicated article (art. 8*a*) and other modifications in the political rights Act (PRA) [6] and a dedicated chapter (art. 27*a* ff.) in the political rights ordinance (PRO) [5] was introduced in 2002 and has been in force since 1st January 2003. Swiss cantons started internet voting trials in 2003 (cantonal votes) and 2004 (federal votes). The federal Council (federal Government) evaluated the trials in

²⁸ <https://www.bk.admin.ch/bk/en/home/dokumentation/medienmitteilungen.msg-id-75615.html>

²⁹ *Ibid.*

2006 [8] and 2013 [9].³⁰ In 2006 it decided to continue to experiment internet voting and extend the trials to include the Swiss abroad and new cantons. New forms of cooperation developed between cantons with an internet voting system (Geneva, Zurich and Neuchâtel) and those without system. Fifteen out of the 26 cantons have tried internet voting so far; the majority outsources the internet voting service to another canton with a system (Geneva until June 2019) or to a privately held system (currently, the Swiss Post). In its third evaluation report of 2013, the federal Council decided to continue to use internet voting but to gradually replace the “black box” first generation systems with “end-to-end verifiable” systems. As a result, the federal regulation was thoroughly modified in December 2013: the federal Ordinance on Political Rights (PRO) was updated and a new instrument, the Ordinance of the federal Chancellery on e-voting (VEleS) was introduced,³¹ both in force since 15 January 2014. An additional requirement became mandatory as of 1st July 2018: the publication of the source code of the software of complete verifiability as well as the procedure for its publication (see articles *7a* and *7b* VELeS).³²

Regulation is based on the idea that e-voting must respect all principles applicable to democratic votes and elections and the ensuing legal requirements.³³ The federal regulatory framework for e-voting has a cascade structure that includes the Constitution and the higher-level formal law (PRA), the federal Council ordinance (PRO) that implements the PRA and, further down, the federal Chancellery ordinance on electronic voting (VELeS) and its Appendix which contain detailed provisions that implement the higher level requirements to e-voting. This structure allows for a relatively quick adaptation of the detailed provisions (VELeS) to reflect technical developments and good practices which are considered important in the security area.³⁴ Generally speaking, the federal regulation requires that e-voting systems and their security are state of the art, as stated in art. 27/ para. 1 let. b PRO.

According to federal regulation, use of internet voting at federal votes is furthermore subject to authorization by the federal Council and agreement by the federal Chancellery.³⁵ Different levels of compliance and respective limitations are foreseen.³⁶ The Swiss Post system became the first to be formally certified compliant with regulation for systems providing individual verifiability, potentially allowed to cover up to 50% of the electorate. End 2018 it was expected to become fully compliant with the federal regulation for systems providing complete (individual and universal) veri-

³⁰All evaluations can be found at <https://www.bk.admin.ch/bk/fr/home/droits-politiques/groupe-experts-vote-electronique/rapports-et-etudes-concernant-le-vote-electronique.html>

³¹ RO 2013 5365 and RO 2013 5371

³² RO 2018 2279

³³ At the federal level e-voting must comply namely with the principle of free elections of art. 34.2 of the federal Constitution (Cst., RS 100), the principles mentioned in article *8a* PRA, which is the legal basis for introducing e-voting, and the detailed provisions of articles *27a* ff PRO, of VELeS and its appendix.

³⁴ The VELeS Appendix contains several references to good/best practice.

³⁵ Art. *8a* para. 1 and *1^{bis}* PRA, art. *27a* and *27e* PRO.

³⁶ Art. *27f* PRO. See also the discussion in Puiggali/Rodriguez-Pérez (2018).

fiability which opens the door to authorization to cover up to 100% of the electorate.³⁷ At this point, it was required to pass two important tests: a public intrusion test (PIT)³⁸ and the publication of the source code of the software for complete verifiability in compliance with the VELeS requirements for doing so.³⁹

We do not refer to cantonal legislation as it is less detailed and mainly a repetition of federal provisions. In principle, cantons have important autonomy in the electoral field [10]. However, with respect to internet voting, the main requirements, namely those related to security, are defined at the federal level and are the same across the country and the systems.

2.2 Federal requirements on security and transparency

The federal regulation of internet voting introduced in 2002 was quite a detailed piece of legislation which also inspired the development of the Council of Europe 2004 Recommendation on e-voting [11, 12]. Electoral authorities controlled the implementation of security related requirements. External audits were conducted but the findings were not published.⁴⁰ Privileged players, i.e. federal authorities in the context of the authorization procedure and the electoral commission in cantons where it existed, had access to the documentation. Political parties represented at the electoral commissions, namely in Geneva and Neuchâtel, could access the documents, which is a good practice.⁴¹ A form of peer-control was provided by federal groups accompanying each cantonal e-voting project whose members are e-voting specialists from other cantons. However, security and transparency of first generation systems, and indirectly the regulation on which they were based, was criticized by research which referred to

³⁷ On its web page, the Swiss Post says that the advantage of its e-voting solution is that it “offers state-of-the-art technology and, in its most advanced phase of development, meets all statutory provisions”. To the attention of cantons and municipalities it says that its solution is “Certified for all eligible voters resident in Switzerland and abroad”, <https://www.post.ch/en/business-solutions/e-voting/the-e-voting-solution-for-cantons>.

³⁸ See fn. 10

³⁹ Art. 7b VELeS

⁴⁰ In its second report on e-voting the federal Government said that “the technical documentation including evaluations of an e-voting system and its security are cantonal confidential documents that are annexed to the request for authorization addressed by a canton to the federal Council. These documents are not public. Cantons that apply the transparency principle can attach conditions to the consultation by the public of these documents and source codes or even refuse access to the extent that they contain sensitive security information or trade secrets. This practice has been upheld by the federal Court” (our translation), FF 2006 5205, 5215. In its third report, the federal Council reminds that only one canton (Geneva) had introduced legislation on limited access to the source code, FF 2013 4519, 4596 f.. The federal Council notes that the mid and long-term objective is to achieve maximum transparency without violating legal or contractual obligations.

⁴¹ Third report of the federal Government on e-voting, point 5.4.4, FF 2013 4519, 4600

them as “security by obscurity” approach [13].⁴² To sum up, first generation systems introduced in the beginning of 2000 did not provide for independent, transparent verifications. They were not submitted to federal requirements to divulgate the source code or security relevant documents.⁴³

The thoroughly revised regulation introduced in December 2013 resulted from close cooperation with research.⁴⁴ It has the following general approach to security and transparency issues. The higher-level principles that e-voting should satisfy⁴⁵ are as many objectives that an e-voting system and program should fulfil to receive federal authorization. The objectives take into account the weaknesses inherent to the underlying technology, as well as main threats, both internal and external ones. Threats include malware on the client or server side, DNS spoofing, MITM attacks, administrator attacks both on the content of votes and on the secrecy of the information, criminal organisations’ attacks, DOS etc. Switzerland having already a generalised system of distant postal voting, threats related to “family voting” are not considered as they are not specific to e-voting [9].⁴⁶ Risks must be constantly evaluated and kept at an acceptable level by the cantons. A risk arises if a weakness in the system can be exploited by a threat and therefore the fulfilment of a security objective is potentially jeopardised. Threats and vulnerabilities inherent to e-voting should be monitored permanently and appropriate countermeasures are introduced whenever necessary by federal and cantonal authorities.⁴⁷

⁴² For an overview of major weaknesses that technical research identified in first generation systems and proposals to correct them in second generation systems, see in particular Dubuis, Haenni, Koenig (2012), pp. 10 ff, in particular points 1, 2, 5 and 11.

⁴³ Security was mainly based on measures taken by the voter to protect her own computer, on the discouraging effect of penal law provisions and on the security provided by the system itself at the structural, functional and technical levels. The fact for e-voting to be only a complementary voting method, not an exclusive one, was considered relevant to its security: See the first report of the federal Government on e-voting, FF 2002 612, 632 ss, 640.

⁴⁴ The main novelties of the new regulation introduced in 2013, namely verifiability and formal certification, as well as the source code publication introduced in 2018, reflect proposals by technical research. The federal Chancellery accompanied the publication of the Berner Fachhochschule study on the concept and implications of verifiable e-voting systems of 21 February 2012 with a note saying that, although the full implementation of the proposals of BFH is to be considered in the long term, nothing prevents (the authorities) from integrating them already in the daily work of improving the systems (our translation), <https://www.bk.admin.ch/bk/fr/home/droits-politiques/groupe-experts-vote-electronique/rapports-et-etudes-concernant-le-vote-electronique.html>.

⁴⁵ Mainly found in art. 27b PRO. Some of them, such as the publication of source code are currently to be found in the lower level VELeS. The proposed modification of PRA aimed at bringing the main principles from PRO and VELeS up to the PRA level. As discussed above, the Government decided on 26 June 2019 to postpone the PRA amendment.

⁴⁶ This being so, critique on end-to-end verifiable systems related to secrecy does not affect the Swiss verifiability solution. With respect to such critique, see e.g. Jones D.W.: Some problems with end-to-end voting (2009).

⁴⁷ VELeS Appendix, point 3 “Security requirements”

The regulation admits that absolute security is impossible to achieve in e-voting, or in any other voting channel for that matter.⁴⁸ Optimum security is the objective.⁴⁹ It rests on three pillars: strong requirements (federal regulation of e-voting refers to state of the art solutions), controls by independent and competent bodies of the conformity of the system with requirements (incl. formal certification)⁵⁰ and the possibility to detect possible problems that may still arise during the voting or counting process (plausibility and verifiability checks).⁵¹ If more than 30% of the cantonal electorate are to be authorised to participate in e-voting, the system and its operation must be examined in particular detail with regard to several criteria:⁵² control of the cryptographic protocol which can be done by a highly specialised institution upon approval by the federal Chancellery; control of other aspects (functionalities, security of infrastructure and operation, protection against attempts to infiltrate the infrastructure, requirements for printing offices and control components) which is to be carried out by an institution accredited by the Swiss Accreditation Service (SAS).⁵³

A third provision, important for security, came into force in July 2018: the publication of the source code of systems that offer complete verifiability. The source code should be published only after the system has been certified. In the words of the federal Chancellery, a trustworthy control prior to publication guarantees that the advantages of the publication of the source code outweigh the potential risks associated with it [14].⁵⁴ Further, the publication should be done in line with good practice to make sure that interested persons have effective access to the source code and the time needed to analyse it and to submit remarks. In particular, the source code should be prepared and documented in line with good practice.⁵⁵ Access should be simple and free.⁵⁶ The documentation on the system and its operation must explain the relevance of the individual components of the source code for the security of electronic voting. The documentation must be published along with the source code.⁵⁷ Finally, anyone is entitled to examine, modify, compile and execute the source code for ideational purposes, and to write and publish studies thereon.⁵⁸ This provision integrates and goes beyond good cantonal and international practice.⁵⁹ The legal requirement to publish the source code marks a new approach in e-voting security, in line with good

⁴⁸ Already the first report of the federal Government on e-voting in 2002 noted that « permanent and absolute security is illusory », FF 2002 612, 639.

⁴⁹ VELeS Appendix, point 3 “Security requirements”

⁵⁰ Art. 27/ PRO

⁵¹ art. 27*i* PRO

⁵² Art. 7 para. 2 VELeS

⁵³ Appendix VELeS, chapter 5.

⁵⁴ See in particular comments on art. 7*a*, al. 2, VELeS in reference [14]

⁵⁵ Art. 7*b* para. 1 VELeS

⁵⁶ Art. 7*b* para. 2 VELeS

⁵⁷ Art. 7*b* para. 3 VELeS

⁵⁸ Art. 7*b* para. 4 VELeS

⁵⁹ Canton Geneva introduced legislation on source code publication already in 2016. An important previous milestone was the publication of the source code of the Norwegian system.

practice and suggestions from research: security is no longer linked to secrecy but to openness and independent verification [15].

To summarize, the regulation requires state-of-the-art security measures. Control of compliance with the regulation and detection of problems rely mainly on certification, verifiability and publication of the source code. These controls are expected to prove a system's conformity with requirements and the absence/presence of potential problems during implementation and should themselves be conducted in a state-of-the-art fashion. The Swiss regulation on security and transparency of internet voting is quite detailed and integrates research recommendations and good practice. It is the first standardisation and certification framework for online voting systems [16]. However, the PIT and source code publication revealed lacunae and raise questions.

3 Public intrusion test and publication of the source code of the Swiss Post/Scytl system

Intrusion tests are required by federal regulation to check a system's security. They should be organized at least every three years and be conducted by an accredited organism as part of the certification process.⁶⁰ The federal Chancellery and cantons decided to organize a public intrusion test (PIT), open to anyone, to check the security of the Swiss Post system offering complete verifiability.⁶¹ The PIT took the form of a "bug bounty" with the Swiss Post committing financial compensation to participants who would be the first to reveal a relevant vulnerability. The Confederation contributed a substantive amount (250'000 Swiss francs) to the "bug bounty" fund. The PIT lasted one month, from 25 February to 24 March 2019. Around 3,200 people from 137 countries participated.⁶² The PIT was accompanied and monitored by a management committee composed of members of the Confederation and the Cantons. The management committee should prepare a final report to the attention of the Steering Committee of the federal internet voting project.⁶³ The PIT participants discovered least severe vulnerabilities which include findings that show uncritical optimization opportunities.⁶⁴

The most critical vulnerabilities were discovered by examining the source code of the Swiss Post system, whose publication was done in line with the newest requirements of VELeS. According to researchers of the Berner Fachhochschule, these vul-

⁶⁰ Point 5.5 of Appendix to VELeS

⁶¹ Fn. 10

⁶² Swiss Post press release of 29 March 2019 "Facts and figures on the public intrusion test on the e-voting system", <https://www.post.ch/en/about-us/company/media/press-releases/2019/swiss-post-temporarily-suspends-its-e-voting-system>

⁶³ Federal Chancellery's information on the PIT https://www.bk.admin.ch/bk/en/home/politische-rechte/e-voting/oeffentlicher_intrusionstest.html

⁶⁴ The 16 accepted vulnerabilities published on the PIT page <https://www.onlinevote-pit.ch/stats/> are classified as breaches of best practice which do not constitute major risks. See "Qualifying vulnerabilities" on <https://www.onlinevote-pit.ch/conduct/>

nerabilities were already apparent in the system specification documentation; the PIT and publication of source code only played a secondary role in their detection.⁶⁵

The code was published⁶⁶ on the platform GitLab and made available upon registration and acceptance of the terms of use, among which the requirement to publish the findings only 45 days afterwards.⁶⁷ The published code “leaked” in the sense that researchers who did not accept the terms of use, received it from others and were able to examine it. A group of them detected major vulnerabilities affecting the universal and individual verifiability [3].⁶⁸ They communicated their findings ongoing on Twitter, after a short advance notice to the Post, thus in breach of the 45 days deadline. Although the distribution of the source code to third parties who have not accepted the terms of use is prohibited according to the terms of use, the Swiss Post and the federal Chancellery didn’t mention this detail in their communications and took notice and reacted after each published finding.⁶⁹

The first critical error discovered related to universal verifiability.⁷⁰ A trapdoor was found that would allow the system operator or any person with access to the system to modify any number of votes in a way that cannot be detected by the verifiability mechanisms. According to the Post, this vulnerability had already been pointed out two years earlier by Swiss researchers of BFH but still persisted. They said regretting that the technology partner, Scytl, which is responsible for the source code, had not made the correction in full earlier.⁷¹ The trapdoor was found in the new version of the system (for +50% of the electorate) which has never been used so the vulnerability couldn’t have been already exploited to falsify a vote. This time, according to the Swiss Post, Scytl rectified the error, in full and immediately.⁷²

A second vulnerability was found that affects individual verifiability. Someone could theoretically invalidate votes without being detected. Individual verifiability is part of the system that has already been used. However, the Swiss Post relativized saying that exploiting this vulnerability would have produced invalid votes which cannot be accepted by the system and would have been noticed.⁷³ The question remains: why was it not detected by certification and other tests?

⁶⁵ Dubuis, E.: Schwachstellen im E-Voting-System der Post entdeckt, <https://www.societybyte.swiss/2019/03/25/schwachstellen-im-e-voting-system-der-post-entdeckt/>

⁶⁶ <https://www.post.ch/en/business/a-z-of-subjects/industry-solutions/swiss-post-e-voting/e-voting-source-code?shortcut=evoting-sourcecode>

⁶⁷ La Poste Suisse, Accord d’accès au code source de la solution de vote électronique, Janvier 2019.

⁶⁸ <https://www.bk.admin.ch/bk/fr/home/documentation/communiqués.msg-id-74508.html>

⁶⁹ See the two Press releases of the federal Council of 12 March and 29 March 2019, resp. <https://www.admin.ch/gov/en/start/documentation/media-releases.msg-id-74307.html> and <https://www.bk.admin.ch/bk/en/home/dokumentation/medienmitteilungen.msg-id-74508.html>

⁷⁰ *Ibid.* Press release of 12 March 2019.

⁷¹ <https://www.post.ch/en/about-us/media/press-releases/2019/error-in-the-source-code-discovered-and-rectified>

⁷² *Ibid.*

⁷³ <https://www.evoting-blog.ch/en/pages/2019/new-finding-in-the-source-code>

The group of researchers noted that their control was limited as they could only examine a very small percentage of the source code documents, enough though to find two critical vulnerabilities. They would not be surprised to find others. They question other controls which proved successful such as cryptographic and symbolic proofs of verifiability properties,⁷⁴ the role of trust assumptions,⁷⁵ and suggest solutions to rectify the trapdoor [17].⁷⁶

Lewis, Pereira and Teague also highlighted the extremely complex structure of the source code (6000 documents). Other researchers also mentioned that in addition to the security issues, namely the fact that the code allowed manipulations that could have gone unnoticed, a quick examination of the source code revealed other problems, namely: the code is not clear; the documentation does not comply with the standards; all building blocks (code) must be individually configured (by the Post or cantons) which makes it prone to errors; the documentation must not be cited which makes it impossible for researchers to inform and discuss about errors. This last condition clearly does not comply with the VEeS requirements on publication of the source code.⁷⁷

Eventually, the Swiss Post decided to temporarily suspend e-voting and not provide the service to the cantons for the vote of 19 May. It informed it will correct the source code and have it reviewed again by independent experts.⁷⁸ The federal Chancellery invited the Swiss Post to review its security related procedures. It decided to re-examine the certification and agreement procedures.⁷⁹ On 26 June 2019 the federal Council mandated the federal Chancellery to amend the general conditions for future trials.⁸⁰

4 Lessons learned and questions

The publication the of source code and the PIT were meant to confirm an already certified system and help discover potential errors that certification and other tests could not detect. Instead, examination of the code has shown that certification and other controls had failed to notice some critical vulnerabilities in both individual and

⁷⁴ <https://decryptage.be/2019/03/svote/>

⁷⁵ *Ibid.* See also the Berner Fachhochschule experts' conclusions in <https://www.societybyte.swiss/2019/03/25/schwachstellen-im-e-voting-system-der-post-entdeckt/>

⁷⁶ See also fn. 74

⁷⁷ Kolly, M.-J. based on a discussion with Stiller, B. and Killer, Ch. of the Zurich University: Der Quellcode des E-Voting-Systems ist problematisch, und das hat nicht nur mit Sicherheit zu tun. NZZ, 12 March 2019, <https://www.nzz.ch/schweiz/e-voting-der-quellcode-ist-undurchsichtig-sagen-experten-ld.1461406>

⁷⁸ Swiss Post press release of 29 March 2019: <https://www.post.ch/en/about-us/company/media/press-releases/2019/swiss-post-temporarily-suspends-its-e-voting-system>. See also Scytl press release of 1st April: <https://www.scytl.com/en/statement-related-to-the-recent-decision-to-place-evoting-temporarily-on-hold-in-switzerland/>

⁷⁹ Fn. 69

⁸⁰ Fn. 28

universal verifiability. No complete evaluation of the experience has been published so far. There will certainly be important lessons and conclusions that will be drawn on the technical side. E-voting supporters hope that this will make the system/s more secure. The experience also raises more fundamental legal and policy questions.

4.1 Controls

The Swiss regulation on internet voting is an advanced example of designing internet voting requirements to achieve end-to-end verifiable systems in conformity with good practice on security and transparency. The practical implementation shows that, despite good will and the important means dedicated to this, we have not yet obtained an end-to-end verifiable system free of errors. Of course, experts note that these errors would not be there had state-of-the-art solutions been used [e.g. 3, 13, 17].⁸¹ Yet, experts are puzzled by the fact that other cryptographic proofs and controls (other than certification) failed to notice the vulnerabilities.⁸² And they also question one of the basic building blocks of verifiability as practiced here – the trust assumptions.⁸³ This raises a first question of principle. The lay person considers end-to-end verifiability as the way to verify the result of the election given the impossibility to certify that a system, as implemented during e-voting, can be considered 100% secure. If the control of the end-to-end verifiability solution and its implementation presents difficulties similar to those related to controlling the system itself, is end-to-end verifiability a good solution?

Second, the regulation requires state-of-the-art solutions on all aspects related to security, including its control. A constructive dialogue has taken place between cantons, the Confederation and technical research, who have been actively involved in designing and evaluating the two systems. Yet, introducing state-of-the-art solutions in a timely manner is very challenging, as shown by this experience. Certification procedures for end-to-end verifiable solutions were designed end 2013 and allegedly respected good practice at that time. According to researchers, at least by the end of 2018, it became clear that procedures should be redesigned. However, this has not yet been done and the certification of the Swiss Post system was conducted according to the 2013 regulation. According to researchers of the Berner Fachhochschule, this, among others, explains why the system got the certification, despite the flaws.⁸⁴ This shows that, although the Swiss federal internet voting regulation is built in a cascade structure which allows the federal Chancellery to rapidly adapt VELeS to take into account technical developments, it still takes some time to adapt the regulation and the processes. This is unavoidable. Regulation cannot follow technique without delay and there will always be a time lag. In our case this was very detrimental as it allowed

⁸¹ See also fn. 77

⁸² Fn. 74

⁸³ Dubuis, E.: Schwachstellen im E-Voting-System der Post entdeckt, <https://www.societybyte.swiss/2019/03/25/schwachstellen-im-e-voting-system-der-post-entdeckt/>

⁸⁴ *Ibid.*

certification of a flawed system. State-of-the-art that requires adaptations of regulation cannot be implemented without a time lag. Legality seems to weaken the state-of-the-art requirement. Quid?

A third issue is the definition of good practice and state-of-the-art. Researchers for instance pointed out the complexity and quality of the source code of the solution [3]. Certification bodies are expected, according to regulation, to control that the system and security measures are state-of-the-art and respect good practice.⁸⁵ Is this possible at all? Is certification the right instrument for doing this? If yes, is such a certification possible within reasonable time and financial costs? If not, who should define what is state-of-the-art at a given moment and who should check this? Additional questions relate to partial implementation of state-of-the-art and consequences for doing so.

A fourth, crucial issue, relates to the cost of state-of-the-art security. In Switzerland they are covered by the cantons mainly, who organize and conduct elections, including federal ones. As security requirements are determined at the federal level and can vary according to risk evaluation, there is increasingly a friction which may result, as in the case of Geneva, in a decision to abandon internet voting.⁸⁶ The relation between security requirements, which should be uniform and determined at the federal level, and financial means, which come from cantons (states), needs further clarification.

4.2 Transparency

The publication of the source code was the starting point for discovering the most critical vulnerabilities. This highlights the importance of this transparency exercise. The “leaked code” experience shows that restrictions to publication of source code, such as the 45 days silence period, may be unenforceable.

Despite its importance, the publication of the source code and its examination is not a full and systematic control of a system’s security. Researchers indicated that they could only examine a very small fraction of the code. Time and resources fail to do more. Unlike the PIT, the source code examination was not designed as a “bug bounty”, so incentives to detect and report vulnerabilities may be lower. As publication of the source code of systems offering complete verifiability is permanent, conditions may be reconsidered to integrate lessons learned from this first exercise.

⁸⁵ Art. 27/ para. 1 let. b PRO and references to good practice in VELeS.

⁸⁶ Following the cantonal government’s decision of fall 2018 to stop using their own system and outsource the internet voting service to an external provider from the beginning of 2020, on 14 May 2019 the cantonal parliament voted a draft law, which, requires that the design, management and exploitation of an internet voting system remains in public administration’s hands. The Government of Geneva expressed this position – of an internet voting system in public hands – at the consultation on the proposed amendment of PRA (see fn. 27). On 19 June 2019 the cantonal Government decided to advance the deadline and stop using the Geneva system with immediate effect.

4.3 Future directions?

Putting an end to the experimental phase and transforming e-voting into an ordinary voting channel similar to the postal and polling station voting proves to be very challenging. On 26 June 2019, the federal Government decided to delay its introduction as a regular voting channel and reframe the trial phase. Depending on the outcome of the recent popular initiative to introduce a moratorium on e-voting and the interpretation of its requirements, e-voting may even become impossible until its control by the layman is ensured.⁸⁷

Cooperation with research has been crucial in developing second generation systems that offer verifiability and transparency. However, cooperation is important not only in order to develop and evaluate solutions that respect the federal regulation. More should be done already when defining an e-voting policy and regulation. The last decision of the federal Government announced “greater involvement of scientific specialists”. This seems to point into the right direction and should be welcomed.

The contribution of end-to-end verifiability to the security of the internet voting needs a new reflection. Does researchers’ consensus on developing end-to-end verifiable systems need an update? Are elections appropriate playground to try and test end-to-end verifiability? Are there undisputed techniques to achieve “optimum” security?

References

1. Benaloh, J., Rivest, R., Ryan, P. et al. : End-to-end verifiability (2014), <http://arxiv.org/abs/1504.03778>
2. Federal Chancellery Ordinance on Electronic Voting (VEleS), RS 161.116, <https://www.admin.ch/opc/en/classified-compilation/20132343/index.html>
3. James Lewis, S., Pereira, O. and Teague, V.: Trapdoor commitments in the SwissPost e-voting shuffle proof, <https://people.eng.unimelb.edu.au/vjteague/SwissVote>
4. Federal Chancellery, Catalogue des exigences à remplir pour recourir au vote électronique lors de l’élection du Conseil national en 2019. Version 5 April 2018, https://www.bk.admin.ch/dam/bk/fr/dokumente/pore/Anforderungskatalog%20NRW%202019.pdf.download.pdf/Catalogue_des_exigences_ECN_2019_FR.pdf
5. Federal Ordinance on Political Rights (PRO), RS 161.11, <https://www.admin.ch/opc/fr/classified-compilation/19780105/index.html>
6. Federal Act on Political Rights (PRA), RS 161.1, in <https://www.admin.ch/opc/en/classified-compilation/19760323/index.html>
7. Federal Council, « Rapport sur le vote électronique. Chances, risques et faisabilité » of 9 January 2002, FF 2002 612 (2002). We refer to it as “first report”.
8. Federal Council, « Rapport sur les projets pilotes en matière de vote électronique » of 31 May 2006, FF 2006 5205 (2006). We refer to it as “second report”.
9. Federal Council, « Rapport du Conseil fédéral sur le vote électronique. Evaluation de la mise en place du vote électronique (2006–2012) et bases de développement » of 14 June 2013, FF 2013 4519 (2013). We refer to it as “third report”.

⁸⁷ For the time being however the task of verifying the security of internet voting can be conducted by specialists.

10. Driza Maurer, A.: Internet voting and federalism: the Swiss case. In: Barrat i Esteve, Jordi (Coord.) El voto electrónico y sus dimensiones jurídicas: entre la ingenua complacencia y el rechazo precipitado, Iustel, pp. 261-288, Madrid (2016)
11. Braun, N.: E-Voting: Switzerland's Projects and their Legal Framework in a European Context. In: Prosser, A. and Krimmer, R. (Eds.) Electronic Voting in Europe. Technology, Law, Politics and Society, Lecture Notes in Informatics (LNI) - Proceedings Series of the Gesellschaft für Informatik (GI), Volume P-47, pp. 43-52 (2004)
12. Driza Maurer, A.: Ten Years Council of Europe Rec(2004)11. Lessons learned and outlook. In: Krimmer, R., Volkamer, M. (eds) Proceedings of Electronic Voting 2014 (EVOTE2014), pp. 111–120, TUT Press, Tallinn (2014)
13. Dubuis, E., Haenni, R., Koenig, R.: Konzept und Implikationen eines Verifizierbaren Vote Électronique Systems (im Auftrag der Schweizerischen Bundeskanzlei), <https://www.bk.admin.ch/bk/de/home/politische-rechte/e-voting/berichte-und-studien.html> (see “Von der Bundeskanzlei in Auftrag gegebene Studien”, “Konzept Berner Fachhochschule”) (2012).
14. Federal Chancellery : Vote électronique : publication du code source. Rapport explicatif sur la modification de l'ordonnance de la ChF sur le vote électronique (VEleS), du 30 mai 2018, <https://www.bk.admin.ch/bk/fr/home/droits-politiques/groupe-experts-vote-electronique/criteres-pour-les-essais.html> (see “Adaptation des dispositions légales 2018”, “Rapport explicatif OVotE”).
15. Driza Maurer, A.: E-voting source code publication: a good practice becomes a legal requirement. In: Jusletter IT 26. September 2018
16. Puiggali, J., Rodriguez-Peréz, A.: Designing a national framework for online voting and meeting its requirements: the Swiss experience. In Krimmer et al. (eds) E-Vote-ID 2018 Proceedings, pp. 82-97, TUT Press, Tallinn (2018)
17. Haenni, R.: Swiss Post Public Intrusion Test Undetectable Attack Against Vote Integrity and Secrecy (2019), <https://e-voting.bfh.ch/publications/2019/>
18. Haenni, R.: Swiss Post Public Intrusion Test: Generating Random Group Elements (Best Practice) (2019), <https://e-voting.bfh.ch/publications/2019/>

Models and Schemes

Security models for everlasting privacy

Panagiotis Grontas^[0000–0001–7584–0643], Aris Pagourtzis^[0000–0002–6220–3722],
and Alexandros Zacharakis^[0000–0002–1686–8029]

School of Electrical and Computer Engineering
National Technical University of Athens
pgrontas@corelab.ntua.gr, pagour@cs.ntua.gr, azach@corelab.ntua.gr

Abstract. We propose security models for everlasting privacy, which is a property that protects the content of the votes cast in electronic elections against future and powerful adversaries. Initially everlasting privacy was treated synonymously with information theoretic privacy and did not take advantage of the information available to the adversary and his behavior during or after the election. More recent works provided relaxations of the concept, considering the view of a future adversary as well. We integrate all these approaches into game-based definitions. This allows us to contrast the two main proposals to achieve everlasting privacy, namely perfectly hiding commitment schemes and anonymous channels.

Keywords: Electronic voting, everlasting privacy, perfectly hiding commitment schemes, blind signatures, anonymous channels

1 Introduction

Electronic voting is not simply a digital analogue for traditional elections. It aims to improve the voting process by formally defining, analyzing and seeking to satisfy difficult and conflicting security properties.

Verifiability aims to assure candidates and voters that all votes have been considered and incorporated into the result. Its close relation with the integrity of the election process and the acceptance of its output, makes verifiability a very important property, extensively studied [8] and implemented in many protocols under computational assumptions or unconditionally [1,16].

Privacy hides the choice of a voter from the talliers, other voters or external agents in order to free her from external pressure and enable her to express her true will. Verifiability without privacy makes no sense. If one assumes that the contents of all votes are publicly known and linked to individuals, then they can in effect be dictated by external agents applying emotional, personal, social and economic pressures. As a result, one cannot be sure that a vote represents the true will of a voter, as the voter could have yielded to these external forces. Thus, the vote cast would not be the one that was intended. In that sense, it would not differ that much from a vote altered by a malicious entity, as is the case with the verifiability threat model.

Privacy has been studied in many variations, in relation to the capabilities of an adversary and its duration. A first level of privacy protections aims to guard against passive adversaries that want to learn the behavior of a particular voter (subset). This has been implemented in two ways: by hiding the contents of the vote or by disassociating the voter identity from the ballot. The former is usually achieved using a threshold cryptosystem with homomorphic properties, while for the latter an anonymity primitive such as mixnets [6] or blind signatures [7] is applied to the communication channel between the voter and the election authorities. The actual level of privacy offered depends on the implementation. Homomorphic cryptosystems and mixnets usually provide computational and trust guarantees, as it is generally assumed that there will be an honest subset of participants that will follow the protocol. This means that they will refrain from opening individual votes but will decrypt only the result of the final stage. Blind signatures, on the other hand, can offer information-theoretic protection.

Other stronger types of privacy include *receipt freeness* [3], which protects the voters against ‘themselves’ and discourages vote selling. *Coercion resistance* [14] concerns active adversaries that aim to dictate voter behavior with methods ranging from abstention, to random voting and impersonation. *Perfect ballot secrecy* [15] proposed in the context of boardroom voting schemes, guarantees that knowledge about the partial tally of a subset of the voters can be computed only by a coalition of all the remaining voters. However, in all of these cases the adversary is computationally restricted.

Everlasting privacy. A less researched variation of privacy is *everlasting privacy*. Its study, formally initiated by Moran and Naor in [18], focuses on preventing vote leaks from attacks by powerful future adversaries. It is motivated by the observation that in most cases, vote privacy is only protected by a cryptosystem the security of which is based on computational assumptions such as the intractability of the Diffie-Hellman problem [4]. These assumptions, however, may be broken or rendered obsolete in the (not too) distant future, as both the theory and the practice of cryptographic attacks always gets better. This means that votes encrypted with small keys are in danger of being revealed, even without the computational assumption being broken. As famously conjectured by Shamir, at the 2006 RSA Conference cryptographers’ panel, all cryptographic keys used at that time would remain secure for less than thirty years (cf. [18]).

The situation is made worse, because verifiability requires utilizing public evidence generated by the election system. These pieces of data are meant to be widely available and thus it is easy for an adversary to obtain them, even in part. However, one must bear in mind that the adversaries against voting systems are potentially powerful state agencies with enormous budgets and without time constraints. As a result, they have the capability to collect and store large amounts of election related data. Furthermore, as large-scale elections are organized by the government, these agencies can be considered ‘insiders’, having access to even private parts of the election transcript. Finally, these agencies can obtain information exchanged through computer and communication net-

works, both through mass surveillance as well as with the cooperation of the telecommunication companies.

The problem of privacy is exacerbated, as the information concealed in voting does not lose its value, contrary to protected messages in other common cryptographic scenarios. Indeed, one can easily imagine a future authoritarian regime that tries to gather evidence about its subjects based on past democratic elections in cooperation with the state intelligence agency. This evidence might prompt actions ranging from surveillance to questioning and even more severe repercussions. As noted in [18], such dangers constitute an indirect coercion attempt. In fact, since there are many potential coercers the only rational reaction from a voter fearing all possible adverse scenarios is to abstain. Everlasting privacy seeks to protect the secrecy of individual votes in such scenarios.

Our contribution. In this paper, we propose the first game-based definitions for everlasting privacy. Our definitions are generic, which means that they do not consider the cryptographic primitives that will be used in order to achieve this property. This has not been the case so far (cf. Section 2).

More specifically we consider the adversarial capabilities in terms of both data collection and computational power. To model this, we assume two adversaries: The first is contemporary to the election, where he can participate actively (using corrupted voters) and passively (by monitoring communications between the voters and the authorities). He is computationally bounded, though. The second adversary is computationally unbounded but operates (long) after the election is over. The two adversaries can communicate and as a result the future adversary can obtain election transcripts and auxiliary information.

The motivation for this capability stems from the reasonable assumption that there exist powerful entities (e.g. governmental agencies) that might passively hoard election data (among other things as demonstrated by revelations such as Snowden's). It is realistic to assume that a future totalitarian regime will also take control of these agencies (among other things) and have access to their data collection.

By elaborating on the communication options between the present and the future adversary we define two types of everlasting privacy: *strong* and *weak* everlasting privacy, the latter corresponding to the notion of practical everlasting privacy of [2]. Our approach has the added side effect that it associates everlasting privacy with contemporary privacy, which is a relation that, to the best of our knowledge, has not been explored in the literature.

We then apply these threat models of everlasting privacy, against a generic voting scheme. Our analysis focuses particularly on the information gathering capabilities of both adversaries, in relation to the communication channels used. We reason that perfectly hiding commitment schemes do not offer the same levels of protection as anonymous channels, since they cannot hide auxiliary communication information, that can be utilized by a powerful future adversary with insider information.

2 Related work

The term everlasting privacy was coined in [18]. However, there have been previous works that tackle the same problem, even if they do not use the particular name. For instance, in [9] the voter uses the information theoretically hiding Pedersen commitment scheme to commit to the vote. The openings are then secret shared to the authorities using private channels and homomorphically combined. In order to be verifiable, all exchanged data are stored in a Bulletin Board, modelled as a public broadcast channel with memory. Unfortunately, an adversary that hoards its contents can later use his advanced capabilities to break the privacy of the encrypted shares and reconstruct the votes. Interestingly, in this respect, it can be noted that the blind signature-based protocol of [12], achieves this goal as well, if one assumes a *perfectly* anonymous channel (as Theorem 3 of [12] points). The use of this primitive resembles the shuffling of the ballot box contents, which in traditional elections provides a sense of everlasting privacy to the average voter, who as a human is computationally restricted.

The protocols of Moran and Naor [18,19] further elaborate on providing everlasting privacy through perfectly hiding commitment schemes. They propose a concrete voting system that provides universal verifiability, receipt freeness and everlasting privacy. Additionally, they do not require the voter to perform complex calculations which makes their scheme easily usable by humans. In more details, their proposal consists of two authorities that communicate through a private channel and cooperate in order to produce the commitments that the voter selects. To tally the votes, the authorities work together (privately again) to shuffle the commitments and their openings. The latter are encrypted separately using a homomorphic cryptosystem providing computational secrecy. Consequently, one of them can open the perfectly hiding commitments and count the result. Everlasting privacy is achieved under the assumption that the two authorities do not collude, and the commitment openings are not made public and thus available to the future adversary. If only a single authority is honest, then the scheme of Moran and Naor only provides computational privacy, while if both authorities are corrupted then the system provides only correctness. Despite proving the security of their protocol in the UC framework, the threat model for everlasting privacy isn't formally captured. It merely rests on the perfect secrecy of the commitment scheme and an informal description of the adversary's capabilities. Note that in the future an attacker, that functions as an insider, can have an equivalent effect as if the two authorities cooperated.

Subsequent works further elaborate and generalize this technique of splitting voting data into public and private parts, where the private data are never given to the adversary thus achieving a special version of everlasting privacy - towards the public. For instance, in [11] the authors apply this procedure to the Helios [1] voting system, by replacing the exponential ElGamal encryptions with Pedersen commitments that are published to the Bulletin Board. Their opening values are sent to the tallier encrypted through private channels. In [10], a relevant primitive - commitment consistent encryption (CCE) is introduced. It allows the voters to derive commitments from their encrypted votes. These commitments

are then posted to a public Bulletin Board for verifiability purposes. If they are perfectly hiding, then the voting scheme has everlasting privacy. Tallying takes place in parallel using a private Bulletin Board, where the decryption of the result of the homomorphic combination of the votes takes place. They also provide security definitions for the privacy properties of their particular scheme but not for everlasting privacy in general. Furthermore, in [5] this splitting technique is applied to create two synchronized mixnets that operate in parallel, mixing public commitments and private decommitment values respectively.

The central idea in all the works presented so far is that a future adversary might be more powerful in terms of computing power, but he will lack access to data contemporary to the election or private data available to the authorities. This was noted and formalized in [2] with the notion of *practical* everlasting privacy. However, the formalization used the applied pi-calculus and not the more usual indistinguishability cryptographic games. Using automated tools the authors of [2] proved that the protocols of [19] and [11] possess practical everlasting privacy. However, they did not apply their definition to schemes based on blind signatures and anonymous channels. Moreover, the reliance on private channels assumes an external adversary, an adversary, that is, who has a view of the system similar to the view of the voter. This excludes adversaries that cooperate with the election authorities, who in our opinion are more powerful and more likely to be the perpetrators of a future attack.

More recent works revisit the idea of an anonymous channel as a way to add everlasting privacy to voting schemes. In [17], the voter casts an unencrypted choice to the Bulletin Board along with commitments to their voting credential. The use of an anonymous channel and the fact that the voting credential consists of two parts, prevents a future adversary from associating the choice of a voter with her identity. Along the same lines, in [13], the authors add coercion resistance to the classic protocol of [12]. They also solve the ballot stuffing problem of blind signature based systems using a primitive called Publicly Auditable Blind Signatures, an extension of [20], which forces the election authority to verifiably accept or reject ballots for counting. The advantage of their scheme is that it requires no private channels between voters and authorities as all the election data are found in the Bulletin Board. The blindness of the signatures along with the use of an anonymous channel facilitates everlasting privacy.

3 Voting system syntax

We build our definitions on an abstract election scheme that incorporates ideas from many proposals in the literature in order to be as generic as possible. It is associated with three parameters, the security parameter λ , the number of voters n and the number of possible choices m . The election scheme is controlled by an Election Authority \mathcal{EA} , which is stateful and its state is updated in every step of the protocol. In the description that follows we omit this update functionality for simplicity.

We assume the existence of a publicly accessible Bulletin Board where all the election related data is stored. We refer to the current transcript of the Bulletin Board as \mathcal{BB} and we assume that whenever it is used, it contains all the data already written to it. We note that publicly available information such as parameters and public keys are always appended to the public transcript and thus, the \mathcal{BB} would suffice as the public input in the definitions of the scheme. However, we explicitly include such parameters in order to make the algorithms' and protocols' definitions clearer. When we would like to refer to the Bulletin Board as a functionality and not as a data store we use a method invocation-like syntax and we write $\mathcal{BB}()$. Notationally, we use $:=$ for assignment, $=$ for equality, and \leftarrow for an append operation.

Definition 1. *An election scheme*

$$ES = (\text{Setup}, \text{Register}, \text{SetupElection}, \text{Authorize}, \text{Vote}, \text{Tally}, \text{Verify})$$

is a tuple of algorithms and protocols executed by the election authority \mathcal{EA} , the Bulletin Board \mathcal{BB} and the set of voters $\mathcal{V} = \{\mathcal{V}_1, \dots, \mathcal{V}_n\}$ parametrized by $\lambda, n, m \in \mathbb{N}$ such that:

- $(\text{params}_{\mathcal{EA}}, \text{sk}_{\mathcal{EA}}, \text{pk}_{\mathcal{EA}}) := \text{Setup}(1^\lambda)$
 Setup is an algorithm executed by the \mathcal{EA} which on input 1^λ outputs public parameters of the ES and a key pair of the \mathcal{EA} ($\text{sk}_{\mathcal{EA}}, \text{pk}_{\mathcal{EA}}$). The Bulletin Board transcript \mathcal{BB} is appended with $(\text{params}, \text{pk}_{\mathcal{EA}})$.
- $(\text{pk}_i, (\text{sk}_i, \text{pk}_i)) := \text{Register}(\mathcal{EA}(\text{sk}_{\mathcal{EA}}), \mathcal{V}_i(), i)$
 Register is a protocol executed between a voter \mathcal{V}_i and the \mathcal{EA} . The common input is the voter id, $i \in \{1, \dots, n\}$ and the output is a voter public key pk_i (available to both parties) and a secret key sk_i as private output of the voter. The values (i, pk_i) are appended to the \mathcal{BB} . We must stress here, that is not obligatory for voters to have a key pair. However, its existence will enable our generic voting scheme to model protocols like [14], that utilize voter credentials for remote voting and coercion resistance.
- $(I, C) := \text{SetupElection}(\text{sk}_{\mathcal{EA}}, n, m, \text{params}, \text{Election-information})$
The \mathcal{EA} with input its secret key $\text{sk}_{\mathcal{EA}}$, the number of voters n , the number of choices m and additional election information (e.g. duration) outputs the set of the eligible voters for the election $I \subseteq \{1, \dots, n\}$ and the candidate slate C which contains encodings of the choices. The tuple of lists (I, C) is posted to the \mathcal{BB} .
- $(\perp, (\text{b}_i, \pi_{\text{b}_i})) := \text{Authorize}(\mathcal{EA}(\text{sk}_{\mathcal{EA}}), \mathcal{V}_i(c_i, \text{sk}_i), \text{params}, \text{pk}_{\mathcal{EA}}, \text{pk}_i, I, C, \mathcal{BB})$
 Authorize is a protocol executed between the \mathcal{EA} and a voter \mathcal{V}_i . The private input of the \mathcal{EA} is its secret key $\text{sk}_{\mathcal{EA}}$ and the private input of the voter \mathcal{V}_i is her choice of candidate $c_i \in C$ and her secret key sk_i . The public input consists of the system parameters, the corresponding public keys $\text{pk}_{\mathcal{EA}}, \text{pk}_i$, the set of eligible voters I , the candidate slate C and the contents of the \mathcal{BB} . The protocol outputs the ballot b_i , which is a transformation (i.e. encryption) of c_i and a proof π_{b_i} of the correctness of this transformation, usually a Non

Interactive Zero Knowledge Proof Of Knowledge. The election authority receives no output from this functionality. We again assume that the protocol transcript is appended to the \mathcal{BB} .

- $\mathcal{BB} \leftarrow \text{Vote}(\mathcal{BB}(), \mathcal{V}_i(b_i, \pi_{b_i}))$
 Vote is a protocol executed between the voter \mathcal{V}_i and the Bulletin Board \mathcal{BB} . The voter \mathcal{V}_i essentially appends the authorized ballot b_i to the election transcript.
- $(\mathbf{T}, \pi_{\mathbf{T}}) := \text{Tally}(\text{sk}_{\mathcal{EA}}, \text{params}, C, \mathcal{BB})$
 Tally is an algorithm executed by the election authority with input the secret key of the \mathcal{EA} , the parameters of the scheme params , the candidate slate C and the transcript \mathcal{BB} of the Bulletin Board and outputs the election tally \mathbf{T} and a proof $\pi_{\mathbf{T}}$. The output is appended to the Bulletin Board \mathcal{BB} .
- $\{0, 1\} = \text{Verify}(\mathbf{T}, \text{params}, \text{pk}_{\mathcal{EA}}, \mathcal{BB}, C, I, b_i, \pi_{b_i}, \pi_{\mathbf{T}})$
 Verify is an algorithm executed by any interested party (voters or public interest organizations) with input the election tally \mathbf{T} , the parameters of the scheme params , the public key of the \mathcal{EA} $\text{pk}_{\mathcal{EA}}$, the contents of the Bulletin Board \mathcal{BB} , the candidate slate C , the set of eligible voters for the election I , the authorized ballot b and the two proofs $\pi_{\mathbf{T}}, \pi_{b_i}$. The output is a bit representing the result of the election verification. Verify can indeed be executed by any interested party using all the ballots, for universal verifiability purposes, since all inputs can be found in the \mathcal{BB} .

4 Everlasting privacy formalization

We now formally define a voting's system properties regarding privacy. For this reason we consider an adversary, who can corrupt voters and use them with the aim to learn what the honest voters voted. We examine privacy from two aspects: The first concerns 'normal' privacy, which models the protection that voters require during or shortly after the elections. The second applies to the 'everlasting' variation of privacy and models how the voters will be protected (long) after voting has finished. Previously these two definitions were examined independently in the voting literature. However, we note that these properties are intertwined, as an adversary might be motivated to participate in an election, gather evidence by exploiting the voting system and the corrupt voters and possibly use these pieces of information later in time when various constraints might not hold.

More specifically, our adversary is assumed to have the following capabilities:

- He can actively participate in the elections, corrupt voters and collect all data generated by the voting system. During these interventions he is assumed to have computational constraints, as his first goal is to break the privacy of the honest voters during the original election timeframe.
- In the future, he can passively (as there will be no voting taking place) examine the election transcript and extract information about the voters' choices. This adversary is modelled as having unbounded computational capabilities, reflecting the fact that in the future the computational assumptions

that protect the votes might not apply due to technological improvements. This future adversary might or might not utilize only the publicly available election information, thus performing either an insider or an outsider attack as discussed in Section 2.

We consider all these cases in our definitions, by assuming a pair of algorithms $(\mathcal{A}, \mathcal{A}')$ where \mathcal{A} is a PPT algorithm and \mathcal{A}' is computationally unbounded. The former participates actively in the election by corrupting voters and the latter looks at the election transcript and (possibly) the information gathered by \mathcal{A} denoted by $view_{\mathcal{A}}$.

Privacy. The privacy game is a variation of the one presented in [16]. We assume that \mathcal{A} is stateful and its state is updated whenever he performs some action in the game. We complete the notation introduced in Section 3 with the use of the symbol \leftarrow to denote the output of an algorithm, and \Leftarrow for information interchange using a communication channel. Every such communication *leaks* miscellaneous data that are not essential to the protocol but can be used by the adversary to break the system. Such data include network addresses, timestamping information and more. We denote by Aux such miscellaneous data and stress that they will be included in the view of the adversary $view_{\mathcal{A}}$. To denote the execution of one of the functionalities f defined in Section 3 by an entity \mathcal{E} with parameters $params$ we use the following notation: $\mathcal{E}(params, f)$.

In the privacy game (Algorithm 1) the challenger \mathcal{C} takes the role of the \mathcal{EA} , the \mathcal{BB} and the honest voters. It flips a coin and executes the **Setup** functionality. After appending its output to the \mathcal{BB} it interacts with each voter in order to complete the **Register** protocol. Subsequently \mathcal{A} executes the **SetupElection** functionality and dynamically decides which voters to corrupt. If voter i is corrupted, then the challenger presents to \mathcal{A} the private key sk_i and gives full control to him. The challenger retains control of the honest voters. The adversary schedules concurrent executions of the **Authorize** and **Vote** functionalities for all voters in the most favorable manner to him. If a voter is corrupted, \mathcal{A} executes these functionalities in her place using a choice $c_{i,\mathcal{A}}$ of his own. If a voter is honest, then \mathcal{C} plays her role, receives 2 selections $c_0, c_1 \in C$ and provides to \mathcal{A} the results of **Authorize**, **Vote** as well as the auxiliary information (such as network traffic, timestamps etc.). The vote cast by \mathcal{C} on behalf of the honest voters is defined by a coin flip. When all voters have finished executions of their protocols, \mathcal{C} executes the **Tally** functionality and announces the result. \mathcal{A} then receives the full contents of the \mathcal{BB} and auxiliary information Aux and tries to guess the bit b . Note that the adversary has full access to the \mathcal{BB} during the game and as a result he can retrieve its contents at will and not only when he is challenged to guess.

Definition 2. A voting scheme Π is private if for every PPT algorithm \mathcal{A} there exists a negligible function μ such that for every $n, m \in \mathbb{Z}$ it holds that

$$\Pr[Priv_{\mathcal{A}, \Pi}(1^\lambda, n, m) = 1] \leq \frac{1}{2} + \mu(\lambda)$$

Algorithm 1: Privacy Game $\text{Priv}_{\mathcal{A},\Pi}(1^\lambda, n, m)$

```
Input :  $1^\lambda, n, m$ 
Output:  $result \in \{0, 1\}$ 

/* Challenger selects random bit and executes setup */
1  $b \leftarrow \mathcal{C}(\{0, 1\})$ 
2  $(\text{params}_{\mathcal{E}\mathcal{A}}, \text{sk}_{\mathcal{E}\mathcal{A}}, \text{pk}_{\mathcal{E}\mathcal{A}}) \leftarrow \mathcal{C}(1^\lambda, \mathbf{Setup})$ 
/* Challenger registers voters */
3 for  $i \in [n]$  do
4 |  $(\text{sk}_i, \text{pk}_i, \text{Aux}) \leftarrow \mathcal{C}(\mathcal{E}\mathcal{A}(\text{sk}_{\mathcal{E}\mathcal{A}}), \mathcal{V}_i, \mathbf{Register})$ 
5 end
/* Adversary setups election */
6  $(I, C) \leftarrow \mathcal{A}(\text{params}_{\mathcal{E}\mathcal{A}}, n, m, \text{pk}_{\mathcal{E}\mathcal{A}}, \{\text{pk}_i\}_{i \in [n]}, \mathbf{SetupElection})$ 
/* Voters perform authorization in the order selected by the adversary */
7 for  $i \in I$  do
/* Adversary chooses voters to corrupt */
8 if  $\mathcal{A}(i, \mathbf{corrupt}) = 1$  then
9 |  $V_c \leftarrow \{i\}$ 
/* Adversary performs Authorize for corrupted voters */
10 |  $(b_i, \pi_{b_i}, \text{Aux}) \leftarrow \mathcal{A}(\mathcal{E}\mathcal{A}(\text{sk}_{\mathcal{E}\mathcal{A}}), \mathcal{A}(c_i, \mathcal{A}, \text{sk}_i), \text{params}, \text{pk}_{\mathcal{E}\mathcal{A}}, \text{pk}_i, I, C, \mathcal{B}\mathcal{B}, \mathbf{Authorize})$ 
11 else
/* The adversary presents two choices and challenger performs Authorize */
12 |  $(c_0, c_1) \leftarrow \mathcal{A}()$ 
13 |  $(b_i, \pi_{b_i}, \text{Aux}) \leftarrow \mathcal{C}(\mathcal{E}\mathcal{A}(\text{sk}_{\mathcal{E}\mathcal{A}}), \mathcal{V}_i(c_b, \text{sk}_i), \text{params}, \text{pk}_{\mathcal{E}\mathcal{A}}, \text{pk}_i, I, C, \mathcal{B}\mathcal{B}, \mathbf{Authorize})$ 
14 | end
15 end
16  $V_h := I \setminus V_c$ 
/* Challenger votes for the set of honest voters in an arbitrary order */
17 for  $i \in V_h$  do
18 |  $(\mathcal{B}\mathcal{B}, \text{Aux}) \leftarrow \mathcal{C}(\mathcal{B}\mathcal{B}(), \mathcal{V}_i(b_i), \mathbf{Vote})$ 
19 end
/* Adversary votes for the set of corrupted voters. */
20 for  $i \in V_c$  do
21 |  $(\mathcal{B}\mathcal{B}, \text{Aux}) \leftarrow \mathcal{A}(\mathcal{B}\mathcal{B}(), \mathcal{A}(b_i), \mathbf{Vote})$ 
22 end
/* Tally is executed by the challenger */
23  $(\mathbf{T}, \pi_{\mathbf{T}}) \leftarrow \mathcal{C}(\text{sk}_{\mathcal{E}\mathcal{A}}, \text{params}, C, \mathcal{B}\mathcal{B}, \mathbf{Tally})$ 
/* Define partial tallies for honest voters against  $c_0, c_1$  */
24  $T_0 := T_{\mathcal{B}\mathcal{B}V_h}^{c_0}$ 
25  $T_1 := T_{\mathcal{B}\mathcal{B}V_h}^{c_1}$ 
26  $b' \leftarrow \mathcal{A}(T_0, T_1, \text{params}, \mathcal{B}\mathcal{B}, \text{Aux}, \mathbf{guess})$ 
27 if  $T_0 = T_1$  and  $b = b'$  then
28 | return 1
29 else
30 | return 0
31 end
```

Algorithm 2: StrongEverPriv _{\mathcal{A}', Π} ($1^\lambda, n, m$)	Strong	Everlasting	Privacy	Game
<pre> Input : $1^\lambda, n, m$ Output: $result \in \{0, 1\}$ /* Challenger selects random bit */ 1 $b \leftarrow_R \{0, 1\}$ /* \mathcal{A}' initialises \mathcal{A} with honest voter choices and corruption strategy */ 2 $\mathcal{A} \leftarrow \mathcal{A}'(c_0, c_1, V_c)$ /* Challenger executes the election system against \mathcal{A} */ 3 $(\mathcal{BB}, \text{Aux}) \leftarrow \mathcal{C}(1^\lambda, n, m, c_b, \Pi, \mathcal{A})$ /* tallies for honest voters against c_0, c_1 */ 4 $T_0 := T_{\mathcal{BB}, V_h}^{c_0}$ 5 $T_1 := T_{\mathcal{BB}, V_h}^{c_1}$ 6 $b' \leftarrow \mathcal{A}'(T_0, T_1, \text{params}, \mathcal{BB}, \text{view}_{\mathcal{A}}, \text{guess})$ 7 if $T_0 = T_1$ and $b = b'$ then 8 return 1 9 else 10 return 0 11 end </pre>				

Everlasting Privacy. For the everlasting privacy property, we define two games in order to capture the differences in the strategy and knowledge of the future adversary. In both games the adversary \mathcal{A}' is unbounded and invokes the election system that is controlled by the challenger. The difference is that in the StrongEverPriv game, \mathcal{A}' collaborates with the computationally constrained adversary \mathcal{A} , receiving his full state, and utilizes his view including all the auxiliary data he has collected. On the other hand, in the weak everlasting privacy game \mathcal{A}' operates only on the publicly available election data, assumed to be contained in the \mathcal{BB} . In both cases \mathcal{A}' tries to guess the result of the coin flip b .

Definition 3. A voting scheme Π has the strong everlasting privacy property if for every pair of algorithms $\mathcal{A}, \mathcal{A}'$, where \mathcal{A} is PPT, there exists a negligible function μ such that for every $n, m \in \mathbb{Z}$ it holds that

$$\Pr[\text{StrongEverPriv}_{\mathcal{A}, \mathcal{A}', \Pi}(1^\lambda, n, m) = 1] \leq \frac{1}{2} + \mu(\lambda)$$

Definition 4. A voting scheme Π has the weak everlasting privacy property if for every algorithm \mathcal{A}' there exists a negligible function μ such that for every $n, m \in \mathbb{Z}$ it holds that

$$\Pr[\text{WeakEverPriv}_{\mathcal{A}', \Pi}(1^\lambda, n, m) = 1] \leq \frac{1}{2} + \mu(\lambda)$$

Algorithm 3:	Weak	Everlasting	Privacy	Game
WeakEverPriv $_{\mathcal{A}', \Pi}(1^\lambda, n, m)$				
Input : $1^\lambda, n, m$				
Output: $result \in \{0, 1\}$				
/* Challenger selects random bit */				
1	$b \leftarrow_R \{0, 1\}$			*/
/* Selection of honest voter choices */				
2	$(c_0, c_1) \leftarrow \mathcal{A}'(1^\lambda, n, m)$			*/
/* Challenger executes the election system using c_b */				
3	$(\mathcal{BB}, \text{Aux}) \leftarrow \mathcal{C}(1^\lambda, n, m, c_b, \Pi)$			*/
/* tallies for honest voters against c_0, c_1 */				
4	$T_0 := T_{\mathcal{BB}V_h}^{c_0}$			
5	$T_1 := T_{\mathcal{BB}V_h}^{c_1}$			
6	$b' \leftarrow \mathcal{A}'(T_0, T_1, \text{params}, \mathcal{BB}, \text{guess})$			
7	if $T_0 = T_1$ and $b = b'$ then			
8	return 1			
9	else			
10	return 0			
11	end			

5 Analysis

Having formalized the desired security notions, we now discuss the necessary conditions to satisfy them. In particular, we focus on the data interchanged during the execution of various functionalities.

In Algorithm 1 (line 4) the **Register** functionality generates the voter credentials. We assume that these have private and public parts. All voting systems include a similar functionality, mostly using traditional (i.e. not electronic means). In most cases it does not produce specialized credentials for the voters, except in the case of voting systems based on the JCJ coercion resistance framework [14]. Such systems impose the strictest of requirements for this initial communication between the voter and the authorities, i.e. an untappable channel. The reason for this is that the private data ought to be out of reach for the adversary in case - the coercer - so that the voter can deny a purported private key and successfully apply a coercion resistance strategy. The inconvenience imposed by the untappable channel, is mitigated by the fact that it takes place only once and is later applied to many elections. However, such a channel is not necessary for everlasting privacy.

The more interesting parts of the election system are the execution of the **Authorize** and **Vote** functionalities in Algorithm 1 (lines 13 and 20 respectively). Note that in many systems these functionalities are integrated, as the authorization is assumed to take place ‘outside’ of the election system, in a manner similar to the registration. In any case, the voter will interact with the election system and post her ballot to the \mathcal{BB} using (a variation of) these functionalities. The

output of this process will be the election ballot in encrypted form and auxiliary information (such as network information, timestamps etc.), both of which may be of interest to the future adversary. Its unlimited computational power will enable the decryption of the ballot and in turn the linking of the contents of the ballot to the voter.

A system providing everlasting privacy must act on this transfer of information and prevent the data leak. Two options have been proposed on this matter. The first one is to hide the choice of the voter, inside the ballot b_i , using an information theoretically hiding commitment scheme. This means that the decommitment values must be somehow exchanged in order to tally the result. Usually this takes place using private channels with the \mathcal{EA} . In our model, this approach does not provide any advantage and is doomed not to possess strong everlasting privacy. This occurs because the strong adversary in Algorithm 2 has access to the private channels and the auxiliary information. More specifically, this approach essentially repeats the posting of the ballot, since an encrypted ballot is essentially the same as a commitment opening, exchanged through a private channel. In both cases the auxiliary information Aux provided to \mathcal{A}' will enable linking the voter to the vote. More specifically in the strong everlasting privacy game (Algorithm 2) the view of \mathcal{A} contains both network identifying information valid during the elections as well as the decommitment values. As a result, \mathcal{A}' can win the game by recovering the votes of the honest voters and guessing the bit. This is not the case with the weak everlasting privacy game (Algorithm 3) as \mathcal{A}' views only the publicly available information in the \mathcal{BB} . As a result, he might have access to the vote, but he lacks information about the voter identity.

Another alternative is the use of an anonymous channel. This has the immediate effect that the auxiliary information Aux is in effect nullified, as the network addresses are hidden. Note that the anonymous channel must not only hide identity information, but the casting order as well. If this is not the case than an adversary that schedules casting to his advantage can break the secrecy of the vote. All he needs to do is have the corrupt voters cast first as shown in Algorithm 1. Subsequently as each honest voter posts her ballot, he can decrypt the last vote cast (using his unlimited computational power) and learn how she voted. There are two ways to thwart this attack: Firstly, there can be an explicit separation of the **Authorize** and **Vote** functionalities. In the beginning, all voters authorize their ballot. After this phase has finished, they cast their votes. This in effect uses the authorization phase to build an anonymity set that hides the order of the votes cast in the voting phase. Alternatively, the same effect can be achieved using an anonymous channel that hides the order of its input messages apart from their source and origin. We consider such an assumption within the range of functionalities provided by such a primitive. Finally, one must note that an anonymous channel can be combined with commitment-based schemes, to nullify the data that leaks during the use of the various communication channels.

One might argue that an unconditionally anonymous channel is required, in order to thwart the information-theoretically powerful future adversary from reversing the anonymity. In our view, however, this is not the case. An anonymous channel might not be in (full) control of the future adversary. It might be distributed, operated (in part) by a non-governmental organization and it might even transcend nation boundaries. As a result, the future totalitarian regime represented by the unbounded adversary, will not have access to the anonymous network in its entirety and subsequently there is no need for it to be based on information theoretically secure primitives. Such a system can successfully succeed in thwarting the adversary from guessing the honest voters' choices as it will not be able to associate them with the real identities and thus achieve both strong and weak everlasting privacy.

6 Conclusion

In this paper we introduced security models for everlasting privacy. Our adversary has the strongest capabilities ever defined in the literature as he is both active during the election by collecting data, as well as in the future where he can break the cryptographic schemes used. Based on this we defined two models of everlasting privacy. Our novel contribution was the modelling of the adversarial capabilities both in terms of computational power and in terms of information context. Using this model, we reasoned that a system based on commitments opened using through channels cannot provide the strongest sense of everlasting privacy, as an adversary with internal knowledge (such as a governmental agency) will have access to both the decommitments and network information. The use of an *independent* anonymous channel, however, will be able to thwart such an attempt. While such a channel is not currently practical, especially at a large scale, our model indicates that research for everlasting privacy will be assisted by its existence. Anonymous channels have the added benefit that they resemble the way traditional elections work and as a result such a system will be more accessible to the voter. Therefore, our paper gives one more reason to continue the research in this direction. In future work, we plan to refine our model and to provide more formal evidence based on concrete instantiations of voting systems and anonymous channels.

Acknowledgements. The authors would like to thank the anonymous reviewers for their helpful comments and suggestions.

References

1. Adida, B.: Helios: web-based open-audit voting. In: Proceedings of the 17th conference on Security symposium. pp. 335–348. USENIX Association (2008), <http://dl.acm.org/citation.cfm?id=1496711.1496734>

2. Arapinis, M., Cortier, V., Kremer, S., Ryan, M.: Practical everlasting privacy. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). vol. 7796 LNCS, pp. 21–40 (2013). https://doi.org/10.1007/978-3-642-36830-1_2, http://link.springer.com/10.1007/978-3-642-36830-1_2
3. Benaloh, J., Tuinstra, D.: Receipt-free secret-ballot elections (extended abstract). In: Proceedings of the twenty-sixth annual ACM symposium on Theory of computing - STOC '94. pp. 544–553. ACM Press, New York, New York, USA (1994). <https://doi.org/10.1145/195058.195407>, <http://portal.acm.org/citation.cfm?doid=195058.195407>
4. Boneh, D.: The Decision Diffie-Hellman problem. In: Buhler, J. (ed.) Algorithmic Number Theory, Third International Symposium, ANTS-III, Portland, Oregon, USA, June 21-25, 1998, Proceedings, Lecture Notes in Computer Science, vol. 1423, pp. 48–63. Springer (2006). <https://doi.org/10.1007/bfb0054851>, <https://doi.org/10.1007/BFb0054851>
5. Buchmann, J., Demirel, D., Van De Graaf, J.: Towards a publicly-verifiable mix-net providing everlasting privacy. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). vol. 7859 LNCS, pp. 197–204 (2013). https://doi.org/10.1007/978-3-642-39884-1_16, http://link.springer.com/10.1007/978-3-642-39884-1_16
6. Chaum, D.: Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Commun. ACM* pp. 84–88 (1981)
7. Chaum, D.: Blind Signatures for Untraceable Payments (1982). https://doi.org/10.1007/978-1-4757-0602-4_18, http://link.springer.com/10.1007/978-1-4757-0602-4_18
8. Cortier, V., Galindo, D., Küsters, R., Müller, J., Truderung, T.: SoK: Verifiability Notions for E-Voting Protocols. In: IEEE Security and Privacy Symposium. pp. 779–798 (2016)
9. Cramer, R., Franklin, M., Schoenmakers, B., Yung, M.: Multi-Authority Secret-Ballot Elections with Linear Work. pp. 72–83 (1996). https://doi.org/10.1007/3-540-68339-9_7, http://link.springer.com/10.1007/3-540-68339-9_7
10. Cuvelier, É., Pereira, O., Peters, T.: Election verifiability or ballot privacy: Do we need to choose? In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). vol. 8134 LNCS, pp. 481–498 (2013). https://doi.org/10.1007/978-3-642-40203-6_27, http://link.springer.com/10.1007/978-3-642-40203-6_27
11. Demirel, D., Graaf, J.V.D., Araújo, R.: Improving Helios with Everlasting Privacy Towards the Public. EVT/WOTE'12 Proceedings of the 2012 international conference on Electronic Voting Technology/Workshop on Trustworthy Elections (2012)
12. Fujioka, A., Okamoto, T., Ohta, K.: A practical secret voting scheme for large scale elections (1992). https://doi.org/10.1007/3-540-57220-1_66, http://link.springer.com/10.1007/3-540-57220-1_66https://doi.org/10.1007/3-540-57220-1_66
13. Grontas, P., Pagourtzis, A., Zacharakis, A., Zhang, B.: Towards everlasting privacy and efficient coercion resistance in remote electronic voting. In: Zohar, A., Eyal, I., Teague, V., Clark, J., Bracciali, A., Sala, F., Massimiliano, P. (eds.) Financial Cryptography and Data Security (FC 2018). Lecture Notes in Computer Science, vol 10958, pp. 210–231. Springer (2019). https://doi.org/10.1007/978-3-662-58820-8_15, <https://eprint.iacr.org/2018/215.pdf>

14. Juels, A., Catalano, D., Jakobsson, M.: Coercion-resistant electronic elections. In: Atluri, V., di Vimercati, S.D.C., Dingledine, R. (eds.) *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. vol. 6000 LNCS, pp. 37–63. ACM (2005). https://doi.org/10.1007/978-3-642-12980-3_2, <http://doi.acm.org/10.1145/1102199.1102213>
15. Kiayias, A., Yung, M.: Self-tallying elections and perfect ballot secrecy. In: Naccache, D., Paillier, P. (eds.) *Public Key Cryptography, 5th International Workshop on Practice and Theory in Public Key Cryptosystems, PKC 2002, Paris, France, February 12-14, 2002, Proceedings. Lecture Notes in Computer Science*, vol. 2274, pp. 141–158. Springer (2002). https://doi.org/10.1007/3-540-45664-3_10, https://doi.org/10.1007/3-540-45664-3_10
16. Kiayias, A., Zacharias, T., Zhang, B.: End-to-end verifiable elections in the standard model. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. vol. 9057, pp. 468–498 (2015). https://doi.org/10.1007/978-3-662-46803-6_16, http://link.springer.com/10.1007/978-3-662-46803-6_16
17. Locher, P., Haenni, R., Koenig, R.E.: Coercion-resistant internet voting with everlasting privacy. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. vol. 9604 LNCS, pp. 161–175 (2016). https://doi.org/10.1007/978-3-662-53357-4_11, http://link.springer.com/10.1007/978-3-662-53357-4_11
18. Moran, T., Naor, M.: Receipt-Free Universally-Verifiable Voting with Everlasting Privacy. pp. 373–392 (2006). https://doi.org/10.1007/11818175_22, http://link.springer.com/10.1007/11818175_22
19. Moran, T., Naor, M.: Split-ballot voting. *ACM Transactions on Information and System Security* **13**(2), 1–43 (feb 2010). <https://doi.org/10.1145/1698750.1698756>, <http://portal.acm.org/citation.cfm?doid=1698750.1698756>
20. Zacharakis, A., Grontas, P., Pagourtzis, A.: Conditional blind signatures. In: 7th International Conference on Algebraic Informatics (short version) (2017), full version available on: <http://eprint.iacr.org/2017/682>

Improvements in Everlasting Privacy: Efficient and Secure Zero Knowledge Proofs*

Thomas Haines and Clementine Gritti

NTNU, Trondheim, Norway (`thomas.haines`), (`clementine.gritti`)@ntnu.no

Abstract. Verifiable electronic voting promises to ensure the correctness of elections even in the presence of a corrupt authority, while providing strong privacy guarantees. However, few practical systems with end-to-end verifiability are expected to offer long term privacy, let alone guarantee it. Since good guarantees of privacy are essential to the democratic process, good guarantees of everlasting privacy must be a major goal of secure online voting systems. Various currently proposed solutions rely on unusual constructions whose security has not been established. Further, the cost of verifying the zero knowledge proofs of other solutions has only been partially analysed. Our work builds upon Moran and Naor’s solution—and its extensions, applications and generalisations—to present a scheme which is additively homomorphic, efficient to verify, and rests upon well studied assumptions.

Keywords: Voting · Everlasting Privacy · Zero Knowledge Proofs.

1 Introduction

Electronic voting schemes have been studied extensively and ongoing research has developed schemes with increasingly strong privacy and integrity guarantees. However, at present the literature has few solutions which are simultaneously efficient, practical, and ensure the ongoing—also called everlasting—privacy of elections. By practical we mean solutions which are easy to deploy securely. Much of the existing literature relies on trusted setup or complicated recovery procedures which reduce the trustworthiness of the election.

Many schemes have sketched how to do elections with everlasting privacy. The constructions tend to use perfectly hiding commitment schemes and public key encryption; this is made verifiable by use of Zero Knowledge Proofs (ZKPs) for correct encryption and correct shuffling of ballots. At present, one of the most common commitment schemes used is not proven secure [17]. A possible method of mixing has been suggested but the security proof is missing [9]. Further, the

* The authors acknowledge support from the Luxembourg National Research Fund (FNR) and the Research Council of Norway for the joint project SURCVS.

suggested method of mixing is not sufficiently practical. The importance of everlasting privacy has been widely recognised and prior works present constructions with competing efficiency.

We want an electronic voting system with everlasting privacy, which is also efficient to run. We introduce the following mechanisms that will enable us to design such a solution, namely Pedersen commitments [19], Sigma Protocols [8,10] and mix-nets [7]. A Pedersen commitment [19] is an informational-theoretic hiding and computational binding commitment scheme. It provides privacy regardless of the computational power of the adversary but its binding property reduces to the Discrete Logarithm (DLOG) problem. Pedersen commitments are popular in electronic voting schemes because the binding property is only relevant during the course of the election, but privacy should be assured even after the election.

Multiparty computation [25] allows the secure evaluation of a function without leaking anything more about the inputs than can be derived from the result and the inputs previously known to the adversary. ZKPs [14] are a powerful technique which allows proving the correctness of a statement without leaking any other information. The application of both multiparty computation and ZKPs to voting is obvious and commonly mentioned [9,10]. However, the general strategies for both techniques are too computationally intensive in most real elections. Hence there are tailored solutions (such as those we present here) which take advantage of the particularities of elections to construct more efficient solutions.

Sigma Protocols [8,10] are a class of protocols known to be secure under composition. They tend to be more efficient than zero knowledge protocols. A protocol of the correct form is proved to be a Sigma Protocol by showing it satisfies the following properties: *completeness*, capturing that the protocol will succeed when both parties are honest; *special soundness*, referring to the inability of the adversary to generate proofs without knowing a witness; and *honest verifier zero knowledge*, emphasising that the proof leaks negligible information.

Mix-nets were first proposed by Chaum [7], as a way to provide privacy. In the context of verifiable electronic voting mix-nets are also required to be verifiable. This is achieved by proving the correctness of the shuffle using a ZKP, of which two techniques are dominant; namely those of Bayer and Groth [4] and that of Terelius and Wikström [23]. Both techniques are general in nature and tend to be optimised for the particularities of the system in which they are used.

1.1 Related Work

Much of the everlasting privacy literature relies on and builds upon Moran and Naor's work [17], which was modified as an extension to the web-based voting Helios scheme [13]. This kind of extension reduces privacy attacks on the system (from an external adversary) to information theoretic security rather than computational. Hence, no future breakthrough in computation power, mathematics,

or large-scale quantum computers will put the voters’ privacy at risk. Unfortunately, the bulk of this work relies on primitives which are somewhat unusual. Since Moran and Naor, a Pedersen commitment variant is often used but its security appears never to have been rigorously established. Indeed, there is much literature which states that Pedersen commitments and Sigma Protocols are generally required to be defined in a prime order group, which this variant is not, meaning its security should be rigorously established [3,6,20]. We denote, in the paper, the combination of Paillier encryption [18] and Pedersen commitments [19], pioneered by Moran and Naor, as the MN encryption scheme.

Arapinis *et al.* [2] recently showed in ProVerif, an automatic cryptographic protocol verifier, that various constructions achieve everlasting privacy, some of these solutions lose verifiability properties in exchange for everlasting privacy but are highly practical in those situations where these verifiability properties are not important. Cuvelier *et al.* [9] systematised much of the research by showing how certain types of primitives can be securely combined. They also present an elegant scheme called PPATC based on Abe *et al.*’s [1] commitment scheme on bilinear pairings, which they show has efficient encryption on the order of 40 times faster than existing methods. The efficiency is due to the elliptic curves which are more secure relative to their size than problems based on factorisation.

However, Cuvelier *et al.* [9] do not account for the verification complexity. We show that Moran-Naor suggestion of Paillier encryption and Pedersen commitments—refereed as PPATP in [9]—is at least as fast to verify as PPATC when using the Sigma Protocol and mix-net we will detail later. Further, the MN system supports homomorphic tallying where PPATC does not which is a significant advantage in some situations. We note that Cuvelier *et al.* [9] do sketch the same Sigma Protocol for correct encryption in their paper that we later present, but provide no proof. We also note that recent work of Hazay *et al.* [16], has made threshold key generation in Paillier practical as with PPATC.

Many of the existing solutions—except Cuvelier *et al.* [9]—are unsatisfactory in one of two ways. They complicate practical issues, by detecting issues after they have occurred rather than using ZKPs initially. Alternatively, they rely on cut-and-choose based ZKPs rather than Sigma proofs, resulting in an increase in computation and communication of about six orders of magnitude.

There are efficient mix-nets for both Paillier ciphertexts and Pedersen commitments (e.g., Moran and Naor highlight Groth’s mix-net working for Paillier encryption scheme [15]). However, mixing the commitments and ciphertexts separately significantly complicates the election process and weakens security. Cuvelier *et al.* note that the general construction of Wikström [24] can be applied but do not prove the required Sigma Protocol. Further, this construction is significantly slower than the optimised constructions popular in electronic voting.

1.2 Contributions

- We present the Sigma Protocol for re-encryption of the MN cryptosystem; we also provide the proof for this Sigma Protocol and for the protocol for correct encryption [9] of the MN cryptosystem;

- We provide the first proof of security for the existing modified Pedersen commitment of semi-prime order;

- We present an efficient variant of ballot mixing;

- We give an analysis of verification efficiency of MN cryptosystem and compare with PPATC, showing MN is as fast to verify when using the mix-net and Sigma Protocols from above.

When Moran and Naor first introduced the MN cryptosystem they said “although more efficient (zero knowledge) protocols exist for these applications, for the purpose of this paper we concentrate on simplicity and ease of understanding” [17]. Unfortunately in the decade since the follow up work has continued to rely on cut-and-choose [5,13]; and, has found updating the existing zero knowledge work to the requirements of the MN cryptosystem more difficult than Moran and Naor expected. Our contribution finally closes this gap by providing efficient proofs for encryption, re-encryption and shuffling.

1.3 Road Map

In the next section, we provide the notations and definitions useful for the comprehension of the paper. In Section 3, we present our security proof for the modified Pedersen commitment scheme [17]. In Section 4, we describe our new Sigma Protocol for re-encryption, and give the security proofs for the latter as well for the Sigma Protocol for encryption [9]. In Section 5, we depict our verifiable mix-net, improving the efficiency of the general construction proposed in [24]. In Section 6, we analyse and compare the efficiency of our solution with the similar work of Cuvelier *et al.* [9]. We conclude our paper in the last section.

2 Preliminaries and Building Blocks

Due to lack of space, we let the readers refer to [14] for zero knowledge notions, and specifically to [10] for Sigma Protocols, and to [7] for mix-nets.

Notations Natural numbers are denoted by \mathbb{N} and integers by \mathbb{Z} . The ring of integers modulo n is denoted \mathbb{Z}_n , and its multiplicative group \mathbb{Z}_n^* . Let M denote a square matrix of order N from $\mathbb{Z}_n^{N \times N}$. Let \mathbf{v} be a vector of length N from \mathbb{Z}_n^N . Let $\langle \mathbf{v}, \mathbf{v}' \rangle = \sum_{i=1}^N v_i v'_i$ denote the inner product. Given a finite set S , $s \leftarrow_r S$ means a uniformly random assignment of an element in S to the variable s . A Polynomial-Time Algorithm (PPT) is a probabilistic algorithm running in

time polynomial in its input size. A relationship $\mathcal{R}_*(\circ)(\diamond)$ is a subset of the Cartesian product of the sets \circ and \diamond . We denote by $\mathcal{R}_1 \vee \mathcal{R}_2$ the relationship consisting of the pairs $((x_1, x_2), w)$ s.t. $(x_1, w) \in \mathcal{R}_1$ or $(x_2, w) \in \mathcal{R}_2$. Let $\mathcal{R}_1 \wedge \mathcal{R}_2$ be the relationship consisting of the pairs $((x_1, x_2), w)$ s.t. $(x_1, w) \in \mathcal{R}_1$ and $(x_2, w) \in \mathcal{R}_2$.

Discrete Logarithm Assumption Given primes p, q and $n = pq$, where $kn+1$ is also prime, for $k \in \mathbb{N}$. Let \mathbb{G}_n denote the group of order $n \bmod \mathbb{Z}_{kn+1}^*$ and let $\mathbb{G}_p, \mathbb{G}_q$ denote the groups of order p and q respectively $\bmod \mathbb{Z}_{kn+1}^*$. \mathbb{G}_p and \mathbb{G}_q are called *Schnorr groups*. The Discrete Logarithm (DLOG) assumption is believed to hold for the set of Schnorr groups.

Commitment Scheme

Definition 1. A homomorphic commitment scheme Π is a triple of PPT algorithms $(\Pi.\text{Setup}, \Pi.\text{Com}, \Pi.\text{Open})$, s.t.:

- The Setup algorithm for a given group \mathbb{G} defines a set of valid Commit Keys CK from which one is uniformly selected: $CK \in \mathcal{CK} \leftarrow_r \Pi.\text{Setup}(\mathbb{G})$.
- A given Commit Key CK defines a message space \mathcal{M}_{CK} , randomness space \mathcal{R}_{CK} , commitment space \mathcal{C}_{CK} , and opening space \mathcal{D}_{CK} . The Com algorithm takes these as domain and co-domain: $\forall m \in \mathcal{M}_{CK}, \forall r \in \mathcal{R}_{CK}, (c \in \mathcal{C}_{CK}, d \in \mathcal{D}_{CK}) \leftarrow \Pi.\text{Com}_{CK}(m, r)$.
- The Open algorithm takes a commitment $c \in \mathcal{C}_{CK}$ and opening $d \in \mathcal{D}_{CK}$ and returns either a message $m \in \mathcal{M}_{CK}$ or null \perp : $\Pi.\text{Open}_{CK}(c \in \mathcal{C}_{CK}, d \in \mathcal{D}_{CK}) \rightarrow m \in \mathcal{M}_{CK}$ or \perp .

Correctness: $\forall CK \in \mathcal{CK}, \forall m \in \mathcal{M}_{CK}, \forall r \in \mathcal{R}_{CK}$, we have $\Pi.\text{Open}_{CK}(\Pi.\text{Com}_{CK}(m, r)) = m$.

Homomorphism: $\forall CK \in \mathcal{CK}, \forall m_1, m_2 \in \mathcal{M}_{CK}, \forall r_1, r_2 \in \mathcal{R}_{CK}$, we have $\Pi.\text{Com}_{CK}(m_1, r_1) * \Pi.\text{Com}_{CK}(m_2, r_2) = \Pi.\text{Com}_{CK}(m_1 + m_2, r_1 + r_2)$. The homomorphic property implies the ability to re-randomise commitments: let the ReRand algorithm be defined as $\Pi.\text{ReRand}_{CK}(c \in \mathcal{C}_{CK}, r \in \mathcal{R}_{CK}) = c * \Pi.\text{Com}_{CK}(1, r)$.

Definition 2. Perfectly hiding property of a commitment scheme: Given a group \mathbb{G} , a commitment scheme Π is perfectly hiding if for any adversary \mathcal{A} , it holds that $\text{Adv}^{\text{hiding}}(\mathcal{A}, \Pi, \mathbb{G}) = \text{Pr}[\text{Exp}_{\mathcal{A}}^{\text{hiding}-1}(\Pi, \mathbb{G})] - \text{Pr}[\text{Exp}_{\mathcal{A}}^{\text{hiding}-0}(\Pi, \mathbb{G})] = 0$ (Fig. 1).

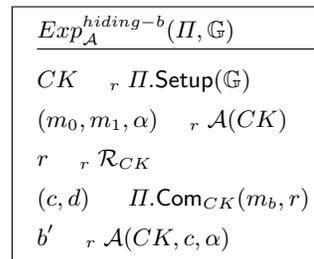


Fig. 1. Hiding experiments

Definition 3. *Binding property of a commitment scheme: Given a group \mathbb{G} , a commitment scheme Π is (t, ϵ) binding if no t -time algorithm \mathcal{A} has $\text{Succ}^{\text{binding}}(\mathcal{A}, \Pi, \mathbb{G}) > \epsilon$ in $\text{Exp}_{\mathcal{A}}^{\text{binding}}(\Pi, \mathbb{G})$ (Fig. 2). For simplicity we will often drop t and ϵ and refer to Π as binding.*

$\text{Exp}_{\mathcal{A}}^{\text{binding}}(\Pi, \mathbb{G})$

$CK \leftarrow_r \Pi.\text{Setup}(\mathbb{G})$
 $(c, d, d') \leftarrow_r \mathcal{A}(CK)$
 $m \leftarrow \Pi.\text{Open}_{CK}(c, d)$
 $m' \leftarrow \Pi.\text{Open}_{CK}(c, d')$
if $m \neq m'$ **return** 1
else return 0

Fig. 2. Binding experiment

Public Key Encryption Scheme

Definition 4. *A homomorphic public key encryption scheme Σ is a triple of PPT algorithms $(\Sigma.\text{KeyGen}, \Sigma.\text{Enc}, \Sigma.\text{Dec})$, s.t.:*

- *The KeyGen algorithm defines a set of valid key pairs (PK, SK) from which one is uniformly selected: $(PK \in \mathcal{PK}, SK \in \mathcal{SK}) \leftarrow_r \Sigma.\text{KeyGen}(1^k)$.*
- *A given public key PK defines a message space \mathcal{M}_{PK} , randomness space \mathcal{R}_{PK} , and ciphertext space \mathcal{C}_{PK} . The Enc algorithm takes these as domain and co-domain: $\forall PK \in \mathcal{PK}, \forall m \in \mathcal{M}_{PK}, \forall r \in \mathcal{R}_{PK}, CT \in \mathcal{C}_{PK} \leftarrow \Sigma.\text{Enc}_{PK}(m, r)$.*
- *The Dec algorithm takes a ciphertext $CT \in \mathcal{C}_{PK}$ and $SK \in \mathcal{SK}$ and returns either a message $m \in \mathcal{M}_{PK}$ or null \perp : $\forall CT \in \mathcal{C}_{PK}, \Sigma.\text{Dec}_{SK}(c) \rightarrow m \in \mathcal{M}_{PK}$ or \perp .*

Correctness: $\forall (PK \in \mathcal{PK}, SK \in \mathcal{SK}) \leftarrow_r \Sigma.\text{KeyGen}(1^k), \forall m \in \mathcal{M}_{PK}, \forall r \in \mathcal{R}_{PK}$, we have $\Sigma.\text{Dec}_{SK}(\Sigma.\text{Enc}_{PK}(m, r)) = m$.

Homomorphism: $\forall PK \in \mathcal{PK}, \forall m_1, m_2 \in \mathcal{M}_{PK}, \forall r_1, r_2 \in \mathcal{R}_{PK}$, we have $\Sigma.\text{Enc}_{PK}(m_1 + m_2, r_1 + r_2) = \Sigma.\text{Enc}_{PK}(m_1, r_1) * \Sigma.\text{Enc}_{PK}(m_2, r_2)$.

We succinctly recall the IND-CPA security concept for a public key encryption scheme as intuitively meaning that the adversary cannot distinguish between the encryption of two known plaintexts.

2.1 Modified Pedersen Commitment Scheme

As we have already noted starting with Moran and Naor [17], Pedersen commitments of semi-prime order have become a significant building block for voting schemes with everlasting privacy. The construction proposed in [17] was to take two safe primes p, q (i.e. to be of the form $2p + 1$ for p prime), let $n = pq$ and work in the subgroup of order n of \mathbb{Z}_{4n+1}^* where $4n + 1$ is also prime.

The modified Pedersen commitment scheme Π is the triple of PPT algorithms $(\Pi.\text{Setup}, \Pi.\text{Com}, \Pi.\text{Open})$, s.t.:

- $CK \leftarrow \Pi.\text{Setup}(\mathbb{G})$ s.t. $CK = \{\mathbb{G}, g, h\}$. Given a group \mathbb{G} of semi-prime order n , let g be any generator of \mathbb{G} and choose $h \leftarrow_r \mathbb{G}$ (with overwhelming probability h will be a generator).
- A given Commit Key $CK = \{\mathbb{G}, g, h\}$ defines the message space $\mathcal{M}_{CK} = \mathbb{Z}_n$, randomness space $\mathcal{R}_{CK} = \mathbb{Z}_n$, commitment space $\mathcal{C}_{CK} = \mathbb{G}_n$, and opening space $\mathcal{D}_{CK} = (\mathbb{Z}_n, \mathbb{Z}_n)$. The $\Pi.\text{Com}_{CK}$ algorithm takes $m \in \mathbb{Z}_n, r \in \mathbb{Z}_n$ and sets $c = g^r h^m$ and $d = (m, r)$.
- The $\Pi.\text{Open}_{CK}$ algorithm takes a commitment $c \in \mathbb{G}_n$ and opening $d \in (m \in \mathbb{Z}_n, r \in \mathbb{Z}_n)$. If $c = g^r h^m$ return m else return \perp .

2.2 A Commitment Consistent Encryption System

The encryption scheme suggested by Moran and Naor [17] is a particular kind of encryption system specialised for everlasting privacy, and commonly used in verifiable electronic voting [12,13]. The standard suggestion, which we describe below, is to use Pedersen commitments of semi-prime order and the generalised Paillier cryptosystem. This notation—while slightly unusual—is useful because it enables the direct application of various existing results, particularly those in the area of mix-nets, as we shall see later. For convenience, we shall refer to this system as the MN cryptosystem.

We now describe MN encryption scheme. Let $\Sigma = (\Sigma.\text{KeyGen}, \Sigma.\text{Enc}, \Sigma.\text{Dec})$ denote a public key encryption scheme. Specifically let $\Sigma.\text{KeyGen}$ be the key generation function of the (generalised) Paillier cryptosystem [18,11] producing $PK = (n)$ and $SK = (d)$, where $n = pq$ is a RSA modulus and d is the lowest common multiple of $p - 1$ and $q - 1$. Choose k s.t. $kn + 1$ is prime, and let g, h be random generators of subgroup of order n in \mathbb{Z}_{kn+1}^* , denoted \mathbb{G}_n . We denote the ciphertext space $\mathcal{C}_{PK} = \mathbb{G}_n \times \mathbb{Z}_{n^2}^* \times \mathbb{Z}_{n^2}^*$, the message space $\mathcal{M}_{PK} = \mathbb{Z}_n$, and the randomness space $\mathcal{R}_{PK} = \mathbb{Z}_n \times \mathbb{Z}_n^* \times \mathbb{Z}_n^*$.

We quickly explain the encryption process. Let $\Sigma.\text{Enc}_{PK}(m \in \mathbb{Z}_n, (r \in \mathbb{Z}_n, r' \in \mathbb{Z}_n^*, r'' \in \mathbb{Z}_n^*))$ produce $CT = (c, ct_1, ct_2) = (g^r h^m \bmod kn + 1, (1 + n)^m r'^m \bmod n^2, (1 + n)^r r''^m \bmod n^2)$. That is we encode the message m in a Pedersen commitment hidden by the randomness r , and we encrypt the opening to this commitment in two Paillier ciphertexts. Let $\Sigma.\text{Dec}_{SK}(CT = (c, ct_1, ct_2))$ be the decryption function. First use the Paillier decryption function to retrieve m, r from ct_1, ct_2 respectively, then if $c = g^r h^m$ the result is m else \perp .

We first make the observation that the Σ scheme is additively homomorphic, that is $\Sigma.\text{Enc}_{PK}(m_0, (r_0, r'_0, r''_0)) * \Sigma.\text{Enc}_{PK}(m_1, (r_1, r'_1, r''_1)) = \Sigma.\text{Enc}_{PK}(m_0 + m_1, (r_0 + r_1, r'_0 * r'_1, r''_0 * r''_1))$. Secondly, that there is a shuffle friendly map [24]: given $CT = (c, ct_1, ct_2)$ and $r = (r_0, r_1, r_2)$, $c' = c * g^{r_0}, ct'_1 = ct_1 * r_1^n, ct'_2 = ct_2 * (1 + n)^{r_0} r_2^n$. We denote this map by $(\phi_{PK}(CT, r) = \mathcal{C}_{PK} \times \mathcal{R}_{PK} \rightarrow \mathcal{C}_{PK})$. The existence of this map is necessary to apply Wikström's general mix-net construction to the cryptosystem [24].

In addition, we preserve the property of Paillier encryption and Pedersen commitments that given a ciphertext $CT = \Sigma.\text{Enc}_{PK}(m_0 \in \mathcal{M}_{pk}, (r, r', r'') \in \mathcal{R}_{pk})$ and a message m_1 it is easy to compute $CT^{m_1} = \Sigma.\text{Enc}_{pk}(m_0 * m_1; (r * m_1, r'^{m_1}, r''^{m_1}))$. In this case the exact effect on the randomness is a combination of multiplication and exponentiation. Lastly, since the Paillier variant we use is the variant of Damgård et al [11], threshold decryption is also available.

3 Security Proof for the Modified Pedersen Commitment Scheme

The sketch of the security proof for the commitment scheme in [17] lacks sufficient detail to be of use in establishing the security of the commitment. Since the group n is not of prime order, given a tuple (m, r, m', r') if $GCD(|m - m'|, n) \neq 1$ and $GCD(|r - r'|, n) \neq 1$ then the sketched reduction to the DLOG problem fails. While it is not particularly surprising that the DLOG problem holds in a group whose order contains a large prime factor, it is important to show that this is indeed true and furthermore does not break any other part of the system. A correct reduction is hence needed. Moreover, we do not require the primes to be safe and thus consider a subgroup of order n of \mathbb{Z}_{kn+1}^* for an integer k . Therefore, the above commitment scheme can be extended to the general case with integers k, n such that $kn + 1$ is prime. We now present the security proof of the generalization of the modified Pedersen commitment scheme.

Proposition 1. *The modified Pedersen commitment scheme Π is a homomorphic perfectly hiding commitment scheme.*

Proof. The correctness of the scheme follows immediately from the definitions of $\Pi.\text{Com}$ and $\Pi.\text{Open}$. The perfect hiding property of the scheme follows in the same way as normal Pedersen commitment schemes: for any two messages m_0, m_1 and commitment c there exist two unique random coins r_0, r_1 s.t. $c = g^{r_0} h^{m_0}$ and $c = g^{r_1} h^{m_1}$, and since the random coins are taken uniformly, the commitment provides no information about which message was committed to.

The key to understanding the next part on the binding property is to recall that for a cyclic group of semi-prime order $n = pq$, there are exactly two non-trivial subgroups: one is of order p and the other q . If we let \mathbb{G} be the subgroup of \mathbb{Z}_{kn+1}^* of order n , where $kn + 1$ is prime, then the two non-trivial subgroups are two Schnorr groups. The reduction we present in the next paragraph reduces the binding property of the modified Pedersen commitment to the DLOG problem in the two Schnorr groups, which we label \mathbb{G}_p and \mathbb{G}_q .

To show that the scheme is binding, we present a reduction in two parts. First, we show that for any t -time adversary \mathcal{A} against the modified Pedersen commitment scheme Π with $\text{Succ}^{\text{binding}}(\mathcal{A}, \Pi, \mathbb{G}) = \epsilon$, we can construct an algorithm which—given a DLOG problem in \mathbb{G}_p , and another in \mathbb{G}_q —outputs the

answer to at least one with probability ϵ . Then having observed against which of the two groups the better success rate is achieved, we construct an adversary against the DLOG problem in that group which succeeds with probability at least $\frac{\epsilon}{2}$. This suffices to show that the binding property of the commitment scheme cannot be broken with probability more than twice that of the DLOG problem in the weakest of the two underlying Schnorr groups \mathbb{G}_p and \mathbb{G}_q .

There exists an efficiently computable isomorphism between the direct product of $\mathbb{G}_p \times \mathbb{G}_q$ and \mathbb{G}_n . The challenger takes the two subgroups of \mathbb{G}_n and a DLOG problem in each. It combines these to construct the commitment key which it gives to the adversary. Since g_p and g_q are generators of their respective groups \mathbb{G}_p and \mathbb{G}_q , if h_p and h_q are random elements (as they are in the DLOG experiment) then this is indistinguishable from the honest run. The successful adversary $\mathcal{A}(\mathbb{G}, g, h)$ outputs $(c, (m, r), (m', r'))$ s.t. $m \neq m'$. If $GCD(|m - m'|, n) = 1$ or $GCD(|r - r'|, n) = 1$ then we extract $\alpha = dlog_g h$ as normal with Pedersen commitments and calculate $dlog_{g_p} h_p = \alpha \bmod p$ and $dlog_{g_q} h_q = \alpha \bmod q$. If this is not the case, then w.l.o.g. $GCD(|r - r'|, n) = GCD(|m - m'|, n) = p$ and hence there exists unique $\delta, \gamma \in \mathbb{Z}_q$ s.t. $\delta p = \alpha \gamma p \bmod n$ and hence $\alpha = \frac{\delta}{\gamma} \bmod q$. By the Chinese remainder theorem $\alpha \bmod q = dlog_{g_q} h_q$ and we successfully answer that.

□

Our solution is not only provably secure (under reasonable assumptions) but also more general with the setting $kn + 1$, with $k \in \mathbb{N}$, rather than $4n + 1$. The homomorphism of the scheme follows immediately from the group properties and the isomorphism of \mathbb{Z}_n and \mathbb{G}_n .

4 Security Proofs for Sigma Protocols

We present two Sigma Protocols, one for correct encryption from [9] and a new protocol for correct re-encryption; we believe that proofs of both Sigma Protocols have never been published before. These proofs allow the realisation of an electronic voting scheme that is secure (compared to without ZKPs) and highly efficient (compared to the cut-and-choose solutions currently in the literature).

4.1 Sigma Protocol for Correct Encryption

The following Sigma Protocol for correct encryption was proposed by Cuvelier *et al.* [9], though they omit the proof. Such protocol is used to prove that given a ciphertext, one knows the inputs and uses them to generate that ciphertext.

Given $CT = (c = g^r h^m \bmod kn + 1, ct_1 = (1 + n)^m r'^n \bmod n^2, ct_2 = (1 + n)^r r''^m \bmod n^2)$, we show that we know $m \in \mathbb{Z}_n$ and $(r \in \mathbb{Z}_n, r' \in \mathbb{Z}_n^*, r'' \in \mathbb{Z}_n^*)$:

- 1) Let t_1, t_2 be random elements in \mathbb{Z}_n and t_3, t_4 be random elements in \mathbb{Z}_n^* . The prover computes $\alpha = g^{t_1} h^{t_2} \bmod kn + 1, \beta = (1+n)^{t_2} t_3^n \bmod n^2, \gamma = (1+n)^{t_1} t_4^n \bmod n^2$ and sends them to the verifier.
- 2) The verifier sends a challenge ξ chosen at random in \mathbb{Z}_n .
- 3) The prover computes $s_1 = t_1 + \xi r \bmod n, s_2 = t_2 + \xi m \bmod n, s_3 = t_3 * r'^{\xi} \bmod n, s_4 = t_4 * r''^{\xi} \bmod n$, and sends these to the verifier.
- 4) The verifier accepts if $\alpha c^{\xi} = g^{s_1} h^{s_2} \bmod kn + 1, \beta ct_1^{\xi} = (1+n)^{s_2} s_3^n \bmod n^2, \gamma ct_2^{\xi} = (1+n)^{s_1} s_4^n \bmod n^2$.
The transcript (with the elements exchanged between the prover and the verifier) is $(\alpha \in \mathbb{G}_n, \beta \in \mathbb{Z}_{n^2}^*, \gamma \in \mathbb{Z}_{n^2}^*, \xi \in \mathbb{Z}_n, s_1 \in \mathbb{Z}_n, s_2 \in \mathbb{Z}_n, s_3 \in \mathbb{Z}_n^*, s_4 \in \mathbb{Z}_n^*)$.

Security Proof

Proposition 2. *The above protocol has perfect completeness, special soundness, and honest verifier zero knowledge and is hence a Sigma Protocol.*

Proof. Completeness follows trivially and is omitted due to lack of space.

Special Soundness Given two accepting transcripts $(\alpha, \beta, \gamma, \xi, s_1, s_2, s_3, s_4)$ and $(\alpha, \beta, \gamma, \xi', s'_1, s'_2, s'_3, s'_4)$, we show that $r = \frac{s_1 - s'_1}{\xi - \xi'}, m = \frac{s_2 - s'_2}{\xi - \xi'}, r' = (s_3/s'_3)^{\frac{1}{\xi - \xi'}}$, $r'' = (s_4/s'_4)^{\frac{1}{\xi - \xi'}}$ must be valid given that two transcripts accept. The difference $\xi - \xi'$ has no inverse with negligible probability. To calculate $r' = (s_3/s'_3)^{\frac{1}{\xi - \xi'}}$, $r'' = (s_4/s'_4)^{\frac{1}{\xi - \xi'}}$, we use our knowledge of the message in ct_1 and ct_2 , extracted from s_1 and s_2 , and the homomorphic property of Paillier encryption to create $ct'_1 = r'^n$ and $ct'_2 = r''^n$. We can directly apply the technique from Damgård *et al.* [11] to extract r' and r'' from the elements s_3, s'_3, s_4, s'_4 .

Honest Verifier Zero Knowledge Consider a transcript $(\alpha, \beta, \gamma, \xi, s_1, s_2, s_3, s_4)$. In the honest run, t_1, t_2 are random elements in \mathbb{Z}_n , t_3, t_4 in \mathbb{Z}_n^* and ξ in \mathbb{Z}_n . To simulate, choose s_1, s_2 from \mathbb{Z}_n , s_3, s_4 from \mathbb{Z}_n^* and ξ at random from \mathbb{Z}_n . Set $\alpha = c_1^{-\xi} g^{s_1} h^{s_2}, \beta = c_2^{-\xi} (1+n)^{s_2} s_3^n, \gamma = c_3^{-\xi} (1+n)^{s_1} s_4^n$, that is a perfect simulation. Moreover, the elements β, γ are uniformly random in the honest run, and the tuple $(\alpha, s_1, s_2, s_3, s_4)$ is uniquely determined by (ξ, β, γ) . In the simulation, the elements s_1, s_2, s_3, s_4 are chosen uniformly at random and consequently β, γ are uniformly at random for fixed elements ξ, c, ct_1, ct_2 . \square

4.2 Sigma Protocol for Correct Re-Encryption

We introduce the following Sigma Protocol for correct re-encryption. It is used to prove that given a pair of ciphertexts, the second is a re-encryption of the first.

Given $CT = (c, ct_1, ct_2), CT' = (c' = c * g^{r_0} \bmod kn + 1, ct'_1 = ct_1 * r_1^n \bmod n^2, ct'_2 = ct_2 * (1+n)^{r_0} r_2^n \bmod n^2)$, we show that we know $(r_0 \in \mathbb{Z}_n, r_1 \in \mathbb{Z}_n^*, r_2 \in \mathbb{Z}_n^*)$:

1) Let t_1 be a random element in \mathbb{Z}_n and t_2, t_3 be random elements in \mathbb{Z}_n^* . The prover computes $\alpha = g^{t_1} \bmod kn + 1, \beta = t_2^n \bmod n^2, \gamma = (1 + n)^{t_1} t_3^n \bmod n^2$ and sends them to the verifier.

2) The verifier sends a challenge ξ chosen at random in \mathbb{Z}_n .

3) The prover computes $s_1 = t_1 + \xi r_0 \bmod n, s_2 = t_2 * r_1^\xi \bmod n, s_3 = t_3 * r_2^\xi \bmod n$, and sends these to the verifier.

4) The verifier accepts if $\alpha(c'/c)^\xi = g^{s_1}, \beta(ct'_1/ct_1)^\xi = s_2^n, \gamma(ct'_2/ct_2)^\xi = (1 + n)^{s_1} s_3^n$.

The transcript (with the elements exchanged between the prover and the verifier) is $(\alpha \in \mathbb{G}_n, \beta \in \mathbb{Z}_{n^2}^*, \gamma \in \mathbb{Z}_{n^2}^*, \xi \in \mathbb{Z}_n, s_1 \in \mathbb{Z}_n, s_2 \in \mathbb{Z}_n, s_3 \in \mathbb{Z}_n^*)$.

Security Proof

Proposition 3. *The above protocol has perfect completeness, special soundness, and honest verifier zero knowledge and is hence a Sigma Protocol for correct re-encryption.*

Proof. Completeness follows trivially and is omitted due to lack of space.

Special Soundness Given two accepting transcripts $(\alpha, \beta, \gamma, \xi, s_1, s_2, s_3)$ and $(\alpha, \beta, \gamma, \xi', s'_1, s'_2, s'_3)$, we show that $r_0 = \frac{s_1 - s'_1}{\xi - \xi'}, r_1 = (s_2/s'_2)^{\frac{1}{\xi - \xi'}}, r_2 = (s_3/s'_3)^{\frac{1}{\xi - \xi'}}$ must be valid given that two transcripts accept. The difference $\xi - \xi'$ has no inverse with negligible probability. To calculate $r' = (s_3/s'_3)^{\frac{1}{\xi - \xi'}}, r'' = (s_4/s'_4)^{\frac{1}{\xi - \xi'}}$, we use our knowledge of the message in ct_1 and ct_2 extracted from s_1 and s_2 , and the homomorphic property of Paillier encryption to create $c'_2 = r''^n$ and $c'_3 = r''^m$. We can directly apply the technique from Damgård *et al.* [11] to extract the randomnesses r', r'' from the elements s_3, s'_3, s_4, s'_4 .

Honest Verifier Zero Knowledge In the honest run, t_1 is chosen at random from \mathbb{Z}_n, t_2, t_3 from \mathbb{Z}_n^* and ξ from \mathbb{Z}_n . To simulate, we instead choose s_1, s_2, s_3, ξ at random and set $\alpha = g^{s_1} (c'_1/c_1)^{-\xi}, \beta = s_2^n (c'_2/c_2)^{-\xi}, \gamma = (1 + n)^{s_1} s_3^n (c'_3/c_3)^{-\xi}$. We get the same distribution in both cases. □

5 A New Efficient Verifiable Mix-Net

Verifiable mixing is an important building block for almost all verifiable voting systems. Given a vector of ciphertexts with known relationships to the voters, mixing allows this link to be broken without allowing ballot modification or substitution.

Wikström's general result [24] shows that verifiable mixing is possible for all cryptosystems on which a homomorphic map exists and an overwhelmingly complete Sigma Protocol is known for re-encryption. However, this generic construction gives an 8-round proof, while a more optimised instance is desirable for

practicality. We can take advantage of special properties from our solution and derive a secure 4-round proof. We illustrate a verifiable ballot mixing process in Fig. 3 with three mixers. We now present our more efficient mixers. While there

Formally we operate two mixes, one on the public bulletin board and on the secret bulletin board. At each step, the election authorities check that the two versions of the Pedersen commitments e_j and c_j match. Our solution is similar to Demirel *et al.* [13], but is actually shown to be secure and far more computationally efficient. The green arrows represent verifiable mixers, the red arrows represent the equality of Pedersen commitments at each stage and the blue arrow represents verifiable decryption.

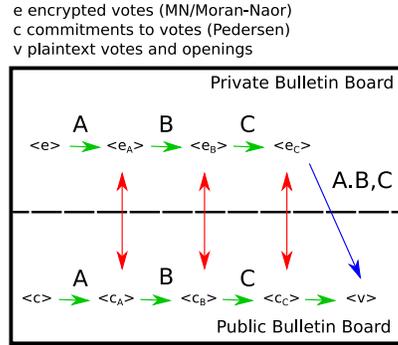


Fig. 3. Mixing with three authorities

are crucial differences (for instance, the composite group order), our optimisations and accompanying proofs are similar to those for the optimised ElGamal version which is presented and proven by Terelius *et al.* [21]. We first detail the mix-net for the public board, see Algorithm 1, and then the mix-net for the private board, see Algorithm 2. We recall that π is permutation function induced by the permutation matrix M and ϕ is the re-encryption map defined in Subsection 2.2. We use $\bar{1}$ to denote the all one vector.

We define \mathcal{R}_{com} to be the relation consisting of pairs of tuples of the form commitment key CK , commitment \mathbf{c} , two distinct messages M, M' and two associated randomness vectors \mathbf{r} and \mathbf{r}' s.t. $\mathbf{c} = \Pi.Com_{CK}(M, \mathbf{r}) = \Pi.Com_{CK}(M', \mathbf{r}')$. We also define \mathcal{R}_π to be the relation consisting of pairs of tuples of the form commitment key CK , commitment \mathbf{c} , message M and associated randomness vector \mathbf{r} s.t. M is a permutation matrix and $\mathbf{r} = \Pi.Com_{CK}(M, \mathbf{r})$. Let $\mathcal{R}_{\phi_{PK}}^{shuf}$ be the relation consisting of pairs of tuples of the form public key PK , two vectors of ciphertexts $\mathbf{CT} = (ct_1, \dots, ct_n)$ and $\mathbf{CT}' = (ct'_1, \dots, ct'_n)$ and a permutation π and randomness vector $\mathbf{r} = (r_1, \dots, r_n)$ such that $ct'_i = \phi_{PK}(ct_{\pi(i)}, r_{\pi(i)})$ for all $i \in [1, N]$. Let $\mathcal{R}_{rerand_{CK}}^{shuf}$ to be the relation consisting of pairs of tuples of the form commit key CK , two commitment vectors $\mathbf{c} = (c_1, \dots, c_n)$ and $\mathbf{c}' = (c'_1, \dots, c'_n)$, a permutation π and randomness vector $\mathbf{r} = (r_1, \dots, r_n)$ such that $c'_i = \Pi.ReRand_{CK}(c_{\pi(i)}, r_{\pi(i)})$.

Algorithm 1: Proof of Shuffle on Public Board

Common Input: Commitment parameters $g, h, h_1, \dots, h_N \in \mathbb{G}_n$, two Pedersen

commitments $\mathbf{e} = (e_1, \dots, e_N) \in \mathbb{G}_n^N$ and

$\mathbf{e}' = (e'_1, \dots, e'_N) \in \mathbb{G}_n^N$, and a permutation matrix commitment

$\mathbf{c} = (c_1, \dots, c_N)$.

Private Input : Permutation matrix $M = (m_{i,j}) \in \mathbb{Z}_n^{N \times N}$, randomness

$\mathbf{r} = (r_1, \dots, r_N) \in \mathbb{Z}_n^N$ s.t. $c_j = g^{r_j} \prod_{i=1}^N h_i^{m_{j,i}}$, and randomness

$\mathbf{r}' = (r'_1, \dots, r'_N) \in \mathbb{Z}_n^N$ s.t. $e'_i = e_{\pi(i)} g^{r_{\pi(i)}}$ for $i, j \in [1, N]$.

1 \mathcal{V} chooses $\mathbf{u} = (u_1, \dots, u_N) \in \mathbb{Z}_n^N$ randomly and hands \mathbf{u} to \mathcal{P} .

2 \mathcal{P} defines $\mathbf{u}' = (u'_1, \dots, u'_N) = M\mathbf{u}$ and then chooses

$\hat{\mathbf{r}} = (\hat{r}_1, \dots, \hat{r}_N), \hat{\mathbf{w}} = (\hat{w}_1, \dots, \hat{w}_N), \mathbf{w}' = (w'_1, \dots, w'_N) \in \mathbb{Z}_n^N$, and

$w_1, w_2, w_3, w_4 \in \mathbb{Z}_n$. \mathcal{P} then defines $\bar{\mathbf{r}} = \langle \bar{1}, \mathbf{r} \rangle$, $\tilde{\mathbf{r}} = \langle \mathbf{r}, \mathbf{u} \rangle$, $\hat{r} = \sum_{i=1}^N \hat{r}_i \prod_{j=i+1}^N u'_j$

and $r' = \langle \mathbf{r}', \mathbf{u} \rangle$. \mathcal{P} hands to \mathcal{V} , where we set $\hat{c}_0 = h$ and $i \in [1, N]$,

$$\hat{c}_i = g^{\hat{r}_i} \hat{c}_{i-1}^{u'_i} \quad t_1 = g^{w_1} \quad t_2 = g^{w_2} \quad t_3 = g^{w_3} \prod_{i=1}^N h_i^{w'_i} \quad t_4 = g^{-w_4} \prod_{i=1}^N (e'_i)^{w'_i} \quad \hat{t}_i = g^{\hat{w}_i} \hat{c}_{i-1}^{w'_i}$$

3 \mathcal{V} chooses a challenge $\xi \in \mathbb{Z}_n$ at random and sends it to \mathcal{P} .

4 \mathcal{P} then responds with:

$$s_1 = w_1 + \xi \cdot \bar{\mathbf{r}} \quad s_2 = w_2 + \xi \cdot \hat{\mathbf{r}} \quad s_3 = w_3 + \xi \cdot \tilde{\mathbf{r}} \quad s_4 = w_4 + \xi \cdot r'$$

$$\hat{s}_i = \hat{w}_i + \xi \cdot \hat{r}_i \quad s'_i = w'_i + \xi \cdot u'_i$$

5 \mathcal{V} accepts if and only if, for $i \in [1, N]$,

$$t_1 = \left(\prod_{i=1}^N c_i / \prod_{i=1}^N h_i \right)^{-\xi} g^{s_1} \quad t_2 = (\hat{c}_N / h \prod_{i=1}^N u_i)^{-\xi} g^{s_2} \quad t_3 = \left(\prod_{i=1}^N c_i^{u_i} \right)^{-\xi} g^{s_3} \prod_{i=1}^N h_i^{s'_i}$$

$$t_4 = \left(\prod_{i=1}^N (e_i)^{u_i} \right)^{-\xi} g^{s_4} \prod_{i=1}^N (e'_i)^{s'_i} \quad \hat{t}_i = \hat{c}_i^{-\xi} g^{\hat{s}_i} \hat{c}_{i-1}^{s'_i}$$

Algorithm 2: Proof of Shuffle on Private Board

Common Input: Commitment parameters $g, h, h_1, \dots, h_N \in \mathbb{G}_n$, two ciphertexts

$\mathbf{e} = (e_1, \dots, e_N) \in \mathcal{C}_{PK}$ and $\mathbf{e}' = (e'_1, \dots, e'_N) \in \mathcal{C}_{PK}$, and a

permutation matrix commitment $\mathbf{c} = (c_1, \dots, c_N)$.

Private Input : Permutation matrix $M = (m_{i,j}) \in \mathbb{Z}_n^{N \times N}$, randomness

$\mathbf{r} = (r_1, \dots, r_N) \in \mathbb{Z}_n^N$ s.t. $c_j = g^{r_j} \prod_{i=1}^N h_i^{m_{j,i}}$, and randomness

$\mathbf{r}' = (r'_1, \dots, r'_N) \in \mathcal{R}_{pk}$ s.t. $e'_i = \phi_{PK}(e_{\pi(i)}, r'_{\pi(i)})$, for

$i, j \in [1, N]$.

1 \mathcal{V} chooses $\mathbf{u} = (u_1, \dots, u_N) \in \mathbb{Z}_n^N$ randomly and hands \mathbf{u} to \mathcal{P} .

2 \mathcal{P} defines $\mathbf{u}' = (u'_1, \dots, u'_N) = M\mathbf{u}$ and then chooses

$\hat{\mathbf{r}} = (\hat{r}_1, \dots, \hat{r}_N), \hat{\mathbf{w}} = (\hat{w}_1, \dots, \hat{w}_N), \mathbf{w}' = (w'_1, \dots, w'_N) \in \mathbb{Z}_n^N$, and $w_1, w_2, w_3, \in \mathbb{Z}_n$

and $w_4 \in \mathcal{R}_{PK}$. \mathcal{P} defines $\bar{\mathbf{r}} = \langle \bar{1}, \mathbf{r} \rangle$, $\tilde{\mathbf{r}} = \langle \mathbf{r}, \mathbf{u} \rangle$, $\hat{r} = \sum_{i=1}^N \hat{r}_i \prod_{j=i+1}^N u'_j$ and

$r' = (\sum_{i=1}^N r'_{i,0} u_i, \prod_{i=1}^N r'_{i,1} u_i, \prod_{i=1}^N r'_{i,2} u_i)$. \mathcal{P} hands to \mathcal{V} , where we set $\hat{c}_0 = h$ and $i \in [1, N]$,

$$\hat{c}_i = g^{\hat{r}_i} \hat{c}_{i-1}^{u'_i} \quad t_1 = g^{w_1} \quad t_2 = g^{w_2} \quad t_3 = g^{w_3} \prod_{i=1}^N h_i^{w'_i}$$

$$t_4 = \Sigma.\text{Enc}_{PK}(1, w_4) \prod_{i=1}^N e_i^{w'_i} \quad \hat{t}_i = g^{\hat{w}_i} \hat{c}_{i-1}^{w'_i}$$

3 \mathcal{V} chooses a challenge $\xi \in \mathbb{Z}_n$ at random and sends it to \mathcal{P} .

4 \mathcal{P} then responds with:

$$s_1 = w_1 + \xi \cdot \bar{\mathbf{r}} \quad s_2 = w_2 + \xi \cdot \hat{\mathbf{r}} \quad s_3 = w_3 + \xi \cdot \tilde{\mathbf{r}} \quad s_4 = w_4 - \xi \cdot r'$$

$$\hat{s}_i = \hat{w}_i + \xi \cdot \hat{r}_i \quad s'_i = w'_i + \xi \cdot u'_i$$

5 \mathcal{V} accepts if and only if, for $i \in [1, N]$,

$$t_1 = \left(\prod_{i=1}^N c_i / \prod_{i=1}^N h_i \right)^{-\xi} g^{s_1} \quad t_2 = (\hat{c}_N / h \prod_{i=1}^N u_i)^{-\xi} g^{s_2} \quad t_3 = \left(\prod_{i=1}^N c_i^{u_i} \right)^{-\xi} g^{s_3} \prod_{i=1}^N h_i^{s'_i}$$

$$t_4 = \left(\prod_{i=1}^N (e_i)^{u_i} \right)^{-\xi} \Sigma.\text{Enc}_{PK}(1, s_4) \prod_{i=1}^N (e'_i)^{s'_i} \quad \hat{t}_i = \hat{c}_i^{-\xi} g^{\hat{s}_i} \hat{c}_{i-1}^{s'_i}$$

Proposition 4. *Algorithm 1 is a perfectly complete, 4-round special soundness, and honest verifier zero knowledge of the relationship $\mathcal{R}_{com} \vee (\mathcal{R}_\pi \wedge \mathcal{R}_{\text{rerand}_{CK}}^{\text{shuf}})$.*

Proposition 5. *Algorithm 2 is a perfectly complete, 4-round special soundness, and honest verifier zero knowledge of the relationship $\mathcal{R}_{com} \vee (\mathcal{R}_\pi \wedge \mathcal{R}_{\phi_{PK}}^{\text{shuf}})$.*

Proof. Due to space limitations we must omit both proofs but they will be present in the full version. Since it is infeasible under the discrete logarithm assumption to find z where $g^z = h$ or to find a pair satisfying \mathcal{R}_{com} , the proposition computationally implies a proof of knowledge of $\mathcal{R}_\pi \wedge \mathcal{R}_{\phi_{PK}}^{\text{shuf}}$. □

6 Comparison and Analysis of Efficiency

We study the efficiency of our solution and compare it with Couvélér *et al.*'s results [9]. In order to accurately confront both schemes, we adopt the similar conventions to Couvélér *et al.* The commitments used by PPATC scheme [9] require an elliptic curve with a type 3 pairing to function. Type 3 pairing is a pairing in which there exist no efficiently computable homomorphism between \mathbb{G}_1 and \mathbb{G}_2 and where the Decisional Diffie-Hellman is hard in both groups. We assume an embedding degree of 16 such that elements of \mathbb{G}_T are of size p^{16} . We, also, associate a unit cost to the multiplication of two 256 bit integers. While Couvélér *et al.* supposed quadratic growth in the length of the operands, we assume $\mathcal{O}(n^{1.5})$, which better reflects that many BigInteger libraries support the optimised multiplication algorithms. We target a security level equivalent to 2048 bits RSA modulus N . We select \mathbb{G}_1 to be taken on \mathbb{F}_p for a 256 bits long prime p and \mathbb{G}_2 to be taken on \mathbb{F}_{p^3} . The size of the target group is then 4096 bits, and for simplicity we take pairing to cost 10 times the effort of a multiplication in \mathbb{G}_1 , this seems to hold for most real implementations.

We count the number of operations in Couvélér *et al.*'s scheme and our solution. Tables 1 and 2 show these numbers for both encryption and opening verification. Let $Exp_{\mathbb{Z}_X^*}$ denote the number of exponentiations in \mathbb{Z}_X^* , and $Mult_{\mathbb{G}_Y}$ the number of multiplications in \mathbb{G}_Y . *Pairing* is defined as the number of pairing operations.

Scheme	$Exp_{\mathbb{Z}_{kn+1}^*}$	$Exp_{\mathbb{Z}_{n2}^*}$	$Mult_{\mathbb{G}_1}$	$Mult_{\mathbb{G}_2}$	Total cost
MN [17]	3.375	4	0	0	1024896 multiplications
PPATC [9]	0	0	9	4	114432 multiplications

Table 1. Total number of operations executed for encryption - Total cost is obtained according to the implementation setting.

Scheme	$Exp_{z_{k_{n+1}}}^*$	$Exp_{z_{n_2}}^*$	$Mult_{G_1}$	$Mult_{G_2}$	Pairing	Total cost
MN [17]	1.125	0	0	0	0	79488 multiplications
PPATC [9]	0	0	1	0	3	119040 multiplications

Table 2. Total number of operations executed for opening verification - Total cost is obtained according to the implementation setting.

While PPATC remains faster for the encryption phase than MN scheme, the latter is 1.5 time faster for the verification phase than PPATC. In regards to mixing, which is of course a very substantial part of the verification cost, we have already shown how an optimised variant of Terelius and Wikström’s approach [22] can be applied to MN cryptosystem.

Cuvelier *et al.* [9] suggested using Terelius and Wikström’s approach as well. However, the efficiency of their general construction is poor compared to the optimised variants (especially when dealing with groups of composite order). The PPATC scheme of Cuvelier *et al.* is a highly elegant construction but contrary to expectations is not more efficient overall than our version of MN scheme [17]. Though, if the voting devices were unusually weak PPATC might still be preferred. In conclusion, while PPATC might still be preferred in some settings, in others where homomorphic properties are desired MN scheme with our optimised ZKPs are of comparable efficiency.

7 Conclusion

Ongoing privacy is fundamental for the proper functioning of elections but significant gaps remained. We fixed several of the outstanding issues. We showed that the modified Pedersen commitment is in fact secure and proved that the Sigma Protocols for correct encryption and correct re-encryption are safe to use. We also provided computational improvements to mixing and examined the feasibility of a secure deployment of our solution. In doing this, we help make everlasting privacy for homomorphic electronic voting a computationally feasible and rigorously secure reality. We show that this approach provides verification efficiency comparable to the most efficient non-homomorphic schemes.

References

1. Abe, M., Haralambiev, K., Ohkubo, M.: Group to group commitments do not shrink. In: EUROCRYPT. LNCS, vol. 7237, pp. 301–317. Springer (2012)
2. Arapinis, M., Cortier, V., Kremer, S., Ryan, M.: Practical everlasting privacy. In: POST. LNCS, vol. 7796, pp. 21–40. Springer (2013)
3. Bangarter, E., Camenisch, J., Krenn, S.: Efficiency limitations for Σ -protocols for group homomorphisms. In: TCC. Lecture Notes in Computer Science, vol. 5978, pp. 553–571. Springer (2010)
4. Bayer, S., Groth, J.: Efficient zero-knowledge argument for correctness of a shuffle. In: EUROCRYPT. LNCS, vol. 7237, pp. 263–280. Springer (2012)

5. Buchmann, J.A., Demirel, D., van de Graaf, J.: Towards a publicly-verifiable mix-net providing everlasting privacy. In: *Financial Cryptography. Lecture Notes in Computer Science*, vol. 7859, pp. 197–204. Springer (2013)
6. Burmester, M.: A remark on the efficiency of identification schemes. In: *EUROCRYPT. LNCS*, vol. 473, pp. 493–495. Springer (1990)
7. Chaum, D.: Untraceable mail, return addresses and digital pseudonyms. *Communications of the ACM* **24**(2), 84–88 (1981)
8. Cramer, R.: Modular design of secure yet practical cryptographic protocols. PhD thesis, Aula der Universiteit (1996)
9. Cuvelier, E., Pereira, O., Peters, T.: Election verifiability or ballot privacy: Do we need to choose? In: *ESORICS*. pp. 481–498. LNCS (2013)
10. Damgård, I.: On Σ -protocols. <http://www.daimi.au.dk/ivan/Sigma.pdf> (2010)
11. Damgård, I., Jurik, M., Nielsen, J.B.: A generalization of paillier’s public-key system with applications to electronic voting. *Int. J. Inf. Sec.* **9**(6), 371–385 (2010)
12. Demirel, D., Henning, M., van de Graaf, J., Ryan, P.Y.A., Buchmann, J.A.: Prêt à voter providing everlasting privacy. In: *VOTE-ID. LNCS*, vol. 7985, pp. 156–175. Springer (2013)
13. Demirel, D., Van De Graaf, J., Araújo, R.: Improving helios with everlasting privacy towards the public. In: *Electronic Voting Technology/Workshop on Trustworthy Elections*. pp. 8–8. USENIX Ass. (2012)
14. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof-systems. In: *Theory of Computing*. pp. 291–304. ACM (1985)
15. Groth, J.: A verifiable secret shuffle of homomorphic encryptions. In: *International Workshop on Public Key Cryptography*. pp. 145–160. Springer (2003)
16. Hazay, C., Mikkelsen, G.L., Rabin, T., Toft, T., Nicolosi, A.A.: Efficient RSA key generation and threshold paillier in the two-party setting. *J. Cryptology* **32**(2), 265–323 (2019)
17. Moran, T., Naor, M.: Split-ballot voting: Everlasting privacy with distributed trust. *ACM Trans. Inf. Syst. Secur.* **13**(2), 16:1–16:43 (2010)
18. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: *EUROCRYPT*. pp. 223–238. Springer (1999)
19. Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: *CRYPTO. LNCS*, vol. 576, pp. 129–140. Springer (1991)
20. Shoup, V.: On the security of a practical identification scheme. *J. Cryptology* **12**(4), 247–260 (1999)
21. Terelius, B.: Some aspects of cryptographic protocols: with applications in electronic voting and digital watermarking. Ph.D. thesis, KTH Royal Institute of Technology (2015)
22. Terelius, B., Wikström, D.: Proofs of restricted shuffles. In: *AFRICACRYPT*. pp. 100–113. Springer (2010)
23. Terelius, B., Wikström, D.: Efficiency limitations of Σ -protocols for group homomorphisms revisited. In: *SCN. Lecture Notes in Computer Science*, vol. 7485, pp. 461–476. Springer (2012)
24. Wikström, D.: A commitment-consistent proof of a shuffle. In: *Information Security and Privacy*. pp. 407–421. Springer (2009)
25. Yao, A.C.: Protocols for secure computations (extended abstract). In: *FOCS*. pp. 160–164. IEEE Computer Society (1982)

Modeling Requirements Conflicts in Secret Ballot Elections

Aaron Wilson^[0000-0002-6712-5284]

Center for Internet Security, New York, United States
aaron.wilson@cisecurity.org

Abstract. Secret ballot elections are a notoriously difficult problem given their unique set of mandated requirements. One of the primary difficulties observed is the number and intensity of conflicting requirements. Unlike other domains, secret ballot elections do not permit the customer or users to select which set of conflicting requirements is preferred when all are necessary for a fair and legitimate election. While many papers mention these conflicts, little work to systematically identify, analyze, and resolve these conflicts has been undertaken. In this paper, we identify and model the conflicts in secret ballot election requirements, including end-to-end verifiability requirements, by using a goal-oriented approach. Four primary goals are decomposed into sub-goals and defined in this paper. Finally, this paper uses goal modeling to graph and discuss the conflicts between sub-goals. As a part of this conflict identification effort, each conflict is explained along with examples of real and theoretical voting solutions which demonstrate the conflict. Future research in conflict analysis and resolution is also proposed.

Keywords: secret ballot elections · electronic voting · requirements engineering · requirements conflicts · goal modeling.

1 Introduction

Secret ballot elections are often performed today with electronic voting systems. An electronic voting (e-voting) system is a system for casting, collecting, and reporting voter intent using predominately electronic means. The Council of Europe defines e-voting as “the use of electronic means to cast and/or count the vote” [1]. While building a computer system to collect votes seems simple at first, creating a system to run binding elections under a rigid set of technical specifications and complex legal frameworks is riddled with difficulties due to conflicting requirements. A system developer or researcher who begins without a full understanding of the goals and requirements for an e-voting system is likely to become frustrated by rework and unexpected sacrifices. A person with a better understanding the requirements is likewise likely to become frustrated by the overwhelming number of conflicts and a lack of a natural starting spot. This is where conflict identification and modeling can help. For our conflict identification and modeling effort, we have chosen to use a goal-oriented approach [23].

Managing conflicts at the goal level provides more freedom for an implementer to find adequate ways to handle conflicts, such as alternative goal refinements and operationalizations, which may result in different system proposals [24]. By applying a goal-oriented approach to identifying e-voting requirement conflicts, this paper hopes to enable eventual conflict analysis and full conflict resolution leading to better and more robust electronic voting solutions for secret ballot elections.

Electronic voting can refer to various implementation models ranging in the level of assistance and reliance on electronic components. [26] identifies at least eight different elections forms ranging from paper-based electronic voting systems to remote electronic voting. In nearly all cases, electronic voting systems are systems of systems that work together to define the election, design the ballots or voting interface, deliver ballots, collect voter intent, tabulate and report results. Moreover, a jurisdiction may deploy more than one electronic voting system or merge various implementations to meet the demand of their diverse voting population. This creates complex interactions which are difficult to evaluate without a proper identification of the high-level goals for an electronic voting system. Relying upon related work in requirements engineering, this paper presents four primary goals as the basis for its analysis. These goals are universal for any secret ballot election and must be met by any viable electronic voting system. These are:

- Voters are afforded a **Secret Ballot**
- **One Person** is afforded and limited to **One Vote**
- Voters are provided **Universal Access** to the voting process
- The voting process is **Transparent and Auditable**

Further descriptions for these goals and a mapping to Volkimer’s principles [25] are provided in Section 4.

2 Our Approach

Our approach builds upon these main goals to identify sub-goals which we then use to identify and classify conflicts. We selected a goal-oriented approach because it provides flexibility for implementation while showing the natural tension between requirements. According to [23], goals are recognized to provide the roots for detecting conflicts among requirements and for eventually resolving them. Modeling the interaction of goals supports a requirements elaboration process that is more accurate and more likely to yield a viable product.

We selected a manual approach to conflict identification, as opposed to automated techniques. A manual approach is where conflicts are identified by a requirements engineer or subject matter expert, which the author is both. An automated approach applies conflict identification techniques using software tools. Automated techniques typically require the use of formal specifications for requirements and any mistake that occurs during the formalization may lead to

incorrect conflict detection. Our specific conflict identification approach analyzes the goals as implemented in real and theoretical e-voting applications and identifies examples of where trade-offs were necessary due to goal conflicts.

Finally, we classify conflicts as either interference or divergence. Interference is defined in [16] as the negative contribution of one requirement on another. Interference causes tradeoffs in satisfying a set of requirements and often means the requirements cannot be satisfied at the same time. Divergence between requirements refers to situations where some combination of circumstances can be found that makes the goals/requirements conflicting [24]. Divergence is a frequent, weaker form of conflict [24].

Managing conflicts is a requirements engineering activity that consists of three main activities: conflict identification, conflict analysis, and conflict resolution. Conflict identification detects the potential conflict. Conflict analysis evaluates and investigates potential conflicts and their tradeoffs. Conflict resolution resolves the potential conflict [16]. Our effort focuses on conflict identification with some discussion of conflict analysis. Conflict identification is visualized using a goal modeling graph. This model can be used to evaluate current implementations, to assist in making trade-off decisions for future implementations, to assist researchers in conducting more focused conflict analysis, and to identify the primary areas in need of conflict resolution. Future research is proposed for conflict analysis and conflict resolution.

The rest of this paper is organized as follows. Section 3 discusses related work. Section 4 provides the decomposed sub-goals stemming from the four primary goals previously introduced. Section 5 reveals the goal model and discusses the various conflicts it depicts. Section 6 presents future work and concludes the paper.

3 Related Work

Identification and evaluation of requirements conflicts in secret ballot elections is lacking in the literature. There is, however, significant work in requirements engineering with a particular focus on requirements specification. Perhaps most influential, Evaluation of Electronic Voting [26] contributes a standardized, consistent, and exhaustive list of requirements for electronic voting systems. These include system requirements, organizational requirements, and assurance requirements. Volkamer provides these requirements for both stand-alone direct recording electronic voting machines and for remote electronic voting machines. Of note, the book provides a detailed review of many sources of e-voting requirements.

In [25], Volkamer, et al. discuss the importance of requirements engineering for e-voting and discuss the development of a new catalog of e-voting requirements and corresponding assessment procedures. The authors introduce 6 principles for which all of their requirements are associated. These are Secret, Free, Equal, Universal, Direct, and Trust. These are used in this paper to help identify goals and sub-goals.

In [15], the authors analyze a subset of e-voting requirements and establish a framework for considering various degrees to which these requirements can be met. The authors consider that in some cases it may be desirable to vary the degree to which these requirements must be met. In identifying multiple satisfaction levels for these requirements, the authors contemplate that it may not be possible to achieve the maximum level for all security requirements in parallel for each election. Consequently, the authors propose potential future work to “identify tradeoffs and incompatibilities among individual levels of different security requirements”. Our effort in this paper makes significant strides to achieve this proposal.

Researchers of secure voting protocols and end-to-end verifiability have commented many times about the inherent conflicts in the requirements. Public Evidence from Secret Ballots [4] states that there is obviously tension between convincing evidence that outcomes are correct and privacy. In this work, Bernhard, et al. discuss many of the factors that make electronic voting a difficult problem. These factors include *no one is trusted*, *the need for evidence*, and *the secret ballot*. Notably, the paper explains the important requirements for secure elections, the solutions already available from research, and then identifies the most important directions for research. The authors present 31 open questions which provide the basis for countless areas of future research, many of which are questions of how to satisfy multiple conflicting requirements simultaneously.

An Overview of End-to-End Verifiable Voting Systems [2] is a survey of attempted implementations of end-to-end verifiable systems. The authors also discuss the current challenges to the deployment of these systems. While not an evaluation of conflicts, this paper identifies areas of conflicts and challenges such as voter privacy and verifiability, verifiability and usability, and vote privacy and accessibility. Similarly, [19] comments that it would be easy to design a system in which the public knew how everyone voted but it is much more difficult to design an end-to-end verifiable system which provides ballot confidentiality, is easy to use, and is accessible to voters with disabilities.

4 Electronic Voting Goals

In this section, we elaborate on our 4 primary goals and how they decompose into sub-goals.

4.1 Secret Ballot Goal

A Secret Ballot election is fundamental to a voting process where voters feel they have the freedom to vote their true conviction. This is achieved by giving voters confidence the process will not divulge their vote and by preventing voter coercion. Our Secret Ballot goal maps to the Secret and Free principles from Volkamer [25]. The Secret Ballot goal is decomposed into the following two sub-goals:

Voter Anonymity Voter’s selections must only ever be known to the voter and anyone he/she willingly shares them with. There must be no means for anyone to obtain the identity of the voter who cast an individual vote or ballot, now and in the future. The concept of everlasting privacy was expressed by Moran and Noar in [17] to define the future portion of this requirement.

Coercion-Resistance The system must protect against vote selling and voter coercion. This is achieved by allowing a given voter to cast their ballot as they truly wish even in the event someone is coercing them to vote a certain way. Voters must also not be able to prove how they voted to anyone even if they wish to do so [4]. Systems achieving coercion-resistance often do so with the property of *receipt-freeness* [5]. This property requires that voters not be allowed to retain possession of anything that can be used as proof to another person of how she voted.

4.2 One Person One Vote Goal

In democratic elections, each voter’s vote has equal weight with every other voter’s vote. This is represented in the Equal principal in Volkamer’s work [25]. This is achieved by treating all votes equally and by ensuring a voter only cast one valid ballot. This is often referred to as the One Person One Vote concept. There are some special elections where entities, such as corporations and property owners, may be given multiple votes. Thus, this goal could alternatively be written as “one person has the number of votes specified by law”. The goal for One Person One Vote is achieved by satisfying the following two sub-goals:

Voter Authenticity The voting process must only be utilized by legitimate and authenticated voters. The identity of the person must be established with proof that the person is who they claim to be.

Ballot Accountability Known in [25] as the Direct principle, the system must record who has received a ballot and prevent attempts to introduce more than one binding ballot per voter. The system must provide a means to audit the number of binding ballots compared to the number of authenticated voters.

4.3 Universal Access Goal

Giving all voters who wish to vote an equal opportunity to vote is critical to a fair and legitimate election. This is achieved by executing a voting process which does not introduce any undue burden on any voter. This concept of Universal Access is codified in various international laws and incorporated by Volkamer as the Universal principal [25]. The following sub-goals are required for Universal Access:

Voter Usability The voting process must be simple and intuitive for voters of various cognitive abilities and cultural sensitivity [1]. Voters must be able to negotiate the process effectively, efficiently, and comfortably [9].

Voter Accessibility The voting process must allow voters with various physical impairments to vote [1]. Accessibility is measured by the degree to which

the system is available to, and usable by, individuals with disabilities [9]. The most common disabilities include those associated with vision, hearing and mobility, as well as cognitive disabilities [9].

Provisional Voting Voters with questionable authenticity or eligibility at the time of voting should be allowed to cast their ballots and prove their eligibility later. Once voted, such ballots must be kept separate from other ballots and are not included in the tabulation until after the voter's eligibility is confirmed [9]. This type of voting is not deployed everywhere.

Transparent and Auditable Goal Transparency and auditability ensure that the public can verify the process was accurate and reliable. This is covered by Volkamer with her Trust principle [25]. Our goal covers the actual accuracy and reliability in the voting process as well as the ability to prove the process was accurate and reliable. This leads to correct results which are accepted by all parties, including and especially the party which lost. To achieve this goal, we lean on end-to-verifiability principles from [19] which are included as sub-goals below:

Cast as Intended Verifiability The voter must be provided the opportunity to verify the voting system correctly interpreted her selections on the ballot [19]. This verifiability is individual and must be done while the voter can still revoke or choose not to cast her ballot.

Recorded as Cast Verifiability The voter must be provided the opportunity to verify that her vote or ballot was received and correctly recorded by the voting system [19]. In addition, the public must be able to verify that each recorded ballot is subject to the recorded as cast check.

Tallied as Recorded Verifiability The public must have the option to verify the vote was correctly tabulated from the same set of ballots which were cast by voters, and that only ballots from eligible voters were included in the final tally [19]. Verifying that the same set of ballots subject to the recorded as cast check is the same as the set of ballots subject to the tallied as record check is referred to as Consistency by Popoveniuc, et al. We incorporate this concept into the Tallied as Recorded verifiability goal for simplicity.

The 3 sub-goals provided above are considered the minimum requirements for an end-to-end verifiable voting system. End-to-end verifiable voting systems are one type of software-independent voting system discussed by Rivest and Wack in [20]. A voting system is software-independent if an undetected change or error in its software cannot cause an undetectable change or error in an election outcome [20].

5 Electronic Voting Goal Conflict Model

Identification of goal conflicts was performed by literature review, review of e-voting implementations, and reasoning with the definitions provided in the previous section. In this section, we discuss each conflict by providing our rationale

and citing implementation examples which demonstrate the conflict. It is important to mention that conflict identification is an ever-changing task. As new implementation ideas surface, the potential for additional conflicts may exist.

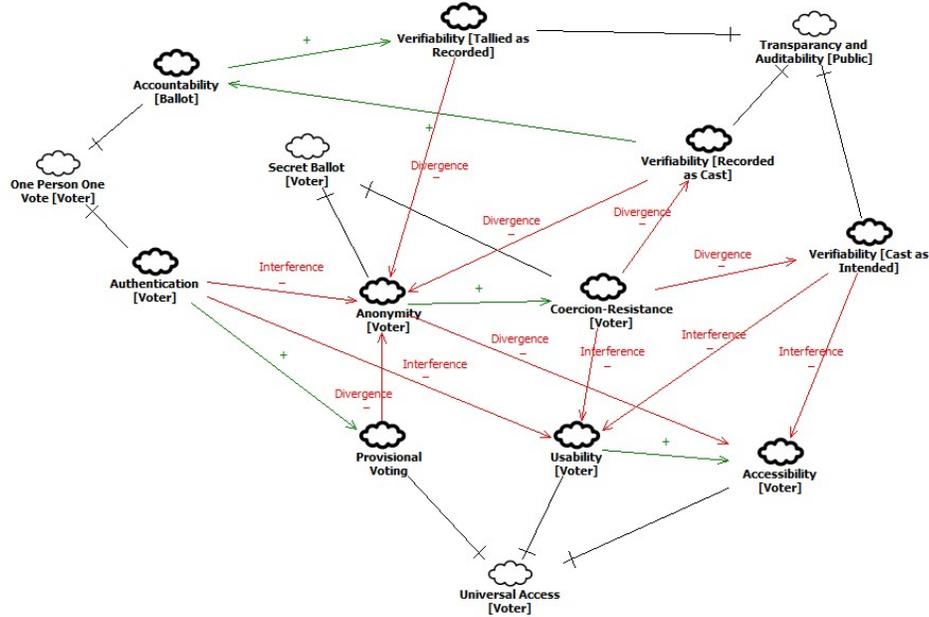


Fig. 1. Secret Ballot Elections Goal Conflict Graph.

Drawing from goal modeling [23], the graph in Figure 1 depicts the main goals, their sub-goals, and the edges between sub-goals. Edges between sub-goals capture the conflict classification as Interference or Divergence. Conflicts (negative relationship) are represented by minus (-) signs and red color. Some positive relationships are depicted with the plus (+) sign. Discussion of positive relationships is outside the scope of this paper.

5.1 Voter Authentication Interferes with Voter Anonymity

Voter anonymity is best accomplished in a system where the identity of the voter is never introduced for any purpose. Thus, the goal of voter anonymity conflicts with voter authentication since voter authentication requires the identity of the voter be known and proven to the system.

Direct or indirect links between the voter’s selections and the voter’s identity could lead to a compromise of the voter’s anonymity. Any voter authentication process will require a record of the voter along with other attributes necessary

to authenticate the person as that voter. This information, along with metadata such as the time of the authentication, is ripe for intentional or unintentional misuse and connection with the voter's selections. This is especially true when voting system artifacts are retained for many months or years depending on local law. The longer the artifacts are retained, the more likely the procedural or technological mitigations put in place to limit the interaction of voter identities and voter selections are to be circumvented.

The most common resolution for mitigating this conflict in electronic voting systems is separation of responsibilities. In this approach, the capturing of voting selections is done in a system which has no knowledge of the voter's identity. The voter authentication is done on a separate system and there is no connection between the systems. This is achievable in scenarios such as in-person voting where procedural controls ensure that only authenticated voters can use the vote capture component of the system. In other scenarios, such as absentee balloting or internet voting, separation of responsibilities is much more difficult, if even possible. In those scenarios, weaker resolutions are deployed such as the two-envelope system used in postal voting. In the two-envelope system, the ballot is inserted inside an inner envelope which is then inserted in an outer envelope. The outer envelope is used for postal delivery. The voter is authenticated using his signature on the inner envelope. Authenticated voters' ballots are removed from the envelope and separated to preserve voter privacy. Digital implementations encrypt and digitally sign the voter's selections and don't allow them to be decrypted in the same system or process where the digital signature is still associated. These systems use decryption mixnets or homomorphic encryption to ensure the plain text vote and voter identities are not associated [2]. Since even the best of these mitigations rely on procedural controls and lack of collusion to protect voter anonymity, we classify this conflict as interference.

5.2 Voter Authentication Interferes with Voter Usability

Determining the best and most usable way to evaluate a voter's identity and eligibility to vote is a matter of intense political and technical debate. The debate centers around the balance of permitting every eligible voter the ability to vote and with requiring a certain level of proof. The more proof required, the lower the usability and the more likely eligible voters are to be incorrectly rejected. Since nearly all voter authentication implementations hurt voter usability to some level, this conflict is classified as interference.

Voter authentication can take many forms. The most common form of voter authentication in American election processes are polling place check-ins with government identification and signature-based authentication for absentee ballots. In remote electronic voting systems, voters use digital forms of authentication such as username/password, requests for personal identifiable information (name, government issued ID numbers, date of birth), smart cards, etc.

Authentication and usability are both essential in the voting process. However, access control requirements and adequate usability are frequently in conflict with each other [6]. According to [6]'s comparative analysis, the authentication

methods which achieve the highest security rating only achieve a moderate usability rating at most. Electronic voting does not present any unique aspects to this conflict, so we intentionally abbreviate the discussion of this conflict.

5.3 Voter Anonymity Diverges with Voter Accessibility

The 2002 Help America Vote Act (HAVA) requires all American election jurisdictions to provide the same opportunity for access and participation, including privacy and independence, equally to all voters [10]. Providing privacy is difficult while providing universal access to voters with various disabilities. For instance, one of the easiest ways to provide voters with disabilities a full voting experience is to provide qualified help. However, doing this will violate that voter's right to a secret ballot. As discussed in [2], adapting voting systems to provide audio/visual aids or human assistance for voters with disabilities may create situations where the voter's candidate choice is revealed to a third party. Likewise, remote voting systems, such as postal or internet voting, are significantly more accessible compared to in-person voting at polling stations but result in a marked deterioration in voter anonymity. We classify this conflict as divergence because the techniques discussed do reduce the risk to voter anonymity to a level accepted by most voters with accessibility needs.

While outside of the scope of this paper, not all requirement conflicts are pair-wise. In reality, there are many complex interactions that occur when 3 or more requirements are combined. One example is the interaction between voter anonymity, voter accessibility, and voter verifiability. For example, Voteegrity [7] was one of the very first voter verifiable, privacy preserving end-to-end verifiable voting schemes. In this scheme, Chaum uses visual cryptography to split an image into multiple shares. Individual shares do not yield any meaningful information about the original image. While in the voting booth, the voter can see her vote by combining two strips of paper. The voter randomly chooses one of the strips as a receipt to take home [2]. This provides voter verification and privacy protection but is not accessible since the use of visual cryptography is not usable by persons with visual impairments. While this limitation is unique to the use of visual cryptography, non-accessible mechanisms are often used to provide voter verification in a private way. Another example is the provision of ballot receipts which can be compared to a public bulletin board later. These are often long, seemingly random, and otherwise meaningless strings which are difficult for voters with cognitive or other impairments to use.

5.4 Coercion Resistance Interferes with Voter Usability

Many of the methods used to achieve coercion resistance ultimately hurt voter usability. This is intuitive because voting in the most straight forward, usable manner allows a coercer the opportunity to simply observe the act of voting. This is especially true in any remote voting method, such as postal or internet voting. Only supervised in-person voting can truly provide coercion resistance without hurting voter usability. In many remote voting proposals, the approach

to coercion resistance is multiple voting. These proposals allow the voter to vote multiple legitimate ballots and take the last one while other schemes allow the voter to cast fake ballots which look real to a coercer. Each of these schemes, however, present usability challenges for voters who are often confused about which ballot was cast or – worse – if they cast a real ballot at all. Therefore, we classify this conflict as interference.

JCJ and its successor Civitas, for example, are known for their high level of coercion resistance but have significant usability issues as detailed in [18]. JCJ/Civitas allows voters to generate fake credentials which are indistinguishable from real ones and can be used to cast dummy votes. Dummy votes show up on the bulletin board for verification but are ultimately not counted. Usability is hurt with this scheme because the voter is forced to manage multiple tokens and know which one is the right one and make sure to use the right one in the free moment(s) she may have away from her coercer. The usability of JCJ/Civitas was improved by [14] by adding smart card support. If the voter enters his correct Pin, the correct ballot is created with his real credentials. Otherwise, a fake credential is used. This could create a scenario where the voter mistypes his PIN, thinks he cast his legitimate ballot, but he did not cast a ballot at all.

5.5 Coercion Resistance Diverges with Cast as Intended Verifiability

Cast as intended verifiability is intended to prove to the voter that the system properly interpreted how they voted. The simplest way to do this is to show the voter the results of the interpretation, or tabulation. This is not possible, however, because it would give the voter proof that she can take away and provide to a coercer. Therefore, providing cast as intended verification must be done in a more complex and less straight forward manner. It must provide proof to the voter in such a way that the voter can't take that proof to others. There are known techniques to achieve both of these goals, such as the Benaloh Challenge [5], so we classify this conflict as divergence. It is critical to mention that many of the implementations which meet both of these goals negatively impact voter usability.

A good example of this is seen in the Norwegian Internet Voting Protocol [11]. In this internet voting system, cast as intended verification is provided by means of distributing voting cards to voters and then transmitting return codes to the voters via text message. If the return codes calculated by the system and provided to the voter via text message match the values on the voting card, the voter can be assured the system interpreted her vote correctly. This practice, however, introduces concerns over vote buying as discussed in [3]. Text messages simplify the task of coercers and vote-buyers because they need only ask the voter to provide the appropriate proof generated by the internet voting system itself [3]. Supporters of this protocol point to the multi-voting support as the means by which the voter can provide proof to the coercer while casting a different, non-coerced ballot. As discussed earlier, multiple voting often leads to usability

issues so while this approach may satisfy this conflict, it moves the concern to a different conflict.

5.6 Coercion Resistance Diverges with Recorded as Cast Verifiability

As first highlighted in Benaloh’s seminal work on receipt-freeness [5], cryptographic election schemes have the potential to suffer from a deficiency which allows the voter to prove to a third party how their vote was cast. Benaloh points out that this deficiency is a result of attempting to produce correctness proofs of the election tally.

Recorded as cast implementations must provide proof to the voter that her ballot was received in its correct form by the election authority. Since this operation is often not immediate, implementations of this verification typically rely on a public bulletin board mechanism. Further, this verification approach is not a simple ballot tracking verification – which would verify the system received a ballot from the voter – but a verification that proves the system received the ballot the voter cast. This distinction is important because it means the system must provide proof to the voter that the contents of the ballot are the same as when the voter cast it. Because of voter coercion concerns, this proof must not actually convey the contents of the voter’s ballot. This forces the implementation to use a more complex and less straight forward mechanism to provide proof to the voter but no one else. Here again, this conflict is satisfiable with current technology so we classify it as divergence with the note that its satisfaction is often at the expense of voter usability.

One example that failed to sufficiently address this conflict was the Rijnland Internet Election System (RIES), used by about 20,000 expatriate voters in the 2006 Dutch parliamentary elections [13]. The system works by publishing a reference table before the elections, including (anonymously) for each voter the hashes of all possible votes, linking those to the candidates. The original votes are only derivable from a secret key handed to the voter. After the elections, a document with all received votes is published. This allows for two important verifications: a voter can verify his/her own vote, including the correspondence to the chosen candidate, and anyone can do an independent calculation of the result of the elections, based on this document and the reference table published before the elections [12]. The voter is provided recorded as cast verification because if the voter’s vote has been registered incorrectly, or not at all, the voter can detect it. This system’s fundamental flaw is that the voter verification scheme can also be used to sell votes. If the voter lets someone else verify their vote, he or she could pay the voter for making the right choice [12].

5.7 Provisional Voting Diverges with Voter Anonymity

Provisional voting is sometimes referred to as conditional or second-chance voting. As these names indicate, provisional voting provides voters an opportunity

to vote if something goes wrong on Election Day under certain conditions. Often, the voter has appeared at the wrong precinct or forgotten a required form of identification. While the ballot is cast on Election Day, the ballot is not counted until the conditions are met. This means the ballot must be held separate from counted ballots and stay associated with the voter herself, so it can be identified. In fully electronic systems, this creates a significant risk to voter anonymity. In paper ballot systems, the voter's identity is only associated with the ballot until the eligibility is determined and then the ballot is separated from the identification. In a digital system, the ballot and voter identity may never be fully separated even though the ballot is counted. Since this concern is limited to digital systems, we classify this conflict as divergence.

One example of this conflict is blockchain voting schemes. To support provisional voting, records of the ballot must be added to the blockchain which are not completely anonymous. There must be a link indirectly or direct back to the voter. There may be another block added to the chain later to denote the ballot's ultimate status, but the original block still contains the ballot and the voter's identity. In one blockchain implementation we are aware of and likely others, the ballot is encrypted and the voter identification is a unique ID related only to the voter through an offline system. These techniques mitigate the conflict, but they do not solve it. This conflict is also present in offline e-voting devices.

5.8 Cast as Intended Verifiability Interferes with Voter Usability

The desire to add cast as intended verifiability to the voting process requires that voters perform some action. This action may be as simple as deciding whether to verify the ballot. As expressed in [2], requiring voters to verify their vote negatively impacts usability by adding extra steps to the process and possibly confusing the voter. Since there is no current implementation that provides voter verification without an extra, undesirable step in the voting process, this conflict is classified as interference.

One example of this conflict is in the Prêt à Voter [21] scheme. The key idea behind the Prêt à Voter approach is to encode the vote using a randomized candidate list. The randomization of the candidate list on each ballot form ensures the secrecy of each vote while providing one half of the ballot as a receipt for cast as intended verification. Because the scheme used in Prêt à Voter alters the printing of the ballot, the voter must also have a way to verify the ballots are well-formed [19]. This check is provided by giving the voter the ability to request two ballots, one to audit and spoil and one to cast. Similarly, in the PunchScan system, the voter can detect maleficence by choosing either the top or the bottom page to keep as her receipt [19]. In both cases, the cast as intended check requires the voter to perform actions well outside the nominal voting experience. Collecting more than one ballot or choosing which receipt to keep contributes to voter confusion and detracts from usability.

It is also important to note that usability is the main factor in whether voters choose to perform the verification process. When given the choice, voters will not choose to perform the verification if it is confusing, inefficient, and seemingly

unimportant. This is also reported in a usability study performed on end-to-end verifiable internet voting systems by the US Overseas Voting Foundation [22]. The less voters who choose to audit the system, the less effective the verification becomes. Therefore, conflicts such as this one not only hurt usability but also negatively impact the verification goal. This dynamic is discussed in [4] and leads to many questions about how to best perform the verification, how to track whether voters performed the verification, and whether the verification can be abstracted away from voters.

5.9 Cast as Intended Verifiability Interferes with Voter Accessibility

Cast as intended schemes rely on generating proof to the voter that the voting system has correctly handled her ballot. Most schemes fall into one of three buckets: independent encryption and compare, Benaloh challenge, and return codes. In each of these cases, the voter is asked to compare values between the voting system and a definitive source which is often on physical media or a separate system. This is very difficult for voters with disabilities, particularly with visual or dexterity limitations. We, therefore, classify this conflict as interference.

A good example of this is code voting schemes first introduced by Chaum in [8]. Code Voting gives each voter a sheet of codes with one for each candidate. Assuming the code sheet is valid, the voter can cast a vote on an untrusted machine by entering the code corresponding to her chosen candidate and waiting to receive the correct confirmation code [4]. This scheme and its successors require voters to manage the interaction between a code sheet and the system and then confirm the confirmation codes and the code sheet. This is not possible for voters with visual impairments.

5.10 Recorded as Cast Verifiability Diverges with Voter Anonymity

In many voting schemes, voter identities are maintained along with the ballot until late into the voting process. This is true in postal voting and many forms of internet voting. This makes it difficult to provide recorded as cast verification while maintaining voter anonymity. Recorded as cast verifiability provides voters the assurance that their ballot has reached the voting authorities without compromise or deletion. This is unique proof that the ballot fidelity remained intact but must be given in such a way that no one can determine from the proof how any voter voted. Each individual proof must be free from any evidence linking the voter to a vote and there must be no way to take the collection of proofs to determine a voter's selections.

Since several schemes have been proposed which address this conflict, we classify it as divergence. These schemes provide the voter only a confirmation code which is tied somehow to the content of the ballot. In Scantegrity, Three-Ballot, and other paper ballot-based systems, the voter takes part of the ballot while the corresponding part is posted to a public bulletin board [2]. The voter can reconstruct the ballot and be assured her ballot was recorded as cast. In

electronic systems like Helios, a probabilistic ballot hash or ballot encryption is posted to the public bulletin board for the voter to review [2].

5.11 Tallied as Recorded Verifiability Diverges with Voter Anonymity

As discussed in the prior conflict, voter identities are often maintained late into the voting process to provide recorded as cast verification and to generally prove the legitimacy of the ballots in the ballot box. This is necessary to show that the ballots being tallied are in fact the same ballots which came from legitimate voters. This creates tension for the actual tabulation process which must strip away voter identification information to preserve privacy while also producing a tabulation result which is verifiably calculated from the original set of ballots. This conflict was first addressed by Chaum using decryption mixnets [2]. These mixnets rely on multiple rounds of decryption each owned by a separate election official. Each mix can be verified but it takes all trustees to accomplish the full decryption. This provides voter anonymity and tallied as recorded verifiability, but voters must trust that all trustees are not colluding.

Several improvements have been proposed which significantly address this problem using homomorphic encryption. Homomorphic encryption is a cryptographic primitive which enables ballot tabulation while still in encrypted form. The result of a homomorphic addition on a set of cipher texts is equivalent to an addition operation performed on the set of plaintexts. Only the result of the addition operation is decrypted, thereby preserving the individual voter's privacy [2]. Scratch and Vote, VoteBox, and STAR-Vote are examples of systems which use homomorphic encryption. We, therefore, classify this conflict as divergence.

6 Future Work and Conclusions

We took a fresh look at the goals for secret ballot elections to properly and fully identify the requirements conflicts which make electronic voting solutions difficult to build. We did this by presenting four primary goals which we decomposed into ten sub-goals. After defining each sub-goal, we used goal-modeling to present the conflicts between the sub-goals in a graphical model. Finally, we discussed the conflicts and included examples of these conflicts materialized in various electronic voting solutions. While many researchers have referenced the inherent conflicts and difficulties with conducting secret ballot elections on electronic voting systems, this is the first work to identify and model these conflicts.

This work is intended to be foundational work in e-voting requirements conflict analysis. Starting with this conflict identification model, we hope to encourage the development of practical and innovative e-voting solutions through further research in conflict analysis and resolution. Our work here will help researchers and engineers understand the difficulties in developing electronic voting solutions for secret ballot elections in a way that does not restrict implementation approaches. Future work will further the conflict analysis provided here

and more precisely identify which conflicts have not been fully resolved in current solutions. This future work will show researchers and practitioners which techniques have more promise than others.

References

1. of Ministers on 14 June 2017, C.: Recommendation cm/ref (2017)5 of the committee of ministers to member states on standards for e-voting. https://search.coe.int/cm/Pages/result_details.aspx?ObjectId=0900001680726f6f
2. Ali, S.T., Murray, J.: An overview of end-to-end verifiable voting systems. CoRR **abs/1605.08554** (2016), <http://arxiv.org/abs/1605.08554>
3. Barrat, J., Chevalier, M., Goldsmith, B., Jandura, D., Turner, J., Sharma, R.: Internet voting and individual verifiability: The norwegian return codes. In: *Electronic Voting* (2012)
4. Benaloh, J., Bernhard, M., Halderman, J.A., Rivest, R.L., Ryan, P.Y.A., Stark, P.B., Teague, V., Vora, P.L., Wallach, D.S.: Public evidence from secret ballots. CoRR **abs/1707.08619** (2017), <http://arxiv.org/abs/1707.08619>
5. Benaloh, J., Tuinstra, D.: Receipt-free secret-ballot elections (extended abstract). In: *Proceedings of the Twenty-sixth Annual ACM Symposium on Theory of Computing*. pp. 544–553. STOC '94, ACM, New York, NY, USA (1994). <https://doi.org/10.1145/195058.195407>, <http://doi.acm.org/10.1145/195058.195407>
6. Braz, C., Robert, J.M.: Security and usability: The case of the user authentication methods. In: *Proceedings of the 18th Conference on L'Interaction Homme-Machine*. pp. 199–203. IHM '06, ACM, New York, NY, USA (2006). <https://doi.org/10.1145/1132736.1132768>, <http://doi.acm.org/10.1145/1132736.1132768>
7. Chaum, D.: Secret-ballot receipts: True voter-verifiable elections. *IEEE Security Privacy* **2**(1), 38–47 (Jan 2004). <https://doi.org/10.1109/MSECP.2004.1264852>
8. Chaum, D.: Surevote: Technical overview. In: *Proceedings of the workshop on trustworthy elections (WOTE 2001)* (2001)
9. Commission, E.A.: 2005 voluntary voting system guidelines
10. 107th Congress (2001): Help america vote act of 2002, <https://www.govtrack.us/congress/bills/107/hr3295>
11. Gjøsteen, K.: The norwegian internet voting protocol. In: Kiayias, A., Lipmaa, H. (eds.) *E-Voting and Identity*. pp. 1–18. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
12. Jacobs, B., Pieters, W.: *Electronic Voting in the Netherlands: From Early Adoption to Early Abolishment*, pp. 121–144. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)
13. Jones, D.W.: Some problems with end-to-end voting. In: *End-to-End Voting Systems Workshop*, Washington DC (2009)
14. Juels, A., Catalano, D., Jakobsson, M.: Coercion-Resistant Electronic Elections, pp. 37–63. Springer Berlin Heidelberg, Berlin, Heidelberg (2010), https://doi.org/10.1007/978-3-642-12980-3_2
15. Langer, L., Schmidt, A., Buchmann, J., Volkamer, M., Stolfik, A.: Towards a framework on the security requirements for electronic voting protocols. In: *2009 First International Workshop on Requirements Engineering for e-Voting Systems*. pp. 61–68 (Aug 2009). <https://doi.org/10.1109/RE-VOTE.2009.9>

16. Mairiza, D., Zowghi, D., Nurmuliani, N.: Managing conflicts among non-functional requirements. In: 12th Australian Workshop on Requirements Engineering (01 2009)
17. Moran, T., Naor, M.: Receipt-free universally-verifiable voting with everlasting privacy. In: Dwork, C. (ed.) *Advances in Cryptology - CRYPTO 2006*. pp. 373–392. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
18. Neumann, S., Volkamer, M.: Civitas and the real world: Problems and solutions from a practical point of view. In: 2012 Seventh International Conference on Availability, Reliability and Security. pp. 180–185 (Aug 2012). <https://doi.org/10.1109/ARES.2012.75>
19. Popoveniuc, S., Kelsey, J., Regenscheid, A., Vora, P.: Performance requirements for end-to-end verifiable elections. In: *Proceedings of the 2010 International Conference on Electronic Voting Technology/Workshop on Trustworthy Elections*. pp. 1–16. EVT/WOTE'10, USENIX Association, Berkeley, CA, USA (2010), <http://dl.acm.org/citation.cfm?id=1924892.1924903>
20. Rivest, R.L.: On the notion of "software independence" in voting systems. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **366**(1881), 3759–3767 (2008). <https://doi.org/10.1098/rsta.2008.0149>, <https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2008.0149>
21. Ryan, P.Y.A., Bismark, D., Heather, J., Schneider, S., Xia, Z.: Prt voter: a voter-verifiable voting system. *IEEE Transactions on Information Forensics and Security* **4**(4), 662–673 (Dec 2009). <https://doi.org/10.1109/TIFS.2009.2033233>
22. U.S. Vote Foundation, G.: The future of voting end-to-end verifiable internet voting usability study (2015), https://www.usvotefoundation.org/sites/default/files/E2EVIV_usability_report.pdf
23. van Lamsweerde, A.: Goal-oriented requirements engineering: a guided tour. In: *Proceedings Fifth IEEE International Symposium on Requirements Engineering*. pp. 249–262 (Aug 2001). <https://doi.org/10.1109/ISRE.2001.948567>
24. van Lamsweerde, A., Darimont, R., Letier, E.: Managing conflicts in goal-driven requirements engineering. *IEEE Transactions on Software Engineering* **24**(11), 908–926 (Nov 1998). <https://doi.org/10.1109/32.730542>
25. Volkamer, M., McGaley, M.: Requirements and evaluation procedures for evoting. In: *The Second International Conference on Availability, Reliability and Security (ARES'07)*. pp. 895–902 (April 2007). <https://doi.org/10.1109/ARES.2007.124>
26. Volkamer, M.: Evaluation of electronic voting: requirements and evaluation procedures to support responsible election authorities, vol. 30. Springer Science & Business Media (2009)

User Experience Design for E-Voting: How mental models align with security mechanisms

Marie-Laure Zollinger¹, Verena Distler¹, Peter B. Rønne¹, Peter Y. A. Ryan¹,
Carine Lallemand², and Vincent Koenig¹

¹ University of Luxembourg, Esch-sur-Alzette, Luxembourg
{marie-laure.zollinger, verena.distler, peter.roenne, peter.ryan,
vincent.koenig}@uni.lu

² Eindhoven University of Technology, Netherlands
c.e.lallemand@tue.nl

Abstract. This paper presents a mobile application for vote-casting and vote-verification based on the Selene e-voting protocol and explains how it was developed and implemented using the User Experience Design process. The resulting interface was tested with 38 participants, and user experience data was collected via questionnaires and semi-structured interviews on user experience and perceived security. Results concerning the impact of displaying security mechanisms on UX were presented in a complementary paper [7]. Here we expand on this analysis by studying the mental models revealed during the interviews and compare them with theoretical security notions. Finally, we propose a list of improvements for designs of future voting protocols.

1 Introduction

Voting protocols are carefully designed to satisfy certain security properties, most importantly Privacy and End-to-End (E2E) Verifiability. Some notable privacy properties are ballot-secrecy, receipt-freeness and coercion-resistance. E2E verifiability is usually separated into the votes being cast-as-intended, recorded-as-cast, and tallied-as-recorded.

E2E-verifiable schemes often require voters to handle encrypted ballots [3, 5, 10]. The Selene e-voting protocol [24] has been designed in order to hide the cryptographic operations from the voter. Instead, each voter is assigned a private tracking number, which lets them verify that their vote has been correctly included in the tally. In the setup phase, a unique tracker number is secretly associated with each voter and cryptographically committed to the bulletin board. At the end of the election, the votes are posted on a public bulletin board along with the associated tracking numbers. To avoid coercion, the voters are notified of their tracking number only after the vote/tracker pairs have been published. This gives coerced voters the opportunity to identify a tracker that points to the coercer's candidate that they can then claim is theirs. The hypothesis is that this mechanism is more intuitive, transparent and easy-to-use than the usual E2E verifiability: where voters should check the encryption of their vote and then presence of this encryption of their vote on the bulletin board.

User tests on voting protocols have shown that schemes that provide security often have usability issues [1, 17, 13]. According to [27], *usability* measures the effectiveness, efficiency, and satisfaction of a software in a specified context of use. Effectiveness is the accuracy and completeness with which the users achieve their goals. Efficiency represents the resources expended for effectiveness. Satisfaction is defined by the comfort and acceptability of use. In the papers [1, 17, 13], the effectiveness of vote casting, that is the ability to cast successfully a vote, has been at most 81.25% [17]. In addition, the meaning of the verification phase is not always well understood, which can lead to voters not performing the verification task or unintentionally aborting the task. Ensuring system usability is further complicated by the fact that elections occur rarely and voters are expected to understand and use a system they are not familiar with.

User Experience is defined as “a person’s perceptions and responses that result from the use or anticipated use of a product, system or service” [28]. It considers emotions, psychological needs and temporal aspects of the interaction between the system and the user, and can measure a person’s perceptions of system qualities such as attractiveness, ease of use and novelty, in addition to usability aspects. To improve the user experience, user-centred methodologies have been developed in order to include the final users in the development of a product [21, 20, 6, 14]. We will describe the process in detail in section 2.

In this paper we present two main contributions, the first is the development of a prototype interface for smartphones for the Selene e-voting protocol, following a user-centered design process [21, 20, 6] called User Experience Design (UXD) Process [14]. We will discuss the impact of our implementation choices on the initial protocol. Then we did a user study on which the primary goal of the interviews was to retrieve insights and interpretation of behavioural data and to complement and triangulate data from the questionnaires, results can be found in [7]. Our second contribution is to study the gaps between voting research and users expectations for a voting system, by exploring the mental models of voters for Privacy and Verifiability expressed in the semi-structured interviews. We define mental models as the concepts in people’s mind that represent their understanding of how things work [20]. This paper describes the first application of the UXD method for app development in the e-voting context and evaluates on its use.

The paper is organized as follows: Section 2 describes the Selene mechanism and details the development of the mobile application following the user-oriented process. Section 3 describes the steps of the user tests that have been done for this study. Section 4 provides an analysis of participants’ interviews and describes the mental models for Privacy and Verifiability. Finally section 5 discusses the implementation, the mental models found and the limitations of the study. We conclude in section 6 by suggesting design improvements and discuss future work.

Related Work The study of mental models is useful to align the system design with the users’ expectation of a system, reducing the possible interaction errors that could lead to additional security (or safety) issues. The subject has received

little attention in voting, we relate our work to the few publications here and discuss them in detail in section 5.

Mental models of verifiability in postal voting and paper voting have been explored by Olembo et al. through a survey conducted in Germany [22]. They suggested breaches in the procedures that could lead to integrity issues and asked participants about different aspects of verifiability. Our approach is different as we do not mention possible security issues in our interviews but we have let the participants express themselves based on the experience of voting with our application (see section 3).

Another paper from Acemyan et al [1] analyzed mental models for three voting schemes which are Helios, Prêt à Voter and Scantegrity II. The experiment aimed to study the participants' mental models through drawings and interviews after using each of the voting systems. The analysis of participants' feedback showed that many participants did not see the E2E-verifiable schemes as being more secure than a standard paper-based voting method. The authors also highlighted that participants tend to focus more on the voting phase, we noticed a lack of understanding for verifiability as described in section 4.

Human factors in security were highlighted by Kulyk and Volkamer in [13]. They extract five concepts including concern and self-efficacy, as we did here: we noticed a lack of concern for verifiability and a lack of self-efficacy (in the sense of knowledge and understanding).

Trust was pointed out by Schneider et al. in [25] as an important factor for participants, as people are aware of potential security issues. Here trust is also an identified mental model of voters.

2 A mobile application for Selene

2.1 Selene mechanism

Protocol overview Most verifiable voting schemes involve voters seeing and handling cryptographic data which can lead to errors or misuses [1]. Selene [24] is an e-voting protocol that has been designed to provide an easier and more intuitive verification procedure for voters. It lets the voters verify that their votes have been included in the tally using a unique tracking number. To protect against coercion threats, i.e. achieve receipt-freeness and coercion mitigation, voters first learn their private tracking number after the votes have been posted. Selene uses ElGamal encryption, that is homomorphic and can act as a commitment scheme. An ElGamal encryption is a pair (α, β) . For a given voter, the tracking number is encrypted using ElGamal and the β -term is published at the beginning of the election. The α -term is kept secret and shared between several entities called Tellers. After the tally has been published, the α -terms are sent to the voter, who can decrypt the tracking number with her key. Full details about the cryptographic mechanisms can be found in the original paper [24].

Voter experience As in Ryan et al. [24] we assume that the voter already has the cryptographic key material needed for the protocol, i.e. we skip the key

setup phase. The voting ceremony without coercion is as follows:

- (1) The voter receives an invitation to vote.
- (2) The voter makes her vote choice in the provided application, encrypts and signs the vote, and sends it to the Election Server.
- (3) (*Optional*) The voter later receives an invitation to visit the bulletin board when votes and tracking numbers are published.
- (4) The voter receives the α -term, and can retrieve the tracking number to verify her vote.

The third step is optional as it is only necessary if the voter is being coerced. For our implementation, we will assume that no coercion is happening and thus the third step is not available. Moreover, we simplified the fourth step by automating the α retrieval and tracker computation. The detailed methodology deployed during the user tests is described in section 3.

2.2 Application Design

A user-oriented approach We followed a user-centred design methodology, which has originally been described by Norman [21] and then detailed as a design process [6, 20]. In particular, we followed the UXD by Lallemand et al. [14]. The process consists of five steps which are *planning*, *exploration*, followed by an iterative process (shown in figure 1) with *ideation*, *generation*, *evaluation*.

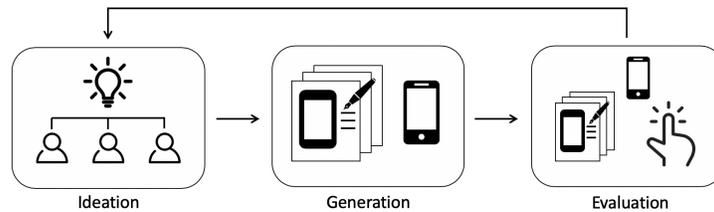


Fig. 1. Iterative process.

The exploration phase includes a collection of user needs, and can be done using various methods, such as the a literature review of previous studies, interviews, focus groups or observations. In our case, we discussed the voting issues mentioned in several papers [1, 2, 8, 9, 11, 17] during meetings with Human-Computer Interaction (HCI) experts who helped us develop and test prototypes of the e-voting application in a user-centred process in close collaboration.

Then we focused on the iterative process: we worked together with the HCI experts for the generation of ideas for the design during group sessions with up to ten group members. We then came up with the concept for a mobile application, that will have both features of voting and verification. We developed a first version that is a low-fidelity paper prototype. We evaluated this version with HCI experts. The received feedback on design and understanding of security allowed us to iterate and develop a second paper prototype, that was tested with

both HCI and security experts. The final iteration was a high-fidelity software version that will be described in detail in subsection 2.3.

A particular challenge for the user experience of Selene is that a certain level of understanding might be necessary to achieve fulfillment of privacy needs: as secure and easy-to-use as the application might be, displaying the plaintext cast vote to the voter after election could seem insecure. Further, the verification phase is not commonly used in standard elections and is largely unknown to users. Following a UXD process, we tried to anticipate users' expectations on voting and their questions on such a protocol, hence we designed an interface that, hopefully, is more understandable.

Cryptography Selene uses several security mechanisms, however, the cryptographic details can be hidden from the user during the voting and verification phases. As mentioned above, we assume here that the voter doesn't have to configure her device with her secret keys or explicitly handle other keys such as the election public key.

As mentioned before, the tracking number retrieval (fourth step in the voting experience) is simplified here and the voter simply has to click on a single button to use the α -term, to download the β -term and to compute the tracking number. It will highlight the result on the bulletin board, displayed in-app, automatically. The voter's trapdoor key won't be explicitly manipulated by the voter and it is embedded in the phone, unlocked by the voter's credentials.

The other primitives used in Selene do not require direct interaction with voters (e.g. zero-knowledge proofs, mixnet, PET tests). Hence these are not mentioned in the conducted user experience test. In a real election implementation all of this can be public and verifiable by observers and interested voters.

We emphasise that this implementation is a first step that provides a user interface, in order to answer our research questions on user experience. This application is not ready to be used in a real election as both software security and the full cryptographic features have not been integrated yet. As described in the protocol, the public key, the encrypted tracker, the commitment and the encrypted vote should be displayed on the bulletin board after every vote update. In this study we simplify and only update the bulletin board in the end.

Trust assumptions Even if not all of the cryptographic primitives have been integrated in this version, we can already discuss the consequences of the design choices on the security properties. Firstly, we assume that the voting device is trusted for privacy. Further, in this test we have used a single device for voting and verification. In real scenarios, we would recommend that different devices, or at least apps, are used for vote-casting and vote-verifying for improved security.

The reason for using only one device was to simplify the experience for the participants focusing on a basic voting and verifying experience and to test this. The tracking number retrieval is also automated: the voter does not have to manually combine the α and β terms and decrypt the tracking number. Since the α term is not shown to the voter, no visible α term needs to be faked, but a coercion-mitigation mechanism stills needs to implemented in the app to fake

the tracking number itself. Further, the level of receipt-freeness in Selene will also depend on the chosen vote-casting method, e.g. a Helios type of electronic ballot will only achieve software-dependent receipt-freeness. However, this is a first iteration in the UX development of Selene, and the feedback from the participants given in section 4 will help us to take the correct direction in the future developments.

Finally, the verification phase was mandatory as a part of the test procedure. But in a real election it is to be expected that not all of the voters will verify their vote. We have not investigated the voters' motivation yet.

2.3 Interface implementation

The final application has been developed with the Android native language (Java) and the back-end server is developed in php and deployed on an Apache server. No security analysis has been performed as the goal of this interface is to run user tests. The security of the application remains basic. We describe below the interfaces provided.

Android application The final application contains the two phases: one for voting and one for verification. The application retrieves flags related to the voter after authentication, that indicates: the voter's state (has voted (**true/false**)), and the election's state (**vote/verify**).

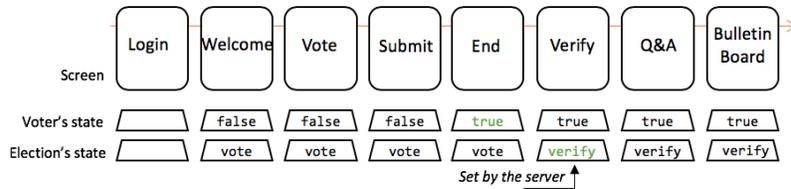


Fig. 2. Application workflow with states.

Figure 2 shows the organization of the application with the corresponding flags. The application has been developed with a linear workflow, the voter only has a minimal choice for navigation, namely going backwards or forwards. One should notice that in the context of a real election, the application might contain more screens with additional information. However, we will discuss in section 5 the advantages of such a linear construction.

Administration page The back-end server is used to verify voter eligibility during authentication, to receive votes sent by the app, and tally the results. When tallying the results, all pairs (tracker, vote) are counted and published on the Bulletin Board. When allowing a voter to verify, the flag for the election's state is set to **verify** and the voter will be able to go through the verification workflow in the application.

Bulletin Board The bulletin board is retrieved during the verification phase by the application. An additional button lets the voter highlight her vote. For this experiment, the bulletin board was accessible on the phone only but can be accessed directly from any browser, however it only contains minimal data needed for the user test.

3 User Testing protocol

Participants We recruited 38 French participants (19 male and 19 female) through social networks, trying to ensure a fair distribution of our sample in terms of gender, age and education level. The average age was 35,4 years old (Min=19, Max=73, SD=12,45). The education level broadly varied as well: no diploma (13%), A-Levels (29%), some college degree (21%), Bachelor (18%), Master (16%) and PhD (3%). The study has been run in French and the data presented has thus been translated into English.

To make their answers consistent and accurate, we selected participants that had participated at least in one political national election in France.

Procedure We provided each participant with a paper sheet explaining the context of the user test, that is a national election in France, together with the candidates' programs. Two personalized letters were distributed to each participant to provide them their individual credentials to access the application. Then the sessions were split up into 4 phases: (1) the voting phase, (2) a semi-structured interview, (3) the verification phase and (4) a semi-structured interview³. Before the verification phase, we gave them a second letter which was an invitation to verify their vote using the application.

Methodology The goal of the present analysis is to identify which mental models participants have of privacy and verifiability in e-voting. The semi-structured interviews entailed the following topics: general opinion about the application, trust, control, understanding of the verification phase and of the bulletin board. The three first topics were addressed after both the voting and verification phases. The two last topics were addressed after the verification phase only. We avoided security priming by not addressing security-related topics (such as privacy) until the very end of the study in order to avoid influencing participants' answers. In most cases, they mentioned by themselves the different security issues they could face with regards to e-voting. We describe in section 4 which mental models we identified. Information about the verification procedure was provided through paper letters and inside the application. The Q&A screens are mandatory in the workflow and the participants have to go through them before verification. We told the participants that the tracking number let them

³ A questionnaire about User Experience and Psychological Needs Fulfilment were also filled by participants during phases (2) and (4). The analysis is discussed in an other paper [7].

verify that their vote has been counted in the final tally, that it helps to validate the election results, that this tracking number is unique and that the count can be verified by anyone. As we did not want to prime participant with possible security issues, we have not mentioned the risks of using one device and the associated trust assumptions.

Data analysis The user test was devised as a between-subjects study, and two versions of the e-voting application have been tested with our participants: half of the participants tested a baseline version and the other half an extended version where security aspects are additionally displayed to the participant. It contains additional information about the ongoing process in the application yet with no extra interaction. The impact of displaying security-related mechanisms was the topic a recently published paper [7] alongside with additional factors impacting UX (attractiveness, novelty, etc.) and psychological needs (autonomy, competence, etc.). Interestingly, we noted that this additional layer of communication remained largely unseen by the participants, with the perceived security being rated as only slightly higher in the elaborated version. The analysis of this paper focuses on interviews only and explore the feedback of participants regarding the security properties of voting, to check their understanding as it will be described in the next section.

To analyze the data retrieved from the interviews, we followed the methodology described in [16]. We coded the data through a theoretical thematic analysis, to look for patterns relating to the voting security properties. We organize the participants' answers into a list of concepts. We classified these concepts in categories given in section 4 to understand voters' mental models of security. The categories were organised to match with the known theoretical models of security of e-voting: privacy properties including ballot-secrecy and coercion-resistance, and verifiability.

The qualitative analysis of answers in the semi-structured interviews leads to similar concepts for both versions, and we will thus analyze the participants' feedback in a similar way without considering the tested version.

One goal of user-centred design is to achieve a better alignment between participants' mental models and researchers' security vision, by "ensuring that products do fit real needs, that they are usable and understandable" [20], we will discuss this in section 5.

Ethics The study follows the guidelines provided by the ethics commission at the authors institution and was conform to GDPR.

4 Mental models

In [19], Norman defined mental models as being "people's views of the world, of themselves, of their own capabilities, and of the tasks that they are asked to perform, or topics they are asked to learn". The interactions they have with the

environment make them form internal models of the system they are interacting with. We here propose a categorization of participants' feedback. An overview is given in figure 3. From the identified concepts, we derive a categorization of the mental models expressed by participants. We will discuss how these results should impact the future development of the application in section 5.

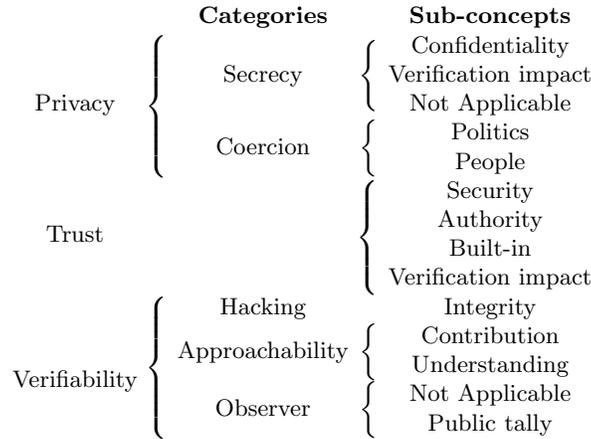


Fig. 3. Mental model categorisation

We now explain this structure in detail.

4.1 Privacy

Secrecy Mental Model Participants mentioned that their vote must be kept *confidential and anonymous*. They questioned the data management, wondering if someone knows the relation between identities and votes. Some participants stressed the importance of the booth, another argued that the booth is not private either, "some people are looking" (P10).

The *verification* could have a negative impact on secrecy of votes as well, "it is like someone else could see it too" (P22).

Finally, another concept was *not applicable*, as some people did not feel concerned by secrecy, as "others know who I am voting for" (P13) or "you have to take responsibility for your political decisions" (P38).

Coercion Mental Model Coercion from *people* (in the sense of a physical attempt to coerce) was mentioned several times. In particular, participants mentioned the advantage of being able to vote at home, as other people won't influence them: "Here I don't have interactions with other people" (P5), or "I am sure to make my own choice [...] I feel less pressure than in polling station, with people behind" (P2). Vote buying was mentioned once "We can be manipulated, one could buy our vote, but we need to evolve" (P23).

Finally, the *political aspect* of coercion was also mentioned a few times, as some parties could try to cheat and to steal credentials from voters: "We must

pay attention to parties, ensure there is no violation, that the elderly or other vulnerable persons do not get their vote stolen” (P10).

4.2 Trust

The concept that appears the most for Trust is *security*. Participants mentioned that their trust in the application is highly dependent on the security provided: ”I don’t trust it, how could we know if it is really secure?” (P29), ”I trust it, there are breaches everywhere but I think we can secure this” (P10). In particular it was reflected on their other mental models related to privacy and efficiency. Hence, we can derive this security concept with the following sub-concepts: coercion, secrecy, and understanding that increased or decreased the security perceived.

Another concept was *authority*, mentioned as trust-transference in [4]. Participants refer to some trusted third party to emphasize their own trust: ”if it is done by an authority, I will trust it” (P2) or ”I trust the government, they will do what is necessary to ensure vote security” (P12). They rejected the verification process arguing with their trust in the authority: ”If I trust the application I don’t see why we should verify that the vote has been taken into account” (P12).

Some participants expressed a *built-in* trust, e.g. ”I always trust technology” (P14) or ”I trust it as I would trust any mobile applications”.

A *verification impact* was raised, mostly decreasing trust, e.g. ”I don’t trust the application after verification, even if the tracking number is private” (P33), even though an opposite positive effect on trust was also mentioned by some users: ”the second phase makes me feel secure” (P4).

4.3 Verifiability

Hacking Mental Model Participants were concerned by the security of internet technologies and had many preconceptions. Even if participants didn’t master the complexity of internet security, they were aware that it could be an issue. For example, they mentioned problems they heard about other voting systems with electronic ballots: ”In United States there was this elections hacking. Paper is more reliable” (P15). Others feared internet technologies in general: ”I think internet is vulnerable, even if the app is secure” (P24). Ballot stuffing was also mentioned as a big problem: ”There are people who can buy hackers’ services to have thousand of votes added, we will never know.” (P28).

Integrity is a concept that often appeared during the interviews. Participants questioned the good behavior of the application as they did not receive any proof of it. The reliability of the system is questioned: ”It does not guarantee that it is really my vote.” (P33). Some participants also expressed the need for a procedure in case of encountering an issue: ”Who should I call in case of problem? And if my vote is not in the list?” (P19) or ”If I voted A and it shows B, what should I do?” (P32).

Approachability Mental Model Some participants were convinced of the good behavior of the system due to the verification phase. It was mentioned as a proof of their personal *contribution* to the elections: "Seeing that my vote is taken into account, seeing others' votes, it lets me believe that I contribute to something" (P18) or "It is important to see that my vote has been counted" (P27).

Most participants understood that they were seeing a confirmation of their tallied-as-intended vote. But they expressed their [lack of] *understanding* in the process of verification: "I feel in control maybe because I can see what I did, I can see my vote again" (P11) or on the contrary "I wonder why this is here, seeing results with percentages is enough for me" (P3).

We tried to rate participants' understanding of Selene's mechanism, through the two last questions stated in section 3. To help participants to answer, we provided some light explanation of the verification phase meaning. However, many participants did not manage to provide a complete description of the verification phase after using the app. Furthermore, the tracking number has not always been understood as such, but rather as a counting of votes: "We can see our candidate and the number of people who voted for him" (P6).

Observer Mental Model Some participants stressed the importance of observation. In France, voters are allowed to go to the polling-station to observe the public count of votes. However, most of our participants did not noticed the link between this real-life procedure and the availability offered by the bulletin board: "The list is not really informative" (P35), "I can't see if there is any interest to see this list with all details" (P17). This can be explained by their lack of understanding of the procedure, as compared to a physical count of votes, in which they can see and understand each step: "In polling stations you can verify by yourself, on internet it's questionable" (P24).

Finally, the *individual* aspect seems to be enough to participants, e.g. "Seeing percentages with general results, and my individual vote is enough" (P17).

5 Discussion

Norman in [19], and Cooper et al. in [6], show that three models must be considered in the design of a user interface: the system or implementation model that reflects how the system actually works, the system image or represented model that reflects what is shown to the user and the mental model that is the projection made by the user. Here we focus on the discrepancies between those three distinct categories. The goal of a user-oriented design process, such as the UXD process, is to provide an interface, a represented model, that is close to the user's mental model and that remains accurate with the system model.

5.1 Comparison between mental models and security properties

The properties on which we base the implementation model of a voting scheme are Privacy properties and Verifiability. Selene provides ballot-secrecy, receipt-

freedom and has a coercion mitigation mechanism. It also provides individual and universal verifiability. However, as mentioned in the previous subsection, the coercion mitigation mechanism has not been implemented.

Despite this, voters were concerned about Coercion and about Privacy during elections in general. Mental models for Privacy were consistent with the properties of the system, and the reason might be that Privacy is a mandatory element required by law during elections in France, and it is taught to people at an early age at school.

On the other side, the novelty of the verification phase seemed to prevent participants from properly explaining their experience. However, indirect properties and potential issues were mentioned, such as hacking and integrity, and public tallying. It appears that participants were able to point out the potential issues of online voting without seeing that the verification mechanism was part of the solution.

Also for privacy, we can argue that an early education on verifiability could lead to a better understanding and acceptance of the concept.

Trust is not a security property of voting protocol. However, it plays an important role for voters and impacts the use of a system. This aspect is important for people to accept the system they use.

Even if the convenience of online voting was mentioned many times, voters stressed their lack of knowledge about internet technologies as a big drawback. Paper-ballot voting contains steps that are understandable and accessible to people, and this is not in general the case for online voting. Even if this aspect is not required by law as in Germany, it seems reasonable that voters are more willing to trust a process they fully understand.

5.2 Impact of a user-centred application on the voting experience

First of all, we observed 100% of effectiveness for vote casting: all participants were able to cast their vote successfully.⁴ The application was designed in order to be easy-to-use and responding to users' expectations, and we can argue that it is the linearity of the vote casting in the Selene implementation that leads to this excellent result. The quality of this straightforward behaviour was mentioned several times by participants, e.g. "we follow the workflow but we can't really make a mistake" (P3).

Another explanation for the observed usability can be the design of the protocol itself. As mentioned in the introduction, Selene was designed in order to reduce complicated interactions with users, and to be more intuitive. Helios is another e-voting scheme that requires, or at least suggests, voters to perform audits of their ballots through a Benaloh's challenge [3]. This often leads to a lower effectiveness rate: a study from Marky et al. [17] has shown that this procedure is considered as counter-intuitive by participants. Indeed, participants

⁴ The verification phase was mandatory in our experiment and everyone managed to go through the verification workflow. But not all participants understood what was happening and we can't ensure that the effectiveness would be as high for verification if it is not mandatory.

who audited their ballot did not understand why they were not allowed to cast the audited ballot after-all. This kind of step does not occur with Selene, as the verification happens after the end of the election. The voting phase is thus not burned with a verification step. Moreover, the authors of [17] showed that automation of the verification feature improved effectiveness. In our application, we automated the retrieval of the tracking number. Instead of asking participants to manipulate *alpha* and *beta* terms, we retrieve them and computed the tracking number automatically.

The questionnaires analyzed in [7] showed that the usability aspects, i.e. efficiency, perspicuity and dependability [26], scored above average. Despite this and the fact that all voters managed to cast their votes, the participants' feedback show that our application needs more development iterations to be better consistent with voters' model of voting.

Also, simply displaying information about security features to the participants was not enough to make them explicitly see it (as discussed in [7]). However, we have seen here in our analysis that the security of such an application is an important factor for trust and it was emphasized for the secrecy and verifiability concepts. Moreover, the information related to security in the expanded version of the app was shown during loading screens. It might be that the progress bar prevented the participants from reading the information displayed below⁵. In a small study [12], the authors found that the voters chose more secure systems as their preferred scheme even if they scored lower on the SUS scale. It is thus interesting whether allowing more cryptographic interactions could increase the acceptance, even if it reduces the usability. One idea could here be to also implement coercion mitigation mechanism. This mechanism allows the voter to ask for a fake tracker in case of coercion. It might be that voters have missed this mechanism to understand the verification phase.

Indeed, the meaning of the verification phase and in particular of the tracking number was explained through Q&A screens. However, many participants did not understand fully, or were not able to describe the verification phase. As Acemyan et. al observed in their study [2], when participants were requested to draw their mental model, they expressed the voting steps for each tested scheme and avoided the verification steps and the verification phases of each system was considered useless in many cases, like in our observations. On the other hand, participants who understood the interest of seeing their vote in the app did not understand why they were seeing others' votes, as their own vote and only this vote was highlighted in the application. The implementation of the coercion mitigation mechanism could also help here, however this assumption needs to be tested in a new iteration of the application.

Olembo et al. [23] showed that specific messages could motivate voters to verify their vote, as they understand better the objective of such a procedure. In particular, they focused on risks, norms and analogies. In our application, the focus has been done on norms only, i.e. we explained what is the purpose of verification and what it brings to society. We emphasized democracy protection

⁵ An other study could verify where the information displayed would be more visible.

and integrity of votes records. Now, according to voters' models for Coercion and their concerns on hacking, a stronger emphasis on the incurred risks and solutions provided by verification might help the voter to understand. Some people understood that the tracking number was instead the number of people voting like they did. A simple improvement is to add letters to the tracking number.

In this version of the application, the bulletin board was not accessible before the individual verification phase⁶. One improvement could be to make the bulletin board available before the individual verification. The possibility to request a fake tracker must also be implemented at this stage. In addition, once the Selene check is doable, we could show the individual vote first (with fake or real tracking number) and let the voter consult the bulletin board on purpose.

Limitations The results of our study are bounded by some limitations.

First, the user tests were done in a laboratory that had a reassuring impact on participants. Some of them admitted that they were not really feeling any threats for their vote as they were part of an experiment. The influence in a lab context on user studies is discussed in [15].

We mention to our participants that the elections were related to the national elections in France, however we did not use real candidate names nor run an election that already happened, as suggested in [18].

Also, the participants had a very limited amount of time to understand the verification procedure, and the novelty of such a protocol might require more time to be understood and accepted. A broader context would be provided in real elections, giving users time to understand the process of verifiability of the application. We also assumed in our study that the configuration of the devices was already done. The ease of use might be questioned if the registration to on-line voting and keys configuration must be performed by voters. However, this configuration could be done only once for several elections.

Finally, we asked the participants to verify their vote right after the vote casting phase. In other protocols and user studies, the verification is done during the vote casting (e.g. Benaloh's challenge, or return codes). In this protocol, the verification is performed after the results have been published and due to experimental constraints the participants had to do it right after vote casting, that could have disconcerted them.

6 Conclusion and Future Work

In this paper, we have provided the first interface prototype for the Selene e-voting protocol. We have followed a user-centered design, UXD, where the usability is enhanced and cryptographic interactions have been hidden. This approach has consequences on trust assumptions for the voting protocol, but has provided insights on the mental models of Privacy and Verifiability. User tests have highlighted possible improvements on our application for Selene, but it has also

⁶ See third phase of the voter experience described in subsection 2.1

raised more general concerns we need to consider in the design of e-voting protocols. We have seen that mental models for Privacy with Secrecy and Coercion were consistent with the voting protocol concepts. However, the understanding of the verification phase has to be facilitated. We have seen that the lack of understanding could lead to trust issues: participants questioned integrity of the elections and the purpose of the verification phase. Voting schemes are developed today to be end-to-end verifiable, but verification is not natural to users and voters need more time to accept it and understand it. An easy-to-perform mechanism for verification like the one described in Selene has been effective but is not enough to convince voters of the security behind the scheme. For future work, the implementation of missing mechanisms for Selene must be performed in order to provide a complete experience to voters. A new iteration of the application (using two devices) based on the received feedback is being developed in order to increase the understanding of voters, and reassure them of the security mechanisms in use.

Acknowledgements We would like to thank the Luxembourg National Research Fund (FNR) for funding, in particular PBR was supported by the FNR INTER-Sequoia project which is joint with the ANR project SEQUOIA ANR-14-CE28-0030-01, MLZ was supported by the INTER-SeVoTe project and VD was supported by FNR grant number PRIDE15/10621687.

References

1. Acemyan, C.Z., Kortum, P.T., Byrne, M.D., Wallach, D.S.: Usability of voter verifiable, end-to-end voting systems: Baseline data for helios, prêt à voter, and scantegrity II. In: 2014 Electronic Voting Technology Workshop/Workshop on Trustworthy Elections, EVT/WOTE '14 (2014)
2. Acemyan, C.Z., Kortum, P.T., Byrne, M.D., Wallach, D.S.: Users' mental models for three end-to-end voting systems: Helios, prêt à voter, and scantegrity II. In: Human Aspects of Information Security, Privacy, and Trust - Third International Conference, HAS 2015 (2015)
3. Adida, B.: Helios: Web-based open-audit voting. In: Proceedings of the 17th USENIX Security Symposium, July 28-August 1, 2008 (2008)
4. Ali, S.T., Murray, J.: An overview of end-to-end verifiable voting systems. CoRR (2016)
5. Benaloh, J., Byrne, M., Kortum, P.T., McBurnett, N., Pereira, O., Stark, P.B., Wallach, D.S.: Star-vote: A secure, transparent, auditable, and reliable voting system. CoRR (2012)
6. Cooper, A., Reimann, R., Cronin, D., Noessel, C.: About face: the essentials of interaction design. John Wiley & Sons (2014)
7. Distler, V., Zollinger, M., Lallemand, C., Rønne, P.B., Ryan, P.Y.A., Koenig, V.: Security - visible, yet unseen? In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (2019)
8. Greene, K.K., Byrne, M.D., Everett, S.P.: A comparison of usability between voting methods. In: 2006 USENIX/ACCURATE Electronic Voting Technology Workshop, EVT'06 (2006)

9. Herrnson, P.S., Niemi, R.G., Hanmer, M.J., Bederson, B.B., Conrad, F.G., Traugott, M.: The importance of usability testing of voting systems. In: 2006 USENIX/ACCURATE Electronic Voting Technology Workshop, EVT'06 (2006)
10. Juels, A., Catalano, D., Jakobsson, M.: Coercion-resistant electronic elections. In: Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society, WPES 2005 (2005)
11. Karayumak, F., Kauer, M., Olembo, M.M., Volk, T., Volkamer, M.: User study of the improved helios voting system interfaces. In: 1st Workshop on Socio-Technical Aspects in Security and Trust, STAST 2011 (2011)
12. Kulyk, O., Neumann, S., Budurushi, J., Volkamer, M.: Nothing comes for free: How much usability can you sacrifice for security? IEEE Security & Privacy (2017)
13. Kulyk, O., Volkamer, M.: Usability is not enough: Lessons learned from 'human factors in security' research for verifiability. IACR Cryptology ePrint Archive (2018)
14. Lallemand, C., Gronier, G.: Mthodes de design UX : 30 mthodes fondamentales pour concevoir des expriences optimales. France, Paris : Eyrolles (2018)
15. Lallemand, C., Koenig, V.: Lab testing beyond usability: Challenges and recommendations for assessing user experiences. Journal of Usability Studies (2017)
16. Lazar, J., Feng, J., Hochheiser, H.: Research Methods in Human-Computer Interaction, 2nd Edition. Morgan Kaufmann (2017)
17. Markey, K., Kulyk, O., Renaud, K., Volkamer, M.: What did I really vote for? In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI 2018 (2018)
18. Markey, K., Zollinger, M.L., Funk, M., Ryan, P.Y., Mühlhäuser, M.: How to assess the usability metrics of e-voting schemes. In: Financial Cryptography and Data Security - FC 2019 (2019)
19. Norman, D.A.: Mental models. In: GENTNER, D., A.L. STEVENS, e. (eds.) Human-computer Interaction, chap. Some Observations on Mental Models, pp. 7–14. Lawrence Erlbaum Associates Inc. (1983)
20. Norman, D.A.: The Design of Everyday Things. Basic Books (2013)
21. Norman, D.A., Draper, S.W.: User Centered System Design; New Perspectives on Human-Computer Interaction. L. Erlbaum Associates Inc., Hillsdale, NJ, USA (1986)
22. Olembo, M.M., Bartsch, S., Volkamer, M.: Mental models of verifiability in voting. In: E-Voting and Identify - 4th International Conference, Vote-ID 2013 (2013)
23. Olembo, M.M., Renaud, K., Volkamer, S.B..M.: Voter, what message will motivate you to verify your vote? USEC (2014)
24. Ryan, P.Y.A., Rønne, P.B., Iovino, V.: Selene: Voting with transparent verifiability and coercion-mitigation. In: Financial Cryptography and Data Security - FC 2016 (2016)
25. Schneider, S., Llewellyn, M., Culnane, C., Heather, J., Srinivasan, S., Xia, Z.: Focus group views on prêt à voter 1.0. In: 2011 International Workshop on Requirements Engineering for Electronic Voting Systems, REVOTE 2011 (2011)
26. Schrepp, M.: User experience questionnaire handbook (2018)
27. Standardization, I.O.F.: Iso 9241-11: Ergonomics of human system interaction part 11: Guidance on usability (1998)
28. Standardization, I.O.F.: Iso 9241-210: Ergonomics of human system interaction part 210: Human-centred design for interactive systems (1999)

Internet Voting Governance: Canada and Estonia

The curse of knowledge? Does having more technology skills lead to less trust towards ivoting?

Mihkel Solvak¹[0000-0003-0179-4036] and Robert Krimmer²[0000-0002-0873-539X]

¹ University of Tartu, Ülikooli 8, 51003 Tartu, Estonia

² Tallinn University of Technology, Akadeemia tee 3, 12618 Tallinn, Estonia
mihkel.solvak@ut.ee
robert.krimmer@taltech.ee

Abstract. The Estonian data shows that people with higher ICT literacy are younger, better educated, have a better income and live in cities compared to people who self-report to have low skills. This is typical for the digital divide that holds almost universally in all societies [3] and the self-branded “digital nation” of Estonia is no exemption here. On top of the digital divide there is however another quite universal “law” for lack of a better word – people with better knowledge of technology tend also to be more aware of the limitations and potential threats it might entail. It holds especially for the digital realm where one can assume, that the more you knowledgeable you are, the more threat averse you also tend to be. Voting technologies are clearly a case in point as demonstrated by the vocal opposition to electronic voting by groups with advanced knowledge in [1,2,5]. In light of this one could say that ivoting diffusion and usage faces a curse of knowledge. The more you know about technology the less trusting towards it you should be given the heightened integrity standards set for elections [4]. We would take it even further and argue that based on the above, ivoting seems to face a class ceiling in terms of usage, it should not be used or trusted among the ones who ideally should be the target audience, i.e. those who know most about the potential of this technology.

Keywords: trust in ivoting, trust in technology, diffusion of ivoting

Survey data on Estonia shows that people with higher ICT literacy are younger, better educated, have a better income and live in cities compared to people who self-report to have low skills. This is typical for the digital divide that holds almost universally in all societies [3] and the self-branded “digital nation” of Estonia is no exemption here. On top of the digital divide there is however another quite universal “law” for lack of a better word – people with better knowledge of technology tend also to be more aware of the limitations and potential threats it might entail. It holds especially for the digital realm where one can assume, that the more you knowledgeable you are, the more threat averse you also tend to be. Voting technologies are clearly a case in point as demonstrated by the vocal opposition to electronic voting by groups with advanced knowledge in technology [1,2,5]. In light of this one could say that ivoting diffusion and usage faces a curse of knowledge. The more you know about technology the less trusting

towards it you should be given the heightened integrity standards set for elections [4]. We would take it even further and argue that based on the above, ivoting seems to face a glass ceiling in terms of usage, it should not be used or trusted among the ones who ideally should be the target audience, i.e. those who know most about the potential of this technology.

We examine this problem from the (potential) user perspective to see how trust towards ivoting technology is dependent on the technological skill level. We also investigate to what degree are technological solutions that should help to build more trust, such as individual vote verification, actually being used depending on the trust and skill level of the user. We do so employing two unique datasets. First survey data from Estonia between 2013-2019 entailing post-election surveys on ivoting from a total of 5 elections all in all with more than 5000 interviews, and second using anonymized ivoting log data from the same period on more than 800 000 ivoting sessions.

The results are puzzling. We do see a clear “ivoting digital divide”. A majority has very high trust levels and a small but persistent minority extremely low trust levels with hardly anyone at mid trust levels - ivoting seems to polarize trust towards itself. But when the trust distribution is examined according to computer literacy the exact opposite holds. The more skills a person has, the more trusting towards ivoting they actually are (see Figure 1).

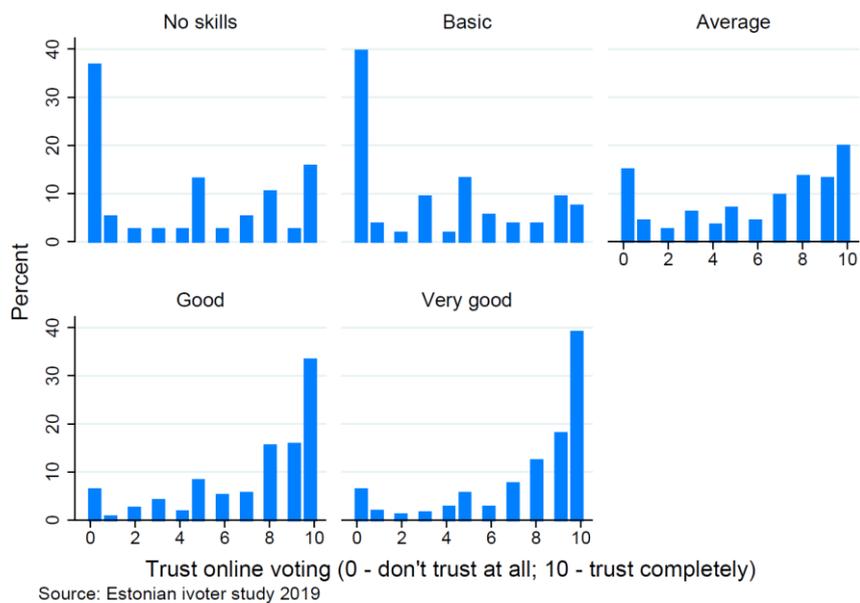


Fig. 1. Trust towards online voting by computer literacy levels.

Even after controlling for other covariates that correlate with better ICT skill levels, such as age, gender, education or income, people with good or very good computer literacy levels have on average between 15 to 19 percentage points higher trust score towards i-voting than people with very poor or non-existing skills. There is a strong association between having voted electronically and trusting this voting mode, so we might be conflating user experience with general attitudes. But the positive association between skill and trust level is still clearly observable also among the non-i-voting and non-voting population, who in general tend to have a much lower trust levels towards online voting than the actual users.

Examining the typical user profiles from the log files we notice that i-vote verifiers are in comparison to non-verifiers on average four years younger and made up of 70% males whereas among non-verifiers males make up less than 50%. Verifiers, who made up an average of 5% of i-voters in the last four elections, have also proportionally four times more Linux and two times more Mac users among them than non-verifiers. In addition, the largest share of i-votes are verified on the very first voting day just as electronic polls open, indicating more active usage among the most eager electronic voters. And finally, the largest share of i-votes being verified at any given time over the seven day i-voting period are votes given late at night between 12pm and 3am. All in all this points towards verification being more likely used by young males, using Linux and voting late in the night. This is indeed a typical profile of a tech savvy i-voter who might use verification exactly because they tend to know more about how technology works and would like to see additional technological developments that makes i-voting more resilient against potential threats.

The paradox reintroduces itself however when we use survey data to connect verification with perceived trust levels. Vote verifiers have by far the highest trust level towards electronic voting out of all possible other groups (non-voters, paper voters, non-verifying i-voters). In 2019 for example the average trust level towards i-voting among verifiers was 9 on a 0-10 trust scale. We also see that being knowledgeable about the possibility of i-vote verification option correlates with high trust level even when people have not used it to verify the vote (or even i-voted) and people who do not know about the option have clearly lower trust levels. In other words, knowing about the option to verify, without necessarily using it, already has a positive effect on trust towards i-voting. The survey evidence therefore clearly points towards verification being used by tech savvy voters who at the same time exemplify excessively high trust towards the technology.

In conclusion we find overwhelming evidence that more knowledge leading to more skepticism does not hold in the Estonian case. I-voting has been available in Estonia since 2005, individual verification was introduced in 2013, the perceived trust towards this technology has over the years balanced around 70% of people having either high or very high trust. One can assume, that the large experience with using electronic services in the public sector, the over 10 years tradition using electronic voting seem to have had an impact on the overall trust of users in electronic voting. Consequently,

Oostveen and van den Besselaar [4] identified in their original study that context plays an important factor, including learning how one can verify the functionality of a technical system such as i-voting.

1 References

1. Gonggrijp, R., & Hengeveld, W. J. (2007, August). Studying the Nedap/Groenendaal ES3B voting computer: A computer security perspective. In *Proceedings of the USENIX workshop on accurate electronic voting technology* (pp. 1-1). USENIX Association.
2. McGaley, M., & McCarthy, J. (2004). Transparency and e-Voting: Democratic vs. commercial interests. *Electronic Voting in Europe*, 47, 153-163.
3. Norris, P. (2001). *Digital divide: Civic engagement, information poverty, and the Internet worldwide*. Cambridge University Press.
4. Oostveen, A. M., & Van den Besselaar, P. (2004). Security as belief: user's perceptions on the security of electronic voting systems. *Electronic voting in Europe: Technology, law, politics and society*, 47, 73-82.
5. Springall, D., Finkenauer, T., Durumeric, Z., Kitcat, J., Hursti, H., MacAlpine, M., & Halderman, J. A. (2014, November). Security analysis of the Estonian internet voting system. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security* (pp. 703-715). ACM.

Online Voting in a First Nation in Canada: Implications for Participation and Governance

Brian Budd¹[0000-0003-3554-430X], Chelsea Gabel²[0000-0002-1007-5351] and Nicole Goodman³(✉)[0000-0002-8607-2595]

¹ Department of Political Science, University of Guelph, Guelph, Canada
buddb@uguelph.ca

² Department of Health, Aging and Society and Indigenous Studies Program, McMaster University, Hamilton, Canada
gabelc@mcmaster.ca

³ Department of Political Science, Brock University, St. Catharines, Canada
Nicole.goodman@brocku.ca

Abstract. Indigenous communities are increasingly adopting technology to create digital opportunities for members and enhance engagement and governance. One recent trend in the adoption of online services is the use of online voting. To date, more than 90 Indigenous communities in Canada and the United States have deployed online voting with many more considering implementation. This article draws upon interviews with local government officials and voter exit surveys as part of community-engaged research with Wasauksing First Nation in Ontario, Canada to explore the specific opportunities and challenges online voting presents for governance and engagement in Indigenous communities and implications for future adoption. Specifically, we examine a 2017 Land Code vote where online voting was introduced to achieve a participation threshold required to pass the framework. Our findings point to online voting as a key tool to modernize Indigenous governance and enhance participatory capacity by making voting more accessible for members. We argue that online voting is an engine that can advance self-determination and support communities seeking an iterative path to self-government.

Keywords: Online voting, Indigenous governance, First Nations, Self-determination, Canada, Community-Engaged Research

1. Introduction

To date online voting has been used by a growing number of Indigenous communities across Canada and the United States for elections or other types of community votes. Online voting is appealing to Indigenous communities as a tool to improve voter accessibility and engagement, especially for communities where large segments of the membership live off of reserve lands [14]. While the engagement of off-reserve members is important to ensure balanced representation of community voice, in the

Canadian context it is crucially important since many communities are subject to federal legislation that has required them to meet participation thresholds in order to pass community laws and gain back autonomy. In this regard, online voting use among First Nations in Canada represents a tool for communities to not only bridge participatory gaps with off-reserve members, but also to increase their capacity to ratify their own legislation and move away from federal control over decision-making and governance processes. Despite optimism about online voting's potential to enhance participation and governance, questions persist about the ability of the technology to do so. There are also important considerations around the cultural appropriateness of online voting and whether its adoption is consistent with Indigenous culture, community visions of self-determination, and local decision-making.

This article examines the opportunities and challenges online voting presents for participation and governance for Indigenous communities through a case study of Wasauksing First Nation in Ontario, Canada. Building upon previous research focused on understanding First Nations' satisfaction with online voting, online voter characteristics and the potential for online voting to improve voter engagement in a First Nation context [14], we draw upon semi-structured interviews with local government officials, participant observation data, and exit survey data from online and paper voters as part of a land code ratification vote in February 2017 to examine how online voting affects perceptions of participation and governance. Specifically, we assess the degree to which online voting contributed to the community's capacity to ratify the land code legislation and its potential to enhance, or limit community capacity and self-determination in the future. The original contribution of this work is its focus on the implications online voting presents for governance and the enhancement of self-determination in First Nations. Findings suggest that online voting is a tool to modernize Indigenous institutions and governance, improve community connectedness, particularly by better connecting off-reserve members in policy discussions, and as an engine to support meeting quorums for critical votes that, if successful, can advance community capacity and enhance self-determination as part of an iterative path to self-government.

This article proceeds as follows. First, we present some brief background about the colonial context in Canada and how online voting can contribute to First Nations gaining back their political power. This is followed by a literature review that explores understandings of Indigenous self-determination and self-government to guide our assessment of the impact of digital technology in Wasauksing First Nation and reviews what studies have found about First Nations' use of online voting. Next, we present an overview of Wasauksing and explain the rationale for, and steps taken to, implement online voting for the ratification vote. In the fourth and fifth sections, we present our methodological approach and analysis. Finally, we discuss what our findings mean for future adoption of online voting in First Nations and Indigenous communities more broadly, notably implications for self-determination, community engagement, and governance.

2. Background

Indigenous communities around the world face many social, economic, political and legal inequalities arising out of ongoing colonial legacies. In the case of Canada, the presence of paternalistic legislation designed to subvert complex systems of Indigenous governance has created enduring conditions of dependency as First Nation, Métis and Inuit communities have been subordinated under the oversight and administration of the federal government. While Indigenous resistance to colonial administration has been constant throughout Canada's history, in the past 50 years there has been a shift in the relationship between Indigenous peoples and the Canadian state. Bolstered by the emergence of an international Indigenous rights movement [19], Indigenous peoples have taken steps to assume greater control over governance and policy-making within their communities, albeit with overarching institutional apparatuses of colonial administration still intact. While the assertion of the right to self-determination has become a prominent feature of Indigenous politics in Canada, variation exists in the approaches Indigenous communities take to achieve self-determination and their capacity to pursue it.

Numerous strategies and approaches have been taken to achieve self-government and enhance self-determination. One strategy is to take larger steps such as the signing of self-government agreements and land claim negotiations, while a second type of approach is more incremental, characterized by smaller steps involving the devolution and decentralization of decision-making into the hands of Indigenous communities and stakeholders. Literature on Indigenous politics has mostly focused on the former set of approaches, exploring the "high politics" of self-determination and self-government, emphasizing formal agreements and negotiations between Indigenous communities and other levels of government [17]. Far less attention has been given to the smaller, less-visible steps that communities are taking to achieve the same ends. Aided in part by legislative opportunities created by the federal government, an increasing number of Indigenous communities across Canada are enacting iterative steps toward the development of self-determination through the ratification of community-based legislation [11]. In recent years the iterative approach has become increasingly possible through the adoption of online voting as a tool to enhance voter engagement and pass key pieces of legislation. We explore this process and its impacts in this article in the context of First Nations in Canada.

3. Indigenous Governance, Digital Technology & Online Voting

3.1. Indigenous Governance

First Nations in Canada face an array of governance and participation challenges stemming from Canada's history of colonization and continued colonialism. While beyond the scope of this article to summarize in full, this history has left previously autonomous Indigenous nations with truncated forms of territorial sovereignty and decision-making authority over their communities [4, 18]. The legislative framework

under which most First Nations are governed is the federally written *Indian Act*. Originally devised by representatives of the Crown in 1867, the *Indian Act* allowed for the forcible removal of First Nation peoples from their traditional territories while replacing previous forms of Indigenous governance with colonially-imposed band councils. Under the terms of the *Indian Act*, First Nations are given the ability to democratically-elect representatives for their communities. However, the *Indian Act's* legal framework effectively limits the governing authority of these representatives by granting the federally-appointed Minister of Indian Affairs the authority to approve or disallow many decisions passed by band councils [16]. Furthermore, the original terms of the *Indian Act* provide the Minister with the ability to replace democratically-elected band council representatives and assume authority over First Nations. This has infused an inherent element of instability into First Nation governance.

While the majority of First Nations in Canada continue to operate under the *Indian Act's* legal framework, substantive efforts have been made to revitalize on-reserve governance and improve political engagement. These efforts have come in many forms including amendments to the *Indian Act* itself and intergovernmental negotiations [2, 1, 27]. While many First Nations have replaced the *Indian Act* in its entirety with self-government agreements, most have pursued a more incremental approach. This incremental approach has focused on increasing the autonomy and self-government capacity of First Nations through targeted reforms to specific provisions of the *Indian Act* and negotiated agreements with the federal government that allow for the decentralization of on-reserve services to First Nation governments [27]. These changes have provided opportunities for communities to take iterative steps toward self-government, negotiating with the federal government to gradual develop their governance capacity and assume control over service delivery in areas such as health, housing, education and social services. This incremental approach does not include systemic reform to broader legal and political frameworks. Change instead comes largely through targeted reforms to specific provisions of the *Indian Act* or the negotiation of agreements between individual or collective groups of First Nations. While these changes tend to be smaller and less visible, they represent important evolutions in governance and Indigenous sovereignty often overlooked by post-colonial or neo-Marxist scholars [3, 5].

To date more than 90 Indigenous communities in Canada and the United States have deployed online voting. Decisions to adopt online ballots have been primarily motivated by a desire to improve political participation [14]. Under the terms of the *Indian Act*, however, voting in elections and referendums is permitted only by paper ballot (either in-person at a poll location or by mail) though it is allowed for other types of votes, such as ratification votes or community polls.

3.2. Digital Technology & Online Voting

In recent years, digital technologies have emerged as important tools for First Nations pursuing incremental approaches toward self-government. Digital technologies have allowed First Nations to strengthen governance capacity while addressing social, political and economic challenges. Scholarly literature has explored the use of tech-

nology in the areas of healthcare, education, social services, economic development and cultural renewal [20, 26, 21, 28, 23]. These studies have demonstrated the resourcefulness of First Nations in overcoming digital divides through the creation of innovative funding and ownership models [19]. This body of research has also drawn connections between digital technology and self-determination, exploring the ways in which technology has been used to support broader political goals. For example, several scholars have noted how technology has been used by First Nations to improve administrative capacity to assume greater control over service design and delivery. For First Nations, digital technologies are not only an avenue to improve service delivery or public outreach, but are also understood as tools within a broader decolonizing struggle to roll back the power of settler governments and realize self-determination [22].

To date, few scholars have examined the effects of online voting on Indigenous communities. Research in this context mostly focuses on First Nations in Canada [9, 10, 13] despite the fact that Indigenous communities in the United States and New Zealand are embracing the technology. Existing work has pointed to a number of benefits in Indigenous contexts. Studies have documented improved political engagement through enhanced voter accessibility [9, 10, 13], albeit with small sample sizes, and improved community connectedness by stimulating intergenerational communication among youth and elders [9]. While studies of online voting use in municipal elections in Canada has shown that the voting method can increase turnout by 3.5 percentage points [15], other analyses in comparative contexts that employ similar methodological approaches find no increase [12]. It is unclear whether we could expect the same effect in First Nations as in Canadian municipalities given the unique context and since online voting is used less frequently for Chief and Council elections given legislative limitations and more so for other types of votes (referenda, agreement votes etc.).

Beyond engagement, online voting has also been shown to positively benefit the governance capacity of First Nations. By allowing communities to engage a larger number of their citizens, online voting has supported communities in reaching difficult quorums required to ratify legislation [8, 9]. Further, online voting has also been shown to benefit local governance and strengthen administrative capacity by expediting tabulation of results and by generally including a wider group of community members [14]. Though self-determination has been an important theme in research on online voting in First Nations, it has been relatively underexplored compared to issues of participation and engagement.

Despite benefits covered in the literature, recent research also highlights the unique challenges First Nations face with online voting deployment. First, access to quality broadband is a concern, especially because many communities are located in rural or remote areas [8]. Second, the tendering of online voting contracts to private sector vendors has raised issues especially with respect to data governance and ownership. In addition, the fact that some communities have limited technical resources and capacity has made adequately vetting suppliers more challenging. In response to this a recent report has called for the development of online voting standards to boost technical capacity in First Nations [8]. Finally, it is important to point out that voters'

unfamiliarity with online voting can pose a challenge to implementation. While issues with the novelty of online voting are not unique to First Nations, there is the potential that unfamiliarity coupled with pre-existing feelings of distrust or suspicion toward government may deter uptake.

4. Methodology

4.1. A Community-Engaged Approach

For this study we employed a Community-Engaged Research (CER) approach which seeks to overcome some of the power inequities that exist between researchers and Indigenous communities through the development of research partnerships which promote empowerment, inclusivity, and respect [6, 7]. Such projects share underlying goals of influencing social change, and equitably involving community partners throughout the research process from the inception of the project through knowledge mobilization [24]. Our study employs a qualitative research design, which is considered ethical, respectful, applicable, authentic, beneficial and relevant to the experiences of Indigenous peoples.

We began building a relationship with Wasauksing First Nation in March 2016. This included a number of phone conversations, contributions to community newsletters, attendance and presentations at community events, and presentations to Chief and Council. In addition to receiving university ethics approval, Chief and Council also approved and were supportive of the research. Community members were provided with information about the purpose of the project, data collection process, responsibilities, risks or inconveniences, benefits, assurances of confidentiality and any additional information requested. The practice of gift giving, whether for ceremony or for community events, is common in Indigenous communities [25]. We felt that gift giving was an important part of maintaining positive community relationships and building trust. As a result, we raffled off door prizes at each community meeting, including a number of gift cards and tablets.

4.2. Data Collection

The data for this article comes from semi-structured interviews undertaken with political and bureaucratic leaders in the community, participant observation on voting day, and voter exit surveys. All questionnaires were constructed in active consultation with the community and specific items were added based on their needs.

Interviews asked questions about the process of adopting online voting, challenges and benefits and the role of digital technology in self-determination and self-governance. Notes taken during the interviews and observation were analyzed using NVIVO qualitative data analysis software. Transcripts were uploaded to the software and analyzed using a multi-stage inductive approach, which involved identifying core themes in the transcripts related to community views toward the introduction of online voting. These core themes were then used as coding categories to sort and analyze our interviews with the community. This inductive method of analysis is con-

sistent with the broad CER approach taken in our overall study and allowed for community perspectives and voices to be expressed clearly in the research findings. Our analysis uncovered 3 core themes related to the introduction of online voting: innovation and community modernization, community connectedness, and self-determination and self-governance.

Voter exit surveys were administered to online and paper voters once they had cast a ballot. All persons who voted online or in-person at the polls were offered the option to participate, and completion was voluntary. To support participant recruitment six youth research assistants and one elder interpreter were hired and completed a training course on data collection at the polls. Online voter surveys were administered online. The survey was open for completion during the online voting period, which lasted from December 10th, 2016 until 8:00am on the primary voting day, February 25th, 2017. Paper voters casting a ballot on voting day were encouraged to complete the survey by iPad, but could also fill out a paper copy or complete the survey orally with the assistance of an interpreter. Paper surveys were also offered to voters on December 10th as part of in-person advanced voting.

Survey questions probed voter satisfaction, rationale for use, concerns, likelihood of future use, digital access and literacy, participation histories and standard socio-demographic items. A total of 15 online voter surveys were completed, representing a response rate of 20 percent, and 66 paper voter surveys for a response rate of 66 percent.¹ Respondents self-selected for both surveys and so may have been more likely to like or dislike online voting. Given the self-selected nature of the sample, and the small N's,² descriptive statistics are used to analyze data where appropriate. These limitations prevent us from drawing broad conclusions about online voting's affect on voter engagement in a First Nations context and should be taken as suggestive evidence that could be further explored in future studies.

As part of the knowledge mobilization strategy for this project, Wasauksing was provided with written reports of survey results. Findings were also presented to Chief and Council and the Lands Management committee. These efforts are in line with a CER approach and intended to ensure that the research process and results are meaningful, respectful and relevant, and that they reflect community concerns and interests.

¹ The paper voter sample includes more women (64 percent) than men (34 percent), with 2 percent identifying as 'other'. Paper voter respondents have a median age of 46 years, household income range of \$20,000 to \$29,000, median education level of "some technical community college", and are likely to reside on reserve. The sample of online voters, by comparison, also contains more women (71 percent) than men (29 percent). This sample also reports a median age of 46, household income range between \$80,000 and \$99,000, median education of completed "technical, community college", and are more likely to live off-reserve.

² While the N's are small, they are very good based on the size of the community.

5. Wasauksing First Nation and their Land Code Ratification

Wasauksing First Nation (WFN) is an Ojibway, Odawa and Pottawatomi community located near Parry Sound, Ontario, Canada. The community has a land base of approximately 7,875 hectares and a total population of 1,090 with 369 community members residing on reserve. The community is currently engaged in two land claim negotiations with the federal and provincial governments to extend the boundary line of the reserve to cover an additional 223 hectares of traditional land. The *Indian Act* presently governs elections and referendums in WFN.

In addition to the land claims, WFN has sought to extend its control over its reserve land by signing on to the *Framework Agreement on First Nation Land Management*. The framework agreement was initially signed in 1996 between 13 First Nations and the Minister of Indian Affairs and Northern Development. Since then, it has been ratified as part of the *First Nations Land Management Act* (1999) and expanded to include an additional 125 First Nation signatories. As a signatory, each First Nation is provided the opportunity to develop land code legislation to replace sections of the *Indian Act* related to the governance and management of reserve lands. The agreement is a sectoral self-government agreement that provides First Nations with the legal status and powers to govern and manage their lands and resources through the passage of laws under their own land code.

Once a community has signed on to the framework agreement, they are tasked with developing and drafting their own land code legislation. This legislation covers a number of areas related to land management including general rules and procedures, occupation of reserve lands by members and non-members, financial accountability measures for revenues, and laws and regulations related to environmental protection. Land codes empower communities by setting out the rules and procedures for making and publishing their own land laws, while also diverting funding and fees collected from the administration of land back to the community instead of to the federal government. Passing a land code represents removal from 25 percent of the *Indian Act*.

The penultimate step once a First Nation has drafted land code legislation is to negotiate and sign an Individual Agreement with the Government of Canada. The Agreement establishes the specific terms of the transfer of management of reserve lands from the federal government to the First Nation. After this has been reached, the First Nation must proceed to ratify the proposed Land Code legislation and Individual Agreement by holding a ratification vote. The ratification vote must include all eligible or registered band members aged 18+. Although legislation passed by parliament in December 2018 has softened the requirements to pass a land code, at the time of WFN's vote at least 50% + 1 of registered voters were required to vote, with at least 25% + 1 of all eligible voters casting a yes vote. Historically, meeting this quorum has been difficult [9] and has resulted in failed votes [14].

WFN signed on to the framework agreement in December of 2013 and held a ratification vote on the Land Code in February 2017. Ratification of the Land Code required 178 yes votes from the community's 725 eligible electors. To bolster engagement, WFN decided to offer online voting as a complementary voting method for advanced voting. Paper ballots at the polls and by mail-in were also offered. The final

tally resulted in 191 ballots cast in support of the Land Code and 60 against its passage, Table 1. WFN was successful in ratifying the proposed land code legislation by meeting the required quorum of registered and eligible voters.

Interestingly, 151 ballots (75 internet and 76 mail) were cast remotely, while 100 were cast in person at traditional poll locations. This suggests that remote voting methods were the preferred voting channel for community members and may be important for community engagement. This is explored more fully in the following section.

Table 1. Total Counts in Wasauksing First Nation Land Code Ratification Vote

Votes Received	Yes	No	Totals
Internet	69	6	75
Paper (76 received by mail)	122	54	176
Totals	191	60	251

6. Findings

Examination of interview and participant observation data using NVIVO reveals three prominent themes: innovation and community modernization, community connectedness, and self-determination and self-governance. In the section that follows we discuss the findings according to these three themes and reflect on implications for governance and participation in Wasauksing and the future well-being of the community.

6.1. Innovation & Community Modernization

One key theme identified was the connection between online voting adoption and the modernization of First Nation governance. The deployment of online voting in the Wasauksing Land Code referendum was discussed not just as a one-time novelty to help the community reach quorum, but also as part of a more generalized approach that First Nations are taking to meet citizens' needs and modernize local governance. Most commonly, interviewees discussed the challenge posed by off-reserve residency. Like many First Nations, a large portion of Wasauksing's members reside off-reserve, posing challenges for political involvement such as potentially being less informed and engaged. As one community leader explained:

“Our demographics of our population, that’s a real challenge to us. It might even be a 60/40 or 50/50 that live on-off reserve. We try to maintain all the data on our off reserve members but that’s a challenge because we find historically that, our people move with the seasons. So a lot of people still move that way and so we might have multiple addresses on an annual basis. So that’s one of the challenges, tracking our people that are off reserve and they, to some degree, they hold the big piece of the

mandate. So we want to ensure that we can reach out to them, have a tool or a method of reaching out.”

Community leaders and administrators discussed the role online voting has played to diminish these challenges by reaching citizens off-reserve or whose residency changes frequently. Specifically, officials remarked that online voting provides an effective low-cost avenue to keep community members informed and engaged.

In addition, interview data clearly reveals the connection between political modernization and online voting. For Wasauksing, online voting represents a natural evolution of the community’s engagement with digital technology as a means of adapting governance to the changing realities of member’s lives. As one senior administrator who has worked with land code issues in several First Nations told us, online voting has become standard procedure when undertaking community referenda:

“So currently, every First Nation that is in the development process will be doing e-voting. There are none that are opting out of it. Every one of them is going that route. The way I see it going is that it will replace mail in ballots, we won’t have that cost anymore. So it will just become that. Post land code, I think all the votes on referendums may have some component of in person voting, whether that be a show of hands, because we can do that. But I think e-voting may become the one avenue where all the votes are done through e-voting. So it will just be a more streamlined, economical, efficient way of doing things and so that’s where I see it’s going. But every community that I’ve been working with, they’re all embracing e-voting. It’s not a question of should we do this, it’s just we’re doing it.”

As this quote illustrates, the adoption of online voting is inextricably linked to the community’s vision for future governance. Online voting represents a continuation of local innovation that First Nations have been experimenting with.

However, while the leaders and administrators we spoke with voiced mostly favourable views of online voting, concerns and issues were also raised. Many interviewees discussed concerns similar to what has been observed in non-Indigenous contexts [13] such as breaches in security or privacy posed by potential hacking or interference. Administrators also pointed to confusion and challenges associated with requiring online registration prior to casting a ballot, especially in situations where quorum involves reaching benchmarks of both registered and eligible voters. Comments also focused on the changing nature of voting verifiers in First Nations votes. As one local administrator explained:

“The registration part of it that’s more of challenge because now online, you registered whereas with the paper registration, you had the name of the witness and both signed. So the mindset, not so much of the administrators, not so much of the ratification officers and not so much of the eligible voters, we have an additional person who’s like a scrutineer, their verifier. And so, in their role, they go from reviewing the land code to making sure it’s compliant with the framework agreement to then ensur-

ing that the community ratification process is followed. So then in their minds it becomes “Okay, if there is a challenge based on e-voting, how do I verify that?”

The issues with registration and voter verification stem from administrative confusion that can arise when tallying both paper and online ballots and breaking from traditional processes. None of the administrators we spoke with viewed issues with registration and voter verification as insurmountable challenges, and most stressed that with education and clear communication among administrators, verifiers and electors potential issues could be avoided.

Perhaps the most serious set of issues regarding trust of online voting have nothing to do with the technology, but rather stem from the history of colonialism and Indigenous politics in Canada. Interviewees stated that many citizens who expressed mistrust toward online voting did so on the basis that any change, even one made independent of non-Indigenous governments, may work against the broader political interests of the community and consolidate colonial power dynamics. When explaining this apprehension, one administrator told us:

“I’ve used online voting a few times, there’s not a large response to it and I think that’s because it’s new. I also find that our communities resist change. And that’s probably because in the past in dealing with the government, whenever we agreed to change, it didn’t work to our benefit, it worked to our detriment. So we’re apprehensive about change and even though online voting can be a good thing, it’s a change from the traditional system on how we did it by voting coming in person and voting on paper. So it will take a while for that.”

This fear of introducing online voting highlights the unique political concerns facing First Nations. These are broader concerns that may not be relevant to other jurisdictions, but which First Nations must be acutely aware of and contemplate deeply. These complex and historically situated heuristics fundamentally alter the calculation for enacting any sort of reforms, even one that seemingly extends participation and legal jurisdiction in the community. As one administrator explained:

“That’s right and we don’t really look at change in the way that it could benefit us, such as online voting. That may not be a detriment, it’s change. And so we’re apprehensive, we don’t jump at the chance to change the way.”

These concerns highlight the tension that exists between innovating voting processes with online voting and traditional norms and practices within a First Nation context. The political fissures created by colonialism and threats posed to Indigenous rights lead many First Nation citizens to view political reform with suspicion out of fear that it will lead the community astray from traditional ways of practicing politics. This poses a challenge to the success of online voting use in First Nations, particularly amongst those who may already be disenfranchised or distrustful of governments. As one administrator communicated, face-to-face interaction and the voter experience plays an important role in First Nations:

“I did hear from some community members that weren’t in favour of the use of online voting, that they didn’t like it because one of the big things with voting in person is that it brings the community together and the people get to see each other again and talk and to have that bit of discussion.”

While none of the tensions between modernization and traditional practices were expressed as fatal to the long-term prospects of online voting, it highlighted that online voting should not be viewed as a direct replacement for pre-existing practices and opportunities to participate. Rather, for online voting to be accepted it must be integrated with traditional cultural structures and norms in a supplementary fashion. This was crystallized by Wasauksing’s Chief when speaking about his experiences using social media:

“Social media is one of the tools that we use. Again, some of our approaches to a big change like this, have to go through our structure, our community structure. So again, it’s not so easy when a Chief and Council have an idea and going to a community or the elders, there is a structure in our community and we try to utilize that. That’s going to our elders first and trying to get an idea of what they’re thinking and how they feel and then they infiltrate their families. So that was one of the approaches as well that we do traditionally. I guess it’s my duty, as Chief, to touch base with our elders in regards to that.”

Overall, interviewees positioned online voting as part of the broader modernization of governance practices in the community. Community leaders and administrators were optimistic about the prospects of online voting to improve community connection in local politics.

6.2. Improving Community Connectedness

The second major theme that emerged from interview and participant observation data is the potential for online voting to improve community connectedness by better engaging some members and enhancing community well-being. Improving voter turnout is a commonly cited motivation for the introduction of online voting [12, 13]. For Wasauksing’s leadership, there was hope online voting would improve voter access and enhance participation, but implementation was more focused on better connecting community members by more fully bringing off-reserve members into the policy discussion. Wasauksing’s Chief emphasized this point when asked to reflect on the community’s experience with online voting:

“It gives us that opportunity to be able to connect with them [off-reserve members] and make them feel like they’re part of the reserve and they’re still part of the voting processes and they’re still part of the governance of the reserve and the community and their people. One of the big things with people living off reserve is of course employment and that’s an issue and it doesn’t mean that they don’t want to be part of the reserve, it just mean that for their own financial gains and for their own life circum-

stances, they, for whatever reasons, aren't living on reserve anymore. But it doesn't mean that they don't want to be connected...the online voting, I think that really helped to increase the involvement and the participation of those off reserve members."

In addition, interviewees stressed hopes of augmenting community connectedness with on reserve citizens. In speaking about her experiences with other First Nations, one senior administrator explained that online voting can be a key tool to engage on-reserve members who may not otherwise make it out because they are less engaged. She commented that uptake was equally high among those living on reserve lands:

"On reserve yeah. And you do door-to-door with the mail in ballots or with the VIN numbers or whatever and they attend to participate that way, because these are people who don't come out to council meetings, they don't come out to community meetings, to the annual general meeting. They just don't like to participate or they're intimidated or they're quieter people or they're shy or whatever, this allows them to participate without having to [I: Put themselves out there.] Yeah, yeah. I find, they participate greatly with the online voting but I think it's the on reserve that we've seen a lot of uptake..."

While the assessment of greater uptake by on reserve members is unexpected, it highlights the compatibility and usefulness of technology within a First Nation context.

Looking at Wasauksing's voter survey data, however, shows that paper voters are more likely to reside on reserve while those that chose to vote online are more likely to live off-reserve. In addition, as outlined above, of the 251 votes cast 151 of these were cast remotely – 75 by internet and 76 by mail-in ballot. This suggests that in Wasauksing online voting appeals to off-reserve members and that remote voting methods are important for enabling engagement [14]. This is reinforced by the fact that the primary reasons internet voters cast a ballot online, and why paper voters would consider doing so, include enhanced convenience and accessibility.

In addition, paper voters were asked if they would use online voting in a future vote. Sixty-three percent said they would. Of these, 28 percent said they would do so "in all circumstances" while 35 percent noted they would use it under "special circumstances" such as in instances where they were sick, away, or too busy to make it to a traditional poll location. This signals that online voting may be increasingly important to provide access to, and equality of, the franchise. Put a different way, paper voters were asked how they would prefer to vote if they could not attend a physical poll location. Forty-nine percent of respondents chose online voting, 22 percent voting by mail, 10 percent said they would vote by proxy, 5 percent by telephone and 5 percent reported that they would abstain in such circumstances.

Finally, reflecting on community connectedness by age among both paper and online voters we see that while paper voting is a preferred voting channel for the youngest and oldest voters, online voting had significant uptake from middle-aged voters. This implies that both voting methods are important to connect and engage

community members in votes. Reported satisfaction with paper and online voting emphasizes this point. One hundred percent of online voters reported being satisfied with the online voting process, while 89 percent of paper voters expressed satisfaction with paper voting at the polls. Overall, available survey data from paper and online voters in Wasauksing supports the finding online voting is a tool to engage off-reserve members. It also indicates that online voting is an option paper voters want to see to improve their future voting access, and suggests it could be a means to better connect middle-aged voters with these types of policy discussions.

The importance of community voice was emphasized by the Chief when he reflected that “Chief and Council cannot do it alone”. He noted it was essential citizens shared and discussed information regarding the land vote as a way to increase the collective knowledge about the proposal and extend the number of participating citizens. Overall, for Indigenous communities that face challenges with members located on and off of community lands findings from Wasauksing’s land code experience suggest that the voting method can improve community connectedness and contribute to enhanced community well-being regardless of where members reside. This is consistent with earlier findings published about Wasauksing and other First Nations in Canada [14].

6.3. Self-Determination and Self-Government

The third and final theme that emerged in our analysis is the connection between online voting use and the community’s pursuit of self-determination and self-government. For Wasauksing, understanding the potential gains in self-determination and self-governance by leveraging online voting was a key motivation in partnering with our project. Their goal was not only to understand how online voting could support the community in reaching the quorum to ratify the Land Code, but also to understand whether online voting could be harnessed to create a vibrant and connected citizenry and contribute to long-term sustainable gains in self-determination. As the Chief explained, one of the foremost challenges to enacting self-determination and decision-making is difficulties consulting with community members:

“I know within our nations, I do speak with the regional chiefs and at multiple levels of leadership within the First Nations; we all have issues with reaching out to our off reserve voters. We, as I mentioned, we might be 50/50. I know of communities that are 75/25, 75 living out of their community. So they have real challenges when they need a ratification vote, they don’t have the statistics to hold the data [I: They need a threshold, yeah.]. So they go nowhere in their development, their governance, their development. I guess, to some degree, it’s really, really tough for them to govern and move their communities forward.”

Interviewees viewed online voting as a strategic tool to overcome these challenges, offering an effective and cost-efficient way to foster inclusiveness and enhance connectivity amongst Wasauksing’s membership. Moreover, the use of technologies such as online voting is understood as a viable and necessary pathway to support First Na-

tions in achieving self-determination and seeking recognition of Indigenous rights. For example, as the Chief commented, “The UN Declaration guarantees our inherent rights. We're taking back jurisdiction on many fronts and developing laws, and asserting our rights but we'll need digital tools to do this.” In the context of the land code ratification vote, online voting played a critical role in enabling the community to reach the quorum necessary to ratify the legislation. More significantly, however, the community's ability to deploy online voting and define the terms of its referendum represented an important enactment of self-determination in its own right. This was made clear to us by a senior administrator:

“It's huge...I'll try to go back to the reasons why we've been able to utilize those platforms is because we're not guided by the Indian Act rules for elections or under the referendum regulations. So under the framework agreement under First Nation Land Management, First Nations have really taken on the role of developing their own community ratification process. So that has allowed them to take advantage of all the emerging technology, e-voting. I've been doing this since 2001. So in 2001, I actually did approach Elections Canada and asked them ‘What do you know about e-voting and where's that at?’ ‘Nowhere.’ And they're still nowhere. But First Nations have really been able to take this on and municipalities too.”

Interviewees discussed the role of online voting and technology more broadly in assisting the community in moving out from under the *Indian Act*. For Wasauksing, the Land Code referendum was viewed as a preliminary step toward future gains in self-governance. The impetus to pursue a Land Code *by* the community and *for* the community was made clear by one administrator:

“The motivation is actually very simple, it's jurisdiction. We're not asking for anything new, this is returning to how things used to be before the *Indian Act* was imposed on us before the oversight of Indian Affairs or Aboriginal Affairs or Indigenous Affairs; however you want to call them.”

Broadly, interviews with political and bureaucratic officials in the community illustrate the usefulness of online voting for enhanced self-determination and progress toward self-government in two distinct respects: (1) enabling the passage of legislation that builds community autonomy and capacity, and (2) that the adoption of online voting is an empowering process in and of itself. On the one hand, passing legislation to gain autonomy and move away from governance under the *Indian Act* is a crucial step as part of an iterative approach to self-government. Yet, as outlined, passing such legislation has often been difficult to accomplish with historically imposed quorums. With growing numbers of community members living off-reserve land, connecting members to participate in these processes has become challenging and resulted in failed votes [8]. In this sense, online voting has played a crucial role in connecting community members and enabling their engagement in such votes, without which the passage of these types of community-oriented legislation would not have been possible. Second, the process of adopting online voting as a complementary voting method

and new digital initiative is itself a process of empowerment that enhances self-determination. The process of the community introducing its own chosen tools to enhance its self-determination results in improved capacity on its own.

7. Conclusion

By drawing on community experiences and narratives, this article contributes to our understanding of how online voting affects Indigenous communities, notably the relationship between governance and digital technology. Findings lead us toward optimistic conclusions about online voting positively advancing Indigenous self-determination and community capacity. In the case of Wasauksing First Nation's Land Code referendum, the community accrued significant capacity to reach out to and involve its membership by leveraging online voting. In addition, community leaders and administrators identified online voting as a strategic tool essential in pursuing an iterative approach to self-government.

Wasauksing's experiences with online voting mirror many of the findings and discussions in the scholarly literature focused on notions of digital decolonization and self-determination [22]. For the community, the land code vote represented an opportunity to decentralize authority and decision-making over lands from the federal government. The use of online voting helped the community ensure that the land code process was inclusive by promoting consultation with members living both on and off-reserve. While there are tensions and challenges with online voting use in a First Nation context, our engagement with Wasauksing finds no evidence of online voting conflicting with traditional norms and decision-making practices and its adoption in no way superseded important cultural protocols of community consultation. Rather, the technology was introduced in a manner consistent with the pre-existing decision-making structures that centered on community deliberation and openness. The success of online voting in this case is largely due to the fact that the community dictated the terms of its introduction. This ensured the voting method would be deployed in a manner consistent with the community's broader political goals and that members could be educated and socialized toward the technology.

While these findings are not meant to argue that online voting can be successful in all First Nations, they suggest that by employing an appropriate approach online voting can be a useful tool in the pursuit of Indigenous self-determination. They also emphasize what previous research has underscored about the potential to engage off-reserve membership [14] and the extent to which this enhances community capacity. While some research has explored good practices regarding Indigenous deployment of online voting [8, 13] additional studies are needed to determine the conditions and steps which best ensure online voting serves the interests of First Nations, particularly in other jurisdictions such as the United States or New Zealand. Future research could also more systematically examine how online voting affects voter turnout in Indigenous communities and whether these differ from findings of local government elections [15, 12]. Finally, comparative assessments of how online voting impacts Indige-

nous governance and self-determination in other contexts could support the results of this article.

Acknowledgement. All authors contributed equally to this research. Authorship is listed alphabetically. We extend our deep thanks to Wasauksing First Nation for taking part in this research. Research undertaken for this article was financially supported by the Social Sciences and Humanities Research Council of Canada and Chelsea Gabel's Canada Research Chair in Indigenous Well-Being, Community-Engagement and Innovation.

References

1. Abele, F., & Prince, M. J. Four pathways to Aboriginal self-government in Canada. *American Review of Canadian Studies*, 36(4), 568-595 (2006)
2. Alcantara, C., & Davidson, A. Negotiating Aboriginal Self-Government Agreements in Canada: An Analysis of the Inuvialuit Experience. *Canadian Journal of Political Science*, 48(03), 553-575 (2015)
3. Alfred, T., & Corntassel, J. Being Indigenous: Resurgences against contemporary colonialism. *Government and Opposition*, 40(4), 597-614 (2005)
4. Asch, M. *On being here to stay: Treaties and Aboriginal rights in Canada*. University of Toronto Press (2014)
5. Coulthard, G. S. *Red skin, white masks: Rejecting the colonial politics of recognition* (2014)
6. Dickson, G., Green, K.L. Participatory action research: Lessons learned with Aboriginal grandmothers. *Health Care for Women International* 22, 471-82 (2001)
7. Ermine, W., Sinclair, R., Jeffery, B. *The ethics of research involving Indigenous peoples*. Saskatoon, Saskatchewan: Indigenous Peoples' Health Research Centre (2004)
8. Gabel, C., Goodman, N. *Indigenous Experiences with Online Voting*. Report (2019)
9. Gabel, C., Goodman, N., Bird, K., Budd, B. Indigenous adoption of internet voting: A case study of Whitefish River First Nation. *International Indigenous Policy Journal*, 7 (2016)
10. Gabel, C., Goodman, N., Bird, K., Budd, B. The impact of digital technology on First Nations participation and governance. *The Canadian Journal of Native Studies* 36, 107-127 (2016)
11. Gabel, C. *Towards Healthier Aboriginal Health Policies? Navigating the Labyrinth for Answers*. PhD Dissertation, McMaster University (2012)
12. Germann, M., Serdült, U. Internet voting and turnout: Evidence from Switzerland. *Electoral Studies*, 47, 1-12 (2017)
13. Goodman, N., Pyman, H. *Understanding the effects of internet voting on elections: Results from the 2014 Ontario municipal elections*. Technical Paper, Centre for e-Democracy (2016)
14. Goodman, N., Gabel, C. and Budd, B., *Online Voting in Indigenous Communities: Lessons from Canada*. In *International Joint Conference on Electronic Voting*, Springer, Cham, 67-83 (2018)
15. Goodman, N., Stokes, L. C. Reducing the cost of voting: an evaluation of internet voting's effect on turnout. *British Journal of Political Science*, 1-13 (2018)
16. Imai, S. *The Structure of the Indian Act: Accountability in Government*. Research Paper for the National Centre for First Nations Governance. Ottawa, ON (2007)

17. Ladner, K. L. Understanding the impact of self-determination on communities in crisis. *International Journal of Indigenous Health*, 5(2), 88-101 (2009)
18. Ladner, K.L. Political genocide: Killing nations through legislation and slow-moving poison. In *Colonial Genocide in Indigenous North America* eds. Alexander Laban Hilton, Andrew Woolford and Jeff Benvenuto. Duke University Press, Durham NC (2014)
19. Lightfoot, S. *Global Indigenous Politics: A Subtle Revolution*. Routledge. (2016)
20. Lockhart, E., Tenasco, A., Whiteduck, T., O'Donnell, S. Information and communication technology for education in an Algonquin First Nation in Quebec. *The Journal of Community Informatics*, 10(2) (2013)
21. McMahon, R., LaHache, T., Whiteduck, T. Digital data management as Indigenous resurgence in Kahnawà:ke. *International Indigenous Policy Journal*, 6(3) (2015)
22. McMahon, R. From Digital Divides to the First Mile: Indigenous Peoples and the Network Society in Canada. *International Journal of Communication*, 8, 25 (2014)
23. McMahon, R., Gurstein, M., Beaton, B., O'Donnell, S., Whiteduck, T. Making information technologies work at the end of the road. *Journal of Information Policy*, 4, 250-269 (2014)
24. Minkler, M., Wallerstein, N. *Community based participatory research for health*. San Francisco, CA, US: Jossey-Bass (2003)
25. Moore, C., Castleden, H. E., Tirone, S., Martin, D. Implementing the Tri-Council Policy on Ethical Research Involving Indigenous Peoples in Canada: So, How's That Going in Mi'kma'ki? *The International Indigenous Policy Journal*, 8(2), 4 (2017)
26. O'Donnell, S., Beaton, B., McMahon, R., Hudson, H.E., Williams, D., Whiteduck, T. Digital Technology Adoption In Remote And Northern Indigenous Communities In Canada. *Canadian Sociological Association Annual Conference*. University Of Calgary, Calgary, Alberta, June (2016)
27. Papillon, M., Bakvis, H., Skogstad, G. *Canadian federalism: performance, effectiveness, and legitimacy*, Oxford University Press, 284-301 (2012)
28. Sweet, M., Pearson, L., Dudgeon, P. IndigenousX: A case study of community-led innovation in digital media. *Media International Australia*, 149(1), 104-111 (2013)

How increasing use of Internet voting impacts the Estonian election management

Iuliia Krivososova^[0000-0001-7246-1373], Radu Antonio Serrano Iova^[0000-0003-2183-0313], David Duenas-Cid^[0000-0002-0451-4514], and Robert Krimmer^[0000-0002-0873-539X].

Tallinn University of Technology, Akadeemia tee 3, 12618 Tallinn, Estonia

{iuliia.krivososova, rasser,
robert.krimmer, david.duenas}@taltech.ee

1 Introduction

Despite the initial expectations that digital divide would create a barrier for wide-scale Internet voting [1], the usage of Internet voting in Estonia has been steadily growing in relative and absolute numbers, reaching a record during the 2019 Parliamentary elections. 44% the votes were casted over the Internet, becoming the most popular voting channel for the first time ever, supporting the claim that Internet voting is habit forming being “more persistent and repetitive than paper voting or non-voting” [5]. However, the growing popularity of Internet voting has raised concerns among stakeholders [2]. This paper aims at highlighting a set of impacts which growing usage of Internet voting has on the election administration, that will be explored in further research.

2 Practical implications of growing usage of Internet voting

- **Impact on security:** The growing usage of Internet voting raises the stakes for interference in elections. That is an underlying assumption of many Internet voting systems. Over years, the system of Internet voting in Estonia has already evolved significantly, with many actors and properties being added [3]. However, unlike some countries which have legal requirements in place to reflect on the increasing usage of Internet voting (e.g. in Switzerland), Estonia has never mentioned them explicitly, hence, representing the alternative approach of regulating less to guarantee the possibility for further innovations.
- **Impact on delivery of paper-based voting channels and financing of elections:** Growing usage of Internet voting inevitably results in decreasing popularity of paper-based voting channels (see Table 1), hence, their lower cost-efficiency in comparison to Internet voting. Previous research [4] reveals that the cost difference between paper and electronic voting channels in Estonia in some cases reaches 10 times. This opens the discussion on optimization of supply of polling stations and variety of voting channels. We present an overview of how supply of polling stations changed since introduction of Internet voting and what the reasons are behind this decision.

Table 1. Votes’ distribution among voting channels in 2017 Local elections and 2019 Parliamentary elections.

Voting channel	2019	2017
Internet voting	43,8%	31,4%
Election day voting	37,4%	47,2%
Advance voting	13,8%	15,7%
Early voting	3,8%	4,6%
Home voting	0,9%	1,1%
Voting in diplomatic missions	0,25%	n/a
Postal voting	0,05%	n/a

- **Impact on supporting electoral infrastructure:** The Estonian electoral system is a combination of a modern Internet voting channel and a set of analog voting channels. Apart from Internet voting, Estonia does not use other sophisticated electoral technologies: 1) paper votes are counted manually; 2) voting machines have never been used at polling stations; 3) candidate registration is not automated, and 4) paper voter lists are manually marked-off. Thus, Internet voting system is surrounded by analogue electoral processes, and consolidation of Internet voting with paper-based voting channels is frequently done manually. While manual consolidation was not seen to be a problem when the number of Internet voters was small, growing usage of Internet voting challenges the further feasibility of manual consolidation. We consider a process of manual consolidation in detail, and reflect on measures taken to address the growing need for update of supporting electoral infrastructure.

Acknowledgments

This work received support from ETAG personal research grant 1361.

References

1. Alvarez, R M, and T E Hall. 2003. Point, Click, and Vote: The Future of Internet Voting Point, Click, and Vote: The Future of Internet Voting.
2. Heiberg, Sven, Peeter Laud, and Jan Willemson. 2012. “The Application of I-Voting for Estonian Parliamentary Elections of 2011.” In *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 208–23.
3. Heinsalu, Alo et al. 2016. Elections in Estonia 1992-2015.
4. Krimmer, Robert et al. 2018. “How Much Does an E-Vote Cost? Cost Comparison per Vote in Multichannel Elections in Estonia.” In *Lecture Notes in Computer Science*, 117–31.
5. Solvak, Mihkel, and Kristjan Vassil. 2018. “Could Internet Voting Halt Declining Electoral Turnout? New Evidence That E-Voting Is Habit Forming.” *Policy and Internet* 10(1): 4–21.

Analysis of Deployed Systems

Technical and Socio-technical Attacks on the Danish Party Endorsement System

Carsten Schürmann and Alessandro Bruni

Center for Information Security and Trust
IT University of Copenhagen
{carsten | brun}@itu.dk

Abstract. In this paper we analyze the security of the online Danish party endorsement system (DVE) and present two attacks: one technical, which we discovered during our study of the system in 2016 and which compromises the integrity of the endorsements stored in the DVE-database and another socio-technical, which allows parties to circumvent mechanisms to protect voters against abuse. To understand these attacks, we introduce the legal and technical frameworks of the DVE-system, analyze its problems, and describe a sequence of events that has led to endorsing three new parties that stood in the 2019 Danish Parliament election.

1 Introduction

Collecting signatures to endorse parties running for parliament is an important part of the democratic process, and for Denmark it is no exception. Prior to March 29, 2014, all endorsements were done on paper. Prospective parties had to go out to shopping malls and libraries, talk to people, and collect hand-written signatures. These signatures were subsequently checked by the municipalities for eligibility and correctness, i.e. that each endorser had the right to endorse and had not already endorsed another party. This was tedious work.

On March 29, 2014, the law governing party endorsements was changed to allow a digital solution to be used for collecting endorsements. The Ministry for Social Affairs and the Interior (in Danish “Social- og Indenrigsministeriet” or SIM, as it was known at the time) would become responsible for providing, running, and maintaining the solution. The prospective parties would be responsible for collecting endorsements using the solution. Prior to the law change, SIM was not involved in the operational aspects of collecting and checking endorsements, but with the new law in place the ministry had to provide a digital solution assisting parties to collect endorsements.

The law requires that to stand for Danish Parliament election, a prospective party must collect at least a number of endorsements equal to $1/175$ of all valid votes cast in the previous Danish Parliament election, which corresponds to 20.109 endorsements in 2019. In comparison, to be admitted to the European Parliament election, a prospective party must collect more than 2% of all valid

votes cast in the previous Danish Parliament election, which corresponds to 70.380 endorsements in 2019.¹

Not long after the law changed, SIM released the corresponding administrative regulation and commenced with the procurement of the Danish party endorsement system (DVE). A feasibility study was conducted, requirement documents were drafted, and the Danish company KMD was tasked with developing the system.² Eventually, the DVE-system was deployed after some delay, and voters were invited to endorse prospective parties.

From an operational point of view, the DVE-system works as follows: any new prospective party that strives to be recognized collects email addresses of potential endorsers and then forwards invitations to endorse the party through the DVE-system. Eligibility is checked using Denmark’s national digital identity system, which allows the DVE-system to verify information about the endorser, such as nationality and age. Moreover, the DVE-system allows endorsers to withdraw their endorsements at a later point in time. The system also supports the collection of paper-based endorsements.

In this paper, we discuss some challenges that accompanied the introduction of the DVE systems in the Danish electoral process. We first describe the legal and technical framework of the DVE-system in Section 2, then present best practices following recommendations of international and non-governmental organizations in Section 3. In Section 4, we show two attacks against the 2016 version of the DVE-system, one technical and one socio-technical. In Section 5, we present our reflections and unsolicited recommendations, which we consider important for a future DVE-system. Finally, in Section 6, we conclude and discuss briefly the impact of the DVE-system on the 2019 Danish Parliament election.

2 The DVE Framework

2.1 The Legal Framework

In Denmark, every resident has a digital ID, called NemID³, a personal identification number, which is called CPR⁴, and access to an authenticated email service, called e-Boks⁵ (Digitale Post). With the availability of these technologies, Denmark’s Parliament passed a law in March 2014 that would allow the collection of party endorsements using a digital online system.⁶ SIM, which changed its name to Ministry of Economic Affairs and the Interior (Økonomi- og Indenrigsmin-

¹ See European Parliament election law §11, Section 1.

² Neither the authors nor their affiliations were involved in this project, neither during procurement, nor development, nor quality assurance.

³ See <https://www.nemid.nu>

⁴ See <https://cpr.dk/>

⁵ See <https://www.e-boks.com/danmark/da>

⁶ See <https://www.ft.dk/samling/20131/lovforslag/1124/index.htm>

isteriet, ØIM) in 2016, was ordered to develop the appropriate administrative regulation, which was published in January 2016.⁷

Compared to the early versions of the legal framework, this administrative regulation focused on the usability of the digital solution. To this end, the authority to check the validity of the endorsements and the eligibility of the endorser was transferred away from the municipalities to the DVE-system, relieving the municipalities of the need to procure their own local systems.

The new regulation defines necessary roles in connection with the digital solution, among them the “party administrator”, who works on behalf of the prospective party to drive the endorsement collection activity. The regulation also uses several technical terms, such as “granting access to the digital solution”, “is contacted by email”, “the email contains a link”, “a voter declaration”, “voter’s declaration’s key”, “archive key” etc. which are underspecified and can be interpreted in many different ways. In addition, the new regulation explicitly requires that seven days must pass between the time of registration to the time of endorsement, granting the prospective endorser time to reflect upon whether or not to endorse the prospective party. The new regulation also prescribes what must happen in case the digital solution is not functional or ceases to work.

We observe that the new regulation does not make any reference to the verifiability of the operation of the DVE-system, the accuracy of the endorsements, the integrity of the database, the confidentiality of the endorser, nor the availability of the DVE-system.

2.2 The Technical Framework

The steps to endorse a party using the DVE-system are depicted in Figure 1. The process is driven by the prospective party. In a first step, the party has to apply with ØIM to initiate the endorsement collection process, and ØIM configures the DVE-system accordingly. Prospective voters can be asked to endorse a new party either by email or by regular mail. As nearly all of the Danish residents have a digital ID, only few will endorse parties by regular mail (this case will not be considered any further in this paper).

The legal framework requires that the party should initiate the process (Init) by sending the email address of the prospective voter to ØIM through the online interface of the DVE-system. Once received, the email address will be stored for seven days (which is referred to as the “period to reflect”) before sending an invitation (Invitation) email to the endorser containing the link to finalize the endorsement and the token to identify it.

After receiving the email, the endorser can follow the link to log into the DVE-system through the website vaelgererklæring.dk using the digital ID NemID. By confirming the information on the screen, the endorsement is then given (Endorsement). The DVE-system will subsequently send a receipt (Receipt) to

⁷ See Vælgererklæringsbekendtgørelsen, <https://www.retsinformation.dk/Forms/R0710.aspx?id=176933>

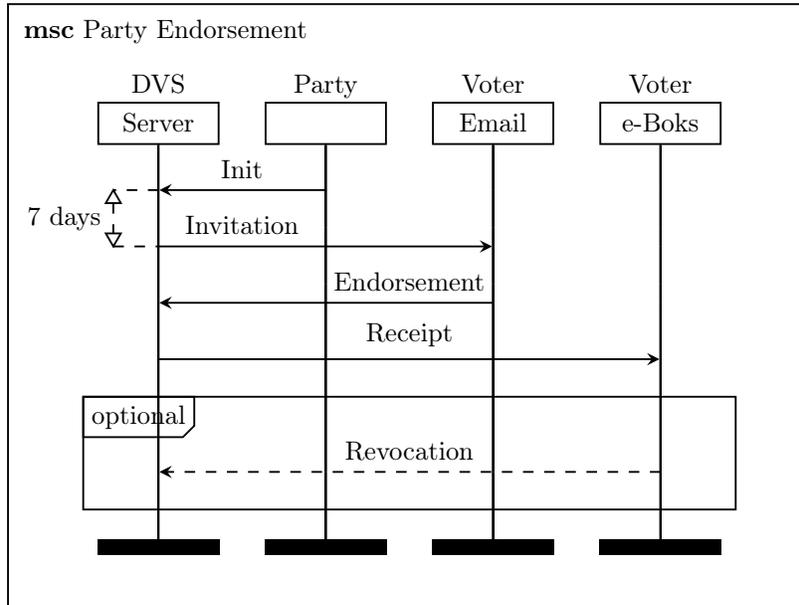


Fig. 1. Process of endorsing a party

the endorser via the authenticated email service e-Boks including a link that give the endorser the possibility to withdraw the endorsement (Revocation).

The DVE-system was implemented in 2016 by the Danish company KMD, and since then has been used by more than 100 prospective parties, four of which applied for recognition, with the result that three applications were granted and one was rejected.

3 Best Practices

In this section we describe our reflections on best practices, before analyzing the legal and technical framework of the DVE-system in the subsequent sections.

Although party endorsement is usually considered secondary to electronic voting, it exhibits all defining characteristics of an Internet voting system: all endorsements are submitted through the Internet; there is no paper trail; the integrity of the database (of all endorsements) is instrumental for creating public confidence in the decisions that follow from it; all voters who are eligible to vote in the election are also eligible to endorse upcoming parties; finally, voters can only endorse one party. Newly approved parties will appear on the ballot and there is likely extensive media coverage of such new parties. Therefore, the DVE-

system qualifies as an election technology that should live up to best practices and standards of any election technology.⁸

Any election technology must be *software independent*, which means that an undetected error in the system software cannot lead to a undetected error in the election outcome [RW06]. One way to achieve software independence is through *verifiability*, which can increase the level of trust in voters, by checking that a system performed correctly by secondary means [BRR⁺15]. The election technology must be *secure*, which means that it can be used safely in adversarial environments, and *accountable*, which means that there are mechanisms to identify misbehaving participants.

4 Findings

We have been following the development of the legal and the technical framework of the DVE-system attentively since 2014 and in particular during its launch in 2016. In this section we share our reflections and observations.

4.1 Legal Framework

Lack of Verifiability The administrative regulation lacks mention of basic requirements usually associated with the introduction of electronic voting technologies, which is contrary to international standards and best practices. The regulation does not mention requirements regarding security, accountability, software independence, or verifiability. Consequently, it is impossible to judge whether a DVE-system is compliant with the regulation, which is particularly worrisome.

A grave oversight of the regulation is that essential elements of the DVE-system are not adequately defined, including what constitutes a *binding endorsement*. The regulation fails to specify how endorsements should be linked to the voter (if at all), and how they are secured. The regulation also omits which mechanisms must be put in place to guarantee the authenticity of an endorsement, and to protect them against copying, tampering, and deletion.

Absence of Endorsement Privacy It is also noteworthy that the regulation fails to take advantage of new opportunities to guarantee higher levels of endorsement privacy. The role of vote privacy in Internet voting systems is well understood as are the mechanisms and the challenges to protect it. We believe that the reason why endorsement privacy is not considered in the regulation is most likely because prior versions of the legal framework did not require it either. Operationally speaking it is not clear how to collect party endorsements on paper, while at the same time verifying the eligibility of the endorser and guaranteeing their privacy.

⁸ See Recommendation CM/Rec(2017)5[1] of the Committee of Ministers to member States on standards for e-voting (Adopted by the Committee of Ministers on 14 June 2017 at the 1289th meeting of the Ministers' Deputies)

The Impossibility to Revert Before the law was changed and the administrative regulation was released, the municipalities had the authority and responsibility to verify the validity of endorsements and the eligibility of the endorser. By transferring this responsibility to the digital solution run by ØIM, the legal framework has forfeit the possibility to revert to a manual process, which would have to be, presumably, carried out by the ministry, which simply does not have the resources to do so. The legal framework therefore did not only create the legal foundation for the use of a digital solution, it made its use mandatory.

Transparency The legal framework refers to the digital solution, as if it was already constructed and as if under no circumstances it could negatively affect the public confidence of endorsers and voters. The administrative regulation does not explicitly require any quality assurance processes (for example software reviews or penetration tests) nor certification with respect to applicable standards (e.g. ISO). It also does not require the source code of the system nor any evaluation or penetration testing reports to be public [EBB⁺15].

4.2 Technical Framework

The protocol depicted in Figure 1 is implemented in the DVE-system. As, apparently, there is no concern to run the DVE-system in a highly contested adversarial environment, it is perhaps not surprising that the implementation that we studied in 2016 did little to defend against cyberattacks and socio-technical attacks, which we discuss next.

Cyberattacks The design documentation and the source code of the DVE-system is closed source, contrary to best practice and international recommendations [CKN⁺14]. It is therefore impossible for any outsider to conduct a security analysis of the DVE-system, in its entirety, i.e. review of requirement documents, design documents, implementation, and deployment plans. Collaborating closely with officials at ØIM, we were asked to endorse a fictional test party (which in our eyes required a limited security analysis, pretending to be an ineligible adversary who was invited by a party administrator to endorse). We submitted our email-address to our contact at ØIM, who set up a test party, and invited us to endorse. By they time we tried to endorse, we observed that the invitation had already expired. Upon request, our contact in the ministry forwarded a second invitation, this time not for the test party but for a real and non-fictional party “De Visionære”. The invitation message is depicted in Figure 1.

Architecture. The outward facing interface of the DVE-system is a website that runs a Java application server and is connected to NemID for the purpose of authentication. Once authenticated, a set of JavaScript files are uploaded to the client computer and the application is then executed client-side.

Following the exchange of messages described in Figure 1, the party owner sends an invitation to the endorser by the way of the DVE-system, who receives

a link and a token (a random looking string of characters) that initiates the endorsement procedure. To endorse, the recipient of the mail follows the link, authenticates via NemID and confirms the party endorsement.

The website follows a three-step process. The first step checks whether the endorser has the proper rights to endorse the party: endorsers can only endorse one party and they must have the right to vote in the election the prospective party stands for; the second step requires confirmation from the voter, showing the party details and the voter identity; finally, the third step returns feedback on the success or failure of the process. Because endorsement invitations and endorsers are not linked, the requirement that seven days must have passed between the initiation of the process and the actual endorsement (“reflection period”) is not checked by the DVE-system.

Adversary model. An election is a contest, where different stakeholder groups, often with conflicting interests, stand as candidates before an eligible voting population to determine who will go in power. Public confidence in the election outcome is paramount.

The party approval process is similar, except perhaps that there is not one but many winners. The adversary model is therefore relatively well defined. It includes every individual who tries on behalf of a party to influence the approval process in any form. Any individual can get an invitation to endorse a party, by contacting the party administrator directly, and therefore any user of the DVE-system must be considered an adversary, which is exactly the adversary group that we focus on in this paper.

In the bigger picture, there are of course other adversaries that must be taken under consideration, including insider attackers (working for example for ØIM or the system vendor), generating spurious endorsements or removing them to influence the result. There are also Nation States who aim to influence the approval process, for example, with the objective to destabilize a country or influence the public discussion on media. Unfortunately, because of ØIM’s policy to keep all information confidential, our security analysis is restricted to an individual attempting to endorse a party by breaking client-side security.

Client-side security refers to the capability of the end-points of a larger distributed system to defend against cyberattacks by keeping adversaries out, protecting the integrity of the data, the confidentiality of sensitive information, and the overall availability of the system. ØIM should be concerned with client-side security because it grants endorsers direct access to sensitive data by the way of end-points, such as laptops, mobile phones, or tablet computers. Weak or no client-side security allows adversaries to gain access to sensitive data, i.e. party endorsement, corrupt the integrity of the databases, or, in the worst case, disrupt service altogether.

Objective of the security analysis. According to Danish law, citizens are permitted to endorse parties for the Danish Parliament, EU-citizens are permitted to endorse parties for the European Parliament elections, and every individual is only allowed to endorse one and only one party. The DVE-system should enforce

this policy and the objective of the security analysis is to determine if it really does. Unfortunately, this was not the case.

Detailed description of an identified vulnerability. On August 12th 2016, 9:00 the authors of this paper commenced the security analysis of the DVE-system by studying the client-side security of the website vælgereklæring.dk. Figure 2

Fra: noreply@vaelgererklaering.dk
Dato: 3. august 2016 kl. 02.02.27 CEST
Til: [redacted]
Emne: Link til at afgive vælgereklæring



Kære [redacted]

Du har tilkendegivet at ville afgive en vælgereklæring til et parti, der søger at blive opstillingsberettiget til Folketingsvalg

For at afgive din vælgereklæring til partiet skal du klikke på linket nedenfor. Hvis linket ikke er klikbart, skal du kopiere det og indsætte det i adresselinjen (øverst) i browseren. Vil du alligevel ikke afgive en vælgereklæring til partiet, eller har du ikke selv oplyst din e-mailadresse, kan du benytte linket til at trække støttetilkendegivelsen tilbage og få slettet oplysningerne om dig. Benytter du ikke linket inden datoen nedenfor, vil alle oplysninger om dig automatisk blive slettet. Ønsker du herefter at afgive en vælgereklæring, skal du henvende dig til det parti, som du vil støtte.

Klik på linket for at komme til den hjemmeside, hvor du kan afgive en vælgereklæring:

<https://www.vaelgererklaering.dk/apos2/eve/vaelger?uuid=2c4e0be1-60b9-4a60-9c98-5f5d0890cb48>

Linket er aktivt til-og-med 30/08/2016

Venlig hilsen
Social- og Indenrigsministeriet

Fig. 2. Invitation Email with token (sensitive information redacted)

depicts the redacted version of the Invitation email that was forwarded to us on August 11, 2016, 15:20, only three hours after our initial request on August 11, 12:10. It is also noteworthy, that this email was not sent to the email address we provided, but to the email of an employee at ØIM (redacted), which means that the ØIM staff was aware that tokens are not bound to the identity of the endorser, but can be forwarded to anyone, effectively circumventing the seven day “reflection period”. The token is clearly displayed as the last argument to the clickable link in the email.

Upon receiving this token by email, we followed the link, and were prompted with the NemID login procedure, which we used to authenticate the first author of this paper. At this point, it was time to launch the tools that any security expert and hacker would use.

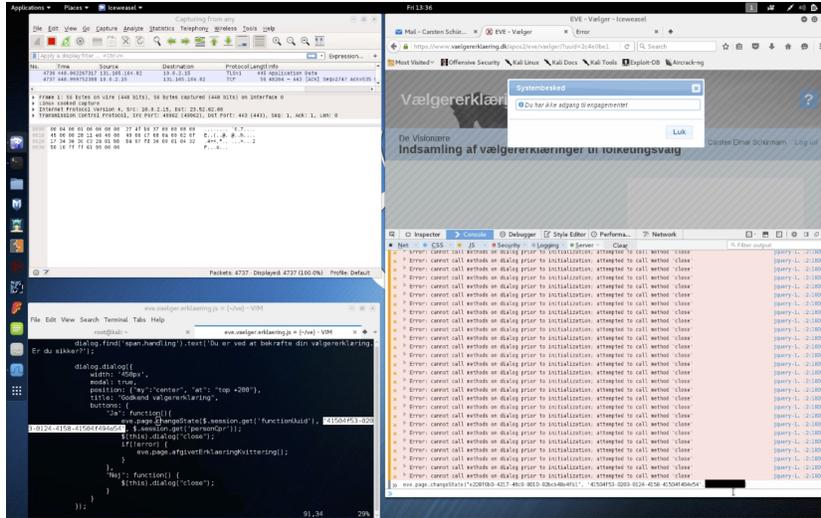


Fig. 3. Screen displaying the attack (sensitive information redacted)

Figure 3 shows a screen shot of our system, running Kali Linux⁹. The window in the upper left corner runs Wireshark¹⁰, a tool that records all network traffic and analyzes it. The window below depicts an editor window showing parts of the JavaScript code application that our machine has downloaded from the website vælgererklæring.dk and which we used to review and understand the code. For example, the highlighted number '41504f53-0203-0124-4158-41504f494e54' is the code word that tells the DVE-system that the endorsement was confirmed by the endorser and is ready to be submitted. To the right, we see an instance of the Firefox¹¹ client window in split screen with the web developers tools running in the lower part.

Our first step was to collect information on how the DVE-system works, so we logged into the website using the first author’s credentials, and proceeded to try to endorse the party “De Visionære”. We observed that the system worked according to specification: as the first author is not a Danish citizen, the system rightly rejected his attempt to endorse the party, which was acknowledged by an “access denied” message on screen. With this information at hand, we started to study the JavaScript code that was running on our computer to learn as much as we could about the implementation.

1 `page.find('a.button').bind('click', function(e){`

⁹ See <https://www.kali.org/>

¹⁰ See <https://www.wireshark.org/>

¹¹ See <https://www.mozilla.org/en-US/firefox/new/>

```

2 e.preventDefault();
3 $.ajax({
4   type: 'GET',
5   url: '!checkVoterRightsProxy',
6   dataType: 'xml',
7   async: false,
8   cache: false,
9   data: {
10    cpr: $.session.get('personCpr'),
11    type: $.session.get('indsamlingsType')
12  },
13  success: function(xml){
14    if($(xml).find('responseMessage').attr('response')
15      == "success"){
16      eve.page.confirmAfgivErklaering();
17    }else if($(xml).find('responseMessage').attr('response')
18      == "dobbelterklaering"){
19      $.session.set('afvistGrund', 'dobbelt');
20      eve.page.afvistErklaeringKvittering();
21    }else if($(xml).find('responseMessage').attr('response')
22      == "valgret"){
23      $.session.set('afvistGrund', 'mangledeValgret');
24      eve.page.afvistErklaeringKvittering();
25    }else{
26      eve.core.confirmationsDialog(
27        $(xml).find('responseMessage').attr('response'));
28    }
29  }
30 });
31 });

```

This piece of code contacts the `!checkVoterRightsProxy` to check if the user has the proper rights to vote. On line 13 we can see how the client reacts to possible responses by the server. There are three cases:

1. all the checks succeed (`response == 'success'`), so the user is allowed to endorse the party; in this case `eve.page.confirmAfgivErklaering()` is called by the client;
2. the user has already endorsed a party (`response == 'dobbelterklaering'`), hence the user is hence denied the option to cast an endorsement;
3. or the user does not have the right to vote at all, (`response == 'valgret'`), and is similarly denied to cast an endorsement for the party.

Since the first author is not Danish, the request is clearly handled by case 3. and so we inspected the code for implementing the necessary functionality. First, the following function is called:

```

1 afvistErklaeringKvittering: function(){ //V9
2   eve.core.clean($('#vContent'));
3   eve.core.clean($('#vInfo'));

```

```

4   eve.page.renderAfvistErklaeringKvittering();
5   eve.page.configHelpButton('V9');
6 },

```

which then in turn calls:

```

1 renderAfvistErklaeringKvittering: function(){
2   var page = $('#vContent'), information = $('#vInfo');
3
4   $.ajax({
5     type: 'GET',
6     url: '!rejectVoterProxy',
7     dataType: 'xml',
8     async: false,
9     cache: false,
10    data: {
11      uuid: $.session.get('functionUuid'),
12      cpr: $.session.get('personCpr')
13    }
14  });
15
16  page.append(eve.page.breadcrumbs(3, 3);)
17
18  var indsamlingsNavn =
19    $.session.get('indsamlingsType') == "FV" ?
20    "folketingsvalg" : "Europa-Parlamentsvalg" ;
21  var partyNameLink = $.session.get('www') != "" ?
22    '<a href="' + $.session.get('www') + '" target="_blank">'
23  + 'Link til ' + $.session.get('navn') + ' hjemmeside</a>' :
24    "" ;
25
26  [...]
27  page.find('a.button').bind('click', function(e){
28    e.preventDefault();
29    eve.page.logout();
30  });
31 },

```

This function contacts the server calling `!rejectVoterProxy`, then constructs a return HTML message using a sequence of appends, which we have abbreviated under `[...]` on line 25. Now we understood and could reconstruct the behavior of the DVE-system. The next step therefore was to try to trick the DVE-system into accepting the first author's endorsement despite him not being a Danish citizen. We traced the code to determine what would have happened if all checks succeeded and inspected the `confirmAfgivErklaering` function, which displays the confirmation of a vote, and it looked as follows:

```

1 confirmAfgivErklaering: function(){
2   var dialog = $('#confirm_dialog');
3

```

```

4   dialog.find('span.handling').text('Du er ved at bekræfte' +
5     ' din vælgererklæring. Er du sikker?');
6
7   dialog.dialog({
8     width: '450px',
9     modal: true,
10    position: {"my": "center", "at": "top +200"},
11    title: "Godkend vælgererklæring",
12    buttons: {
13      "Ja": function(){
14        eve.page.changeState($.session.get('functionUuid'),
15          '41504f53-0203-0124-4158-41504f494e54',
16          $.session.get('personCpr'));
17        $(this).dialog("close");
18        if(!error) {
19          eve.page.afgivetErklaeringKvittering();
20        }
21      },
22      "Nej": function() {
23        $(this).dialog("close");
24      }
25    }
26  });
27 },

```

Line 13 displays the function call that gets executed once the endorsement is confirmed, which is triggered by the user pressing the “Ja” button as the last step. In particular, the following function call is interesting:

```

1  eve.page.changeState($.session.get('functionUuid'),
2    '41504f53-0203-0124-4158-41504f494e54',
3    $.session.get('personCpr'));

```

Digging deeper into the JavaScript code, we observe that the function being called is the following:

```

1  changeState: function(funktionUuid, stateUuid, cpr){
2    var data = {funktionUuid: funktionUuid,
3              stillingsbetegnelseUuid: stateUuid}
4
5    $.ajax({
6      type: 'GET',
7      url: '!changeVoterStateProxy',
8      dataType: 'xml',
9      async: false,
10     cache: false,
11     data: data,
12     success: function (xml) {
13       if($(xml).find('responseMessage').attr('response')
14         != "success") {
15         eve.core.confirmationsDialog($(xml).find('
16           responseMessage'))

```

```

16     .attr('response'));
17     error = true;
18   } else {
19     error = false;
20   }
21 }
22 });
23 },

```

The function contacts the server on `!changeVoterStateProxy` to change the record of the voter.

Returning to the preceding code snippet, we notice that there are two parameters missing from the function call that have to be provided, which are `functionUuid` and `personCpr`. After finding the right `functionUuid` (see Figure 3) in the browser’s cookie store, and knowing the first author’s civil registration number (CPR) we were able to construct and place a request to the DVE-system through the `vælgererklæring.dk` website that would allow us to submit a fraudulent endorsement for the party De Visionære (which we later removed to return the database to an “uncorrupted state”).



Fig. 4. The attack line (sensitive information redacted)

Figure 4 depicts the problematic command with the CPR number redacted that we issued through the Firefox console, while Figure 5 shows the letter sent to the first author’s authenticated e-Boks, confirming a successful attack.

In summary, we have shown that even the simple client-side security objective was violated. After this shocking revelation, we did not attempt further attacks, but many, similar in nature, come to mind. For example: would it be possible to submit or withdraw endorsements impersonating other voters? or to circumvent other restrictions imposed by the law? We concluded at the time that this attack demonstrates a complete lack of integrity in the endorsements database, which also questions the decisions that have been drawn from it.

4.3 Socio-Technical Vulnerabilities

Another grave vulnerability of the DVE-system was pointed out by ØIM, when they sent us the invitation depicted in Figure 2: the tokens for submitting endorsements are not bound to the identity of the endorser. This observation can certainly be turned into another serious attack. A party administrator could, just

Social- og Indenrigsministeriet
Holmens Kanal 22
1060 København

Carsten Elmar Schürmann



Journalnogle
2c4e0be1-60b9-4a60-9c98-5f5d0890cb48

Dato
12. august 2016

Kvittering for afgivelse af vælgererklæring

Du har nu afgivet vælgererklæring til partiet:

De Visionære
Prins Jørgens Gård 1
1218 København K
Tlf.: 22460608
Www.Visionaer.dk

Din vælgererklæring er gyldig i 18 måneder fra datoen for afgivelsen. Du kan ikke afgive vælgererklæring til støtte for et andet partis opstilling til folketingsvalg, så længe vælgererklæringen er gyldig. Når gyldigheden udløber, slettes din vælgererklæring og alle øvrige oplysninger om din afgivelse af vælgererklæring automatisk i den digitale løsning. Du kan dog trække din vælgererklæring tilbage via nedenstående link, hvis du ikke længere ønsker at deltage i anmeldelse af partiet, eller hvis du ønsker at afgive vælgererklæring til et andet parti. Har partiet allerede anmeldt sig for social- og indenrigsministeren, kan du dog ikke trække din vælgererklæring tilbage. Anmeldelsen af et parti er gyldig indtil førstkomende folketingsvalg, dog mindst 18 måneder. Herefter kan du igen frit afgive vælgererklæring til et parti.

Klik på linket for at komme til vælgererklæring.dk, hvis du ønsker at trække din vælgererklæring tilbage:

<https://www.vaelgererklaering.dk/apos2/eve/vaelger?uuid=2c4e0be1-60b9-4a60-9c98-5f5d0890cb48>

Linket er aktivt til og med 3. februar 2018.

Venlig hilsen
Social- og Indenrigsministeriet

48880161

Fig. 5. The receipt (sensitive information redacted)

as our contact at ØIM has done, generate many tokens by specifying a party email. These tokens will arrive in the mailbox after seven days, which can then be forwarded to potential endorsers, effectively bypassing the seven day “reflection period” required by administrative regulation. Endorsements can happen right in front of party representatives, exposing voters to social engineering and coercion attacks.

The only argument that softens the severity of the attack is that endorsers receive a final receipt to their personal e-Boks, and therefore they will always be able to withdraw an endorsement that they may not feel comfortable with.

5 Reflections and Recommendations

The legal framework and the technical framework do not match. The legal framework is weak, because it does not require the digital solution to be software independent or at the very least verifiable. The administrative regulation permits technical solutions that do little to strengthen public confidence in endorsement integrity. The technical framework (that we studied in this paper) is weak, because it is not made to be used in an adversarial environment.

We recommend that the ØIM strengthens the administrative regulation and involves experts in designing a new generation DVE-system as soon as possible, to ensure that the system lives up to best practices and follows international recommendations. ØIM should also consider to strengthen the privacy of endorsers and require coercion resistance beyond just the ability to withdraw endorsements at a later stage.

6 Conclusions

In this paper, we present an analysis of an early and unpatched version of the Danish party endorsement (DVE)-system and argue that the design of the law and the DVE-system are fundamentally flawed. We show that even non-eligible endorsers in possession of an invitation email and a valid NemID could have endorsed the prospective party. According to ØIM this was the only successfully executed attack using this particular vulnerability before the system was fixed in December 2016. A consequence of our demonstration is that the DVE-database of endorsements may lack integrity, and consequently so do all decisions based on it.

Responsible disclosure. After discovering the flaw on August 12th 2016, we informed the ministry of our findings by email. On August 25th the ministry asked the first author to present a detailed document with all our findings. On September 14th 2016 we were invited to present our findings to ØIM and the software vendor, which we did on September 27th and offered our help to fix the fundamental design issues in the protocol. Finally, on October 25th, 2016, we organized a meeting with the vendor at the IT University of Copenhagen, screening a video recording of the attack, and on December 19th we received

notification that the specific issue was resolved. Since then, ØIM has neither invited us to look at the security of the DVE-system again, nor have they invited us to bid or participate in a commissioned security review, nor have they shared with us the results of the security review conducted by an unknown third party.

The aftermath. In 2016, a new party called the “Nye Borgelige”¹² was introduced on the ballot based on the endorsements stored in the database. Another prospective party “Nationalpartiet”¹³ also applied to become a party in 2016, but did not succeed, citing unavailability of paper endorsements and problems with the DVE-system. Their application was eventually rejected. Both decisions were made using the version of the DVE-system discussed in this paper. In 2019, two new parties were added to the ballot, “Klaus Riskær Pedersen”¹⁴ and “Stram Kurs”¹⁵. Both parties stand accused of having exploited the socio-technical vulnerability described in Section 4.3.

In the light of the recent accusations, representative of all parties represented in the parliament have voted that a new DVE-system should be built/procured, citing lack of usability and usable security. The new system should be in place as soon as possible. We pledge our support to ØIM to assist with improving both the regulatory framework and the technology.

Acknowledgements: We would like to thank Christine Boeskov and Søren Stauning from ØIM for their comments on earlier versions of this paper.

References

- [BRR⁺15] Josh Benaloh, Ronald L. Rivest, Peter Y. A. Ryan, Philip B. Stark, Vanessa Teague, and Poorvi L. Vora. End-to-end verifiability. *CoRR*, abs/1504.03778, 2015.
- [CKN⁺14] Michael Clouser, Robert Krimmer, Henrik Nore, Carsten Schürmann, and Peter Wolf. *The Use of Open Source Technology in Elections*. Resources on Electoral Processes. International IDEA, Stockholm, 2014. ISBN 978-91-87729-68-3.
- [EBB⁺15] Jordi Barrat Esteve, Eden Bolo, Alejandro Bravo, Robert Krimmer, Stephan Neumann, Al A. Parreño, Carsten Schürmann, Melanie Volkamer, and Peter Wolf. *Certification of ICTs in Elections*. International Institute for Democracy and Electoral Assistance (IDEA), Stockholm, 2015.
- [RW06] R.L. Rivest and J.P. Wack. On the notion of ‘software independence’ in voting systems, July 2006. <http://vote.nist.gov/SI-in-voting.pdf>.

¹² See <https://nyeborgerlige.dk/>

¹³ See <http://www.nationalpartiet.dk/>

¹⁴ See <https://www.klausriskærpedersen.dk/>

¹⁵ See <https://stramkurs.dk/>

Online Voting in Ontario Municipal Elections: A Conflict of Legal Principles and Technology?*

Anthony Cardillo¹, Nicholas Akinyokun², and Aleksander Essex¹

¹Department of Electrical and Computer Engineering
Western University, London, ON, Canada
{acardill,aessex}@uwo.ca

²School of Computing and Information Systems
The University of Melbourne, Australia
oakinyokun@student.unimelb.edu.au

Abstract. This paper presents the first comprehensive study of the use of online voting technology in the province of Ontario, Canada. Despite having one of the largest concentrations of online voters globally, its use is not governed by any federal or provincial standards. This has left many municipalities to make decisions largely in isolation, relying on for-profit vendors to set their own bar for cybersecurity and public accountability. This study presents important observations about online voting use in the 2018 Ontario municipal election and questions whether the legal principles are being met by the technology deployed in practice.

1 Introduction

In an era characterized by foreign interference in national elections, it can be easy to lose sight of the cybersecurity of elections held at the municipal level. With much of our attention squarely focused on state-level threat actors, we must occasionally remind ourselves of a more fundamental threat to our democracies: loss of confidence in the process itself. This idea is summarized expertly by the Supreme Court of Canada:

Maintaining confidence in the electoral process is essential to preserve the integrity of the electoral system, which is the cornerstone of (our) democracy. ... if (electors) lack confidence in the electoral system, they will be discouraged from participating in a meaningful way in the electoral process. More importantly, they will lack faith in their elected representatives. Confidence in the electoral process is, therefore, a pressing and substantial objective.¹

* This paper is an extended abstract. The full version is available online: <https://whisperlab.org/ontario-online.pdf>

¹ Harper v. Canada (Attorney General), [2004] 1 SCR 827, 2004 SCC 33 (CanLII). Available online: <http://canlii.ca/t/1h2c9>

In this paper, we study online voting in the context of Ontario’s 2018 municipal elections in which as many as one million voters cast a ballot online. In the absence of almost any federal or provincial government standards or oversight, municipalities and their private for-profit vendors are primarily left to set their own bar for cybersecurity and public accountability in their elections.

We present several observations about the election and question whether the associated practices align with the legal principles established in case law. We believe these observations will prove significant to municipalities, since, as the Chief Electoral Officer of Ontario recently pointed out:

As the public becomes more informed about software, malware, and manipulation of technology data systems, they are increasingly interested in knowing exactly how election technology preserves the integrity of our electoral process and the confidentiality of their personal information [5].

This leads to the central thesis of this work: purposeful, malicious interference, or fraud is not necessary to undermine an election. Nor is the honest discharge of an election sufficient to prevent it. Given enough time, a seed of doubt in an otherwise faithfully executed election may eventually grow to accomplish what even the best threat actor cannot. With the goal of preventing this outcome, we hope this work will serve as an encouragement to Ontario municipalities and others contemplating online voting to develop standards to address these issues.

Contribution. We present the first comprehensive study of the cybersecurity of online voting in Ontario’s 2018 municipal elections, including a complete accounting of municipalities, ballot options, vendor partnerships, and the extent of municipalities affected by emergency extensions to the voting period on election night. We present findings showing issues with weak voter authentication; poor transparency of election results; and, a general lack of disaster-preparedness which resulted in nearly one million voters receiving an emergency extension to the voting period due to a misconfiguration in the online infrastructure on election night. We study date of birth as a login credential and show that it could be used to uniquely re-identify up to 50% of online voters in the 2018 election.

2 Background

Canada does not offer online voting at the federal level, and cybersecurity is a significant factor in that position. The parliamentary Special Commission on Electoral Reform (ERRE) reviewed the possibility of online voting in 2016 and recommended against its introduction on cybersecurity grounds [18, 3].

2.1 Online Voting in Ontario Municipalities

Municipalities in the provinces of Ontario and Nova Scotia have held online elections since 2003 [10]. Since then, adoption in Ontario has followed an exponential

trend, nearly doubling with each election cycle. As of the 2018 municipal election, we observed 45% of municipalities (accounting for 29% of the province’s 9.4 million voters) offered online voting. Furthermore, 33% of municipalities (accounting for 16% of all voters in Ontario) eliminated paper ballots completely. While hard numbers of turnout by voting method have not been made publicly available, we estimate the number of Ontario voters casting a ballot online between 2-4 times higher than Estonia (see Section 3.3).

Despite concerns about the use of online voting, the Communications Security Establishment (CSE) assesses threats to municipal elections as “very likely to remain at its current low level,” [3], which is often cited by municipal councils and clerks favoring the adoption of online voting. While the report considers conventional threat actors (nation-states, hacktivists, cybercriminals, terrorist groups, political actors), it overlooks others, such as election officials, system manufacturers, and system operators (cf. [17]). Nor does it consider the inherent threat to confidence posed by the use of non-transparent election technology.

Furthermore, no technical standards currently exist within Canada for designing, testing, or certifying online voting systems, nor auditing or otherwise independently verifying the result they produce. Nor do the federal or provincial governments provide guidance on the procurement and operation of such systems. As we discuss in Section 3.1, Ontario offers almost no oversight to the degree that they do not even track which municipalities offer online voting.

Finally, the population difference between the largest and smallest municipalities in Ontario is *four* orders of magnitude. While some municipalities have the resources to perform security reviews of vendor proposals,² others rely almost entirely on their vendors for cyber-expertise.

2.2 Legal Context

A commonly used expression in Ontario municipal politics is that “cities are creatures of the province,” which references the fact that the province legislates their existence.³ Municipalities are categorized by three tiers: single, lower, and upper. Upper-tier municipalities correspond to counties or regional municipalities, which consist of multiple lower-tier municipalities. Municipal councils exist at all three tiers; however, elections are only conducted by single- or lower-tier municipalities. The composition of upper-tier councils is either determined automatically, e.g., as a council of all the mayors of the constituent lower-tiers (as in Bruce County) or by a direct ballot question in the constituent lower tier-elections (as in the election of the Regional Chair of Durham).

Ontario has 444 municipalities: 30 upper-tier, and 414 lower- and single-tier. In the 2018 Ontario Municipal Election held on October 22nd, each single- and

² Security Assessment of Vendor Proposals, Toronto, 2014. Available online: <https://www.verifiedvoting.org/wp-content/uploads/2014/09/Canada-2014-01543-security-report.pdf>

³ Municipal Act, 2001, S.O. 2001, c. 25. Available online: <https://www.ontario.ca/laws/statute/01m25>

lower-tier municipality was responsible for organizing and delivering its own independent election. This means up to 414 municipal councils made up to 414 individual decisions about the use of online voting in their election.

Municipal Elections Act (MEA). The main piece of legislation governing municipal elections in Ontario is the Ontario Municipal Elections Act (MEA).⁴ Although online voting is not explicitly mentioned in the MEA, it allows a municipal council to pass by-laws authorizing the use of “an alternative voting method, such as voting by mail or by telephone, that does not require electors to attend at a voting place in order to vote,” (MEA sec. 42). Additionally, it grants municipal clerks the power to establish procedures for alternative voting methods.

Whereas the MEA provides extensive language surrounding the delivery of paper-ballot elections and other electoral matters such as the use of rank-choice ballots, it provides no guidance regarding how to deliver an online election. The Act does not even contain the words “online,” or “internet.”

This contrast between specificity for paper-ballot in-person elections on the one hand and ambiguity toward online voting on the other leads to an apparent contradiction in places between the letter of the law, and the technology being used in practice. For example, the Act requires that “no person shall communicate any information obtained at a voting place about how an elector intends to vote or has voted,” (MEA, Sec. 49 (2)c). However, the act of casting a ballot in an online voting system communicates—in the literal network communication sense—information to the online system about how an elector has voted.

Legal Principles. Democratic and legal principles provide an important lens through which to interpret the use of technology in elections (cf. [1]), especially in the absence of technical standards. The principles of the MEA are not included in the MEA itself, but have been inferred from its provisions and set out in case law as follows:⁵

- **Ballot secrecy.** The secrecy and confidentiality of the voting process is paramount,
- **Fairness.** The election shall be fair and non-biased. Voters and candidates shall be treated fairly and consistently,
- **Accessibility.** The election shall be accessible to the voters,
- **Integrity.** The integrity of the voting process shall be maintained throughout the election,
- **Certainty.** There is to be certainty that the results of the election reflect the votes cast,
- **Eligibility.** Valid votes are counted and invalid votes are rejected so far as reasonably possible.

⁴ Municipal Elections Act, 1996, S.O. 1996, c. 32, Sched. Available online: <https://www.ontario.ca/laws/statute/96m32>

⁵ Cusimano v. Toronto (City), 2011 ONSC 2527 (CanLII) at para. 67. Available online: <http://canlii.ca/t/f15pg>

3 Election Statistics

3.1 Initial Survey of Available Data

Several months before the election, we set out to obtain a list of which cities were intending to use online voting. We wrote to the Ontario Ministry of Municipal Affairs and Housing (MAH) in March 2018 and were surprised to discover this list did not exist. Although the MEA requires local municipal councils to formally pass a by-law authorizing the use of an alternative voting method in the year prior to the election, we were informed in an email response that “municipalities are not required to declare their intentions to the province ... the Ministry does not have a list of municipalities that will be using internet voting in the 2018 municipal election.” Several of the vendors had commented publicly on the total number of their municipal clients, but none offered a breakdown. One of our colleagues requested such a breakdown from one of the vendors, but they refused to provide it. It was evident that we would need to collect the data ourselves.

3.2 Data Collection Methodology

Correcting the Municipal List. Our first step was to obtain a complete list of Ontario’s 444 municipalities, their tier-status, and associated URL. We consulted MAH’s online list⁶ and quickly discovered many URLs were incorrect or outdated. For example, many municipalities had switched from the older `city.on.ca` form to the newer `city.ca` form. Some cities no longer owned the URL listed. For example, the URLs listed for Mattawan and Larder Lake directed to Japanese-language websites. We had to inspect each of the 444 URLs for correctness manually. We wrote to MAH around the time of the election and received an acknowledgment that they would undertake to update their list. Six months later, many of the errors we identified remained uncorrected.

Tracking Down Voting Website URLs. Our next step was to determine which municipalities were planning to use online voting, which vendor they contracted, and the URL of the voting website. We were concerned that finding the URLs would be challenging, since many municipalities we observed made it a practice never to list it anywhere online, revealing them only in the voter information package mailed to voters before the election. Sample voter information packages found online used a placeholder URL (e.g., `anytown.election.ca`), and candidate social media fairly consistently respected this approach. We believe the practice of concealing URLs was meant as a cybersecurity protection to make the voting site harder to find by non-residents.

We made inquiries with colleagues in the province about the URL of the voting site in their respective cities and observed a trend in which vendors were encoding a municipality’s voting website either into sub-domain (e.g., Intelivote

⁶ List of Ontario Municipalities. Ontario Ministry of Municipal Affairs and Housing. <http://www.mah.gov.on.ca/page1591.aspx>

used the form `city.evote2018.ca`), or sub-directory (e.g., Dominion used the form `intvoting.com/city`). We then wrote a collection of automated scripts that used the municipal list to search for the existence of voting sites based on the particular URL form a vendor was using. For municipalities encoded into sub-domains, we performed passive DNS lookups. For names encoded as sub-directories, we attempted to fetch the HTTP header from the server and inferred whether the page existed from the response code.

For any municipalities not captured by the bulk search, we conducted a labor-intensive manual web search of online municipal documents, including meeting minutes of councils and voter accessibility documentation. This allowed us to identify municipalities using custom domain names (e.g., `kenoravotes.ca`), and abbreviations (e.g., Elizabethtown-Kitley used `ektwp.evote2018.ca`). The only URL we were not able to find with this approach was Markham's, who were partnered with Scytl, so there was no obvious way to infer the URL from others. Furthermore, staff and candidates made a seemingly flawless effort of not mentioning the URL in online documents, social media, etc. Ultimately, however, we found it (`evote.markham.ca`) by searching certificate transparency logs.

Cross-validation and Corrections. After the election, the Association of Municipalities of Ontario (AMO) published a list of municipalities broken down by election results, number of eligible voters, and voting methods offered.⁷ Rather than being made available as a single downloadable data file, the figures were spread across 444 individual web-pages, which we scraped in order to cross-validate against our list.

We found a few mistakes in the AMO list. For example, the municipalities of Belleville, Bracebridge, and Timmins were reported as not using online voting when, in fact, they did. The township of Machin was reported as using online voting when it did not. We shared this information with the AMO. We also discovered three municipalities with active websites on Intelivote's domain for which no election was held as the races were acclaimed. We also initially falsely concluded that Newmarket had contracted Intelivote since there was an active website on the `evote2018.ca` domain. The Newmarket deputy clerk later confirmed they contracted Scytl instead.

In terms of the correctness of self-declared vendor figures, we observed three of the four vendors reporting more municipal clients than actual elections run. See the full report for further discussion.

3.3 Results: Who Used Online Voting?

Of the 444 municipalities, 30 upper-tier municipalities do not hold elections, and 23 single-/lower-tier municipal councils were acclaimed and therefore did not run an election. In total there were 391 elections involving 9,444,628 eligible voters.

⁷ <https://elections.amo.on.ca>

Voting method	Municipalities	Eligible Voters	
Electronic ballot only	131 (33.5%)	1,512,076	(16.0%)
Electronic and paper	46 (11.8%)	1,230,019	(13.0%)
Paper ballot only	214 (54.7%)	6,702,533	(71.0%)
Total	391	9,444,628	

Table 1. Voting methods offered in the 2018 Ontario municipal election.

Of those, 177 offered an online voting option, of which 131 were completely paperless. Our full dataset is available for download online.⁸

Table 1 shows the number of municipalities and eligible voters by voting method. These consisted of electronic ballot options (online and telephone ballot casting), paper ballot options (incl. optical-scan and postal mail-in), or a combination of options. Combining the AMO’s population data with our observations, our results show that online voting was available to approximately 2.74 million voters, or 29% of the voting population. Of these, approximately 1.51 million voters, or 16% of the voting population experienced a completely paperless ballot, cast either online or by telephone.

Most municipalities did not report turnout categorized by voting method. However, if we combine our numbers with the AMO’s province-wide turnout rate of 38.2%, we estimate the total number of voters who cast ballots online to be between 0.5–1 million, which is approximately 2–4 times the online ballots cast in the 2019 Estonian parliamentary elections.⁹

We observed 4 vendors active in the 2018 Ontario election: Dominion Voting Systems, Intelivote Systems, Simply Voting, and Scytl. Intelivote and Scytl worked together in partnership, although the extent of their business relationship remains unclear to us. Though ostensibly distinct business entities, we observed both Scytl Canada Inc. and Intelivote Systems Inc. have a registered office at the same mailing address in Dartmouth, NS. Additionally, we observed a considerable portion of Intelivote’s web content (Javascript, images) and infrastructure (IPs, domains) appears to have been provided by Scytl. Of the municipalities offering online voting, Table 2 shows the relative market share.

4 Election Observations and Findings

In this section we present three significant findings. Additional findings are presented in the full version.

⁸ <https://whisperlab.org/ontario-online.csv>

⁹ <https://www.valimised.ee/en/archive/statistics-about-internet-voting-estonia>

Vendor	Municipalities	Eligible Voters
Dominion Voting Systems	49 (27.7%)	1,323,194 (48.3%)
Intelivote Systems	98 (55.4%)	860,985 (31.4%)
Simply Voting	28 (15.8%)	304,479 (11.1%)
Scytl	2 (1.1%)	253,437 (9.2%)
Total	177	2,742,095

Table 2. Online voting market share in the 2018 Ontario municipal election.

4.1 Disaster Preparedness

One open question was how municipalities were preparing for the possibility of a disaster in the online voting infrastructure (accidental or otherwise), especially in the absence of standards. Our initial examination of municipal documents found no mention of a disaster recovery plan. We raised this issue in the media six months prior to the election [8]. Several clerks were also interviewed but “could not provide a disaster plan to be implemented in case the election is hacked, or irregularities tip the balance in favor of a candidate who should not have been elected.” The clerk of Sarnia acknowledged, “I don’t have a disaster plan in place right now, I’d have to talk to my vendor about that.” The clerk for St. Thomas added, “We’re hoping nothing does happen.”

Election night emergencies. As it turned out, something significant did happen. Starting around 6 p.m. on election night, the voting websites of 43 municipalities experienced a dramatic slowdown. Just before 6 p.m., we performed a network capture of the login page for Hanover’s voting site, and after 2 minutes the page load timed out. Although the static content appeared to load, the dynamic content loads dragged on, and some eventually timed out.

In the face of an unavailable voting website, and with many affected municipalities without any paper ballot option as a back-up, many clerks made the extraordinary decision to declare emergencies to extend the voting period. In some cases, voting was extended later into the evening by 1-2 hours. The majority of affected municipalities, however, extended voting by a full 24 hours [20, 12].

A statement by Dominion on the night of the election attributed the slowdown to their co-location provider (an IT sub-contractor) “placing an unauthorized limit on incoming voting traffic that was roughly 1/10th of the system’s designated bandwidth.” Dominion did not disclose the names of the affected cities, so we assembled this list manually by examining multiple news sources and municipal websites.¹⁴ The number of municipalities and affected voters are shown in Table 3. A complete list of municipalities who extended voting periods is provided in the full version.

Five months after the election we were invited to present preliminary results of this paper to the Association of Municipal Managers, Clerks and Treasurers

Emergency Extension	Municipalities	Eligible Voters
24-hour extension	35	575,022
Same-evening extension	8	422,085
Total	43	997,107

Table 3. Emergency extensions due to Dominion’s election night slowdown

of Ontario (AMCTO). We spoke to several clerks and a representative from Dominion. None were willing or able to provide any explanation for the events that lead to the co-location provider’s bandwidth restriction, nor even the provider’s identity. According to Sudbury’s post-election report, however, the slowdown was determined to be a “miscommunication between Dominion and the service provider.”¹⁰

Conflict with principles. The outage may contradict the accessibility principle on the basis that the voting websites became inaccessible to voters. The unexpected nature of the outage may contradict the fairness principle on the basis that the emergency extensions to the voting periods allowed some voters an additional day to form a decision relative to those who had cast their ballots just prior to the slow-down.

4.2 Voter Authentication

Voter lists at the municipal level are largely derived from the Municipal Property Assessment Corporation (MPAC), whose primary business is not voter list management. This mismatch of focus has led to inaccurate municipal voter lists over the years, and numerous news stories ran prior to the election on the subject. Because the lists are derived from property ownership, we heard anecdotal accounts of rental tenants who did not receive their online voting login credentials, whereas non-resident adult children away in college did. Other accounts described land owners of multiple properties receiving multiple login credentials. One news story reported a deceased dog in the town of Mono received a PIN [7].

Online voting credentials. The primary credential needed to cast a ballot online consisted of a knowledge factor (a PIN and/or ID) transmitted to the voter in a voter information package via postal mail. To our knowledge, the sole exception was the city of Cambridge, which sent PINs via email. In almost all cases a second knowledge factor (date of birth) was required. See Table 4 for a breakdown of credentials used by the vendor.

¹⁰ City of Sudbury. Post Election Report. Jan 21, 2019. Available: <https://agendasonline.greatersudbury.ca/index.cfm?pg=feed&action=file&agenda=report&itemid=25&id=1312>

Vendor	Primary credential (mailed)	Secondary credential
Dominion	13-digit ID & 8-digit PIN	Date of birth
Intelivote	16-digit PIN	Date of birth
Scytl	16-digit PIN	Date of birth
Simply Voting	9-digit PIN	Date of birth

Table 4. Credentials needed to vote online

The use of single credential for voter authentication is inadvisable since access to the voter information package is sufficient to cast a ballot on another’s behalf. Furthermore, some voters observed that the PINs were legible through the envelope when held up to bright light. See Figure 1. In order to mitigate this risk, most municipalities required a date of birth as a secondary credential. Note that authentication is still considered single-factor (as opposed to multi-factor) authentication since both credentials are knowledge factors.

Dates of birth, however, make a poor login credential for several reasons. Aside from the significant privacy implications (which we discuss in Section 5), they are low entropy, cannot be changed, and typically are not very secret, especially when considering one’s co-habitants (i.e., friends and family) are potential threats. Aside from the widespread practice of sharing dates of birth on social media websites, some US states such as Ohio include dates of birth in voter registries which are freely available for download online.

Much of the voting literature on eligibility and authentication focuses on threats like coercion and vote selling. In practice, however, it appears that a far more pervasive version of these threats is also more casual.

Voting on someone else’s behalf is an offense under the MEA. Nevertheless, we heard anecdotal accounts from several independent sources of parents who voted on behalf of children living in another city, or people who voted on behalf of their spouse while they were at work. We also heard accounts of individuals gifting their unopened voter information packages to friends and family.

Ultimately, knowledge of a PIN or date of birth does not establish a voter’s identity. It merely establishes to the voting server that some entity on the other end of the connection knows a secret. Secrets, of course, can be transferred or intercepted. Indeed, the fraudulent interception of online voting PINs is currently the subject of a criminal investigation in Alberta [6, 15].

Conflict with principles. This form of voter authentication and eligibility verification may contradict a number of principles. The use of dates of birth evidently contradicts the ballot secrecy principle (see Section 5). The multiple anecdotal accounts of individuals voting on behalf of others would seem to contradict the principles of fairness and eligibility.



Fig. 1. Voter login credentials visible through mail envelope

4.3 Transparency and Accountability

The opportunity for an independent evaluation of security claims and implementations is vital to the public interest. There are numerous examples in the academic literature of improperly implemented software leading to critical vulnerabilities in online voting technology (see, e.g., [16, 21, 9, 19]).

As a substantial illustration of this point, academics recently discovered several critical implementation vulnerabilities in ScytI’s software as implemented for the proposed Swiss Post national online voting system [11, 13]. These included, among other things, the possibility of the election provider creating a valid-looking mathematical proof of a fake election result. On March 29, 2019, Swiss Post announced that it would suspend its e-voting system as a result of critical “errors in the source code.” Importantly, these findings were possible because Swiss Post made the system and source code available for independent review not only to the general public but to the international community (Swiss Post reported 3,200 participants from 137 countries).¹¹

No such opportunity for independent review was provided in the election. This fact is troubling, as we found numerous municipal documents in circulation which made security claims which were: short on detail; mostly non-technical; and, largely unverifiable by members of the public.

Result by fiat? For several months after the election, we received phone calls from council candidates from around the province asking how they could verify the correctness of the online vote totals. Many of them had experienced an unexpected loss, and although they all acknowledged there were entirely legitimate possible explanations for the outcome, they were understandably in search of answers.

Unfortunately, however, there appeared to be little objective evidence either supporting or disputing a particular online election result beyond the clerk’s

¹¹ <https://www.post.ch/en/about-us/company/media/press-releases/2019/swiss-post-temporarily-suspends-its-e-voting-system>

declaration of results itself. None of the deployed online systems produced an accompanying paper trail, and there is currently no online equivalent of risk-limiting audits [14].

Based on URLs found in municipal documents obtained under access to information, clerks accessed election results by logging into their vendor’s web admin portal, where they could generate reports of events, activity, and results. The extent of objective evidence the clerks received (if any) remains an open question. Many of the public documents we examined either pointed to the existence of an independent auditor who performed basic logic and accuracy testing, or to third-party firms who performed routine penetration testing of the online system. Aside from neither of these constituting proof of an election outcome, our search of municipal documents uncovered no publicly available reports on the topic. What reassurance do audits provide the public if their scope, methodology and findings are entirely unavailable?

After the election, several residents and former candidates in Wasaga Beach contacted us to share their deep concern about an unexpected election loss. Among other things, we suggested they inquire as to whether there were any IPs responsible for casting an unusually large proportion of ballots in the election. Initially, residents contacted the vendor but were referred to the city clerk. We then helped them write a freedom of information request. The clerk responded that they could not provide this information because the municipality did not have any such records.

Conflict with principles. Our observations point to what we believe is a serious concern over the degree of certainty of results achievable in the current online voting setting. If there ever was evidence of an incorrect result or fault (whether due to error or otherwise), some of the experiences we heard suggest that it would exist beyond the reach of the public.

As Elections Ontario pointed out in its study of alternative voting technologies, unless the implementation of an online voting system provides auditable evidence of the election results, then “the process is open to question” [4]. Perhaps the most pressing issue for Ontario municipal elections is whether online voting in the next election can provide candidates an objective measure of certainty in the results they will have worked so hard to achieve.

5 Analysis of Voter Confidentiality and Ballot Secrecy

A significantly overlooked question in the online voting conversation in Ontario has been to what extent an online voting vendor can associate a voter’s identity with their ballot selection. Recalling the MEA principle stating secrecy of the ballot is paramount, in this section we ask how unique is a voter’s date of birth (DOB) within their particular municipal election.

Data collection. As part of our study leading up to the election we collected basic web data from each of the 180 active voting websites we found. This in-

cluded the IP addresses, TLS certificates, HTTP headers, and static HTML of the login pages. We examined the source code of each web page for elements that indicated the presence of a DOB field. Most voting sites loaded the DOB field dynamically. We did not wish to burden on the election servers by capturing full HTTP sessions of the login pages of every municipality. Loading the login page of a single Dominion municipality, for example, required over 100 separate GET requests, so we opted to capture a single municipality per vendor. As a result do not have a complete accounting of which municipalities used DOB as a login credential, though our sampling of municipal documents suggests a large majority did.

We used a web proxy on the evening of the election to capture HTTP messages sent by the voting client to the election server when the login button was clicked. We used breakpoints so that we could intercept and examine POST messages without actually forwarding them to the server. At the time of capture, we were unable to complete a load of Dominion’s login page (see Section 4.1). We found that within a single web session the server receives information about: the voter’s city (from the URL itself), their date of birth (from the login), and how they voted. We now examine the degree to which this information could be used to associate voter and vote.

5.1 Re-identifying Voters with City and Date of Birth

As a rough estimate, there are approximately 30,000 possible dates of birth in a voting age population (365 days times 80 years). Considering that many of the municipalities who ran online voting had voting populations numbering in the low thousands, it seemed likely that many voters would have a unique DOB in their town. To model this, we used the AMO’s data on eligible voters in each municipality, combined with a sizable real-world DOB dataset to create a distribution from which we could run experiments to study the uniqueness of dates of birth within each municipality.

Modeling Date of Birth distribution. Our experiment required a DOB distribution representative of a general population of voting age individuals. In the US, many states provide public access to voter registries. Most include names and postal addresses, and some even include birth dates. We decided to use the statewide Ohio voter registry, which is a large publicly available dataset (>7 million records) containing voter DOB information.¹²

For each municipality, we ran the following experiment: we uniformly sampled dates of birth from the Ohio voter registry equal to the number of eligible voters in the given municipality. To determine the uniqueness of each record, we counted the frequency of each DOB in the sample, and then counted the number of times each frequency value was recorded. The result was a probability distribution of finite outcome, where the probability of each outcome represented the likelihood

¹² Ohio statewide voter files. Available: <https://www6.sos.state.oh.us>

Vendor	Eligible Voters	$k = 1$		$k = 5$	
		Max Affected	% of Eligible	Max Affected	% of Eligible
Dominion	1,323,194	531,758	(40.2%)	1,181,876	(89.3%)
Intelivote	860,985	613,999	(71.3%)	847,876	(98.5%)
Simply Voting	304,479	190,097	(62.4%)	294,912	(96.9%)
Scytl	253,437	32,880	(13.0%)	123,712	(48.8%)
Total	2,742,095	1,368,734	(49.9%)	2,448,376	(89.3%)

Table 5. Degree to which voters were uniquely identifiable ($k = 1$) or near-uniquely identifiable ($k = 5$) by the use of date of birth as a login credential

that a DOB record would have exactly that many matches in the election. We ran 1,000 trials for each municipality, generating a cumulative distribution where the probability of each outcome represented the likelihood that a particular DOB would have up to that many matches in the election. We estimate the number of re-identified voters within a cell size of k by multiplying the number of eligible voters in a given municipality by the probability of k or fewer matches from its cumulative distribution.

Results. The repeated trial experiment was run for each municipality, determining the maximum number of affected voters that were uniquely identifiable (i.e., $k = 1$). We also considered an *almost* uniquely identifiable case ($k = 5$), which we chose as the smallest cell size found in industry, although a cell size of $k > 20$ is typical. [2]. A breakdown of our findings by vendor is shown in Table 5. Of 9,444,628 eligible voters in the province, 2,742,095 (29.0% of the total voting population) were at some risk of being re-identified by the combination of their city and DOB. Of these, up to 1,368,734 voters (49.9% of the total affected population) could be uniquely identified, and 2,448,376 (89.3% of the total affected population) could be near-uniquely identified. That these numbers are so high is reflective of the fact that much of the 1.4 million voters were spread across numerous small towns, significantly increasing the chance of a unique city/DOB combination. If we were to simulate this effect for the entire province in the scenario where municipalities used online voting, we estimate that up to 2,638,340 voters (27.9%) would be uniquely re-identified and up to 5,302,183 (56.1%) would be near-uniquely identified.

In conclusion, roughly half of the voters eligible to cast online ballots in the 2018 Ontario municipal election were uniquely re-identifiable by their date of birth and town. Given this information is transmitted to the voting server in the same web session as the voter’s cast ballot, there is a strong case to be made that dates of birth as login credentials conflicts with the principle of ballot secrecy.

6 Conclusion

There is significant work to be done in Ontario if online voting is to continue in the long term. As one clerk of a large city acknowledged to us, it may take as little as one successful cyber attack for online voting to be banned permanently. The observations made in this study, however, point to a more likely failure mode without hackers, malice, or fraud. Until the technological practice inhabits the same universe as the legal principles, the absence of standards for online voting in Ontario may lead it to collapse on its own.

Acknowledgments. We are grateful to a many individuals in Ontario and beyond for important insights on technology, policy and law. Special thanks to Jane Buchanan. See the full version of the complete list of acknowledgments.

References

- [1] *Handbook for the Observation of New Voting Technologies*. Organization for Security and Cooperation in Europe (OSCE) Office for Democratic Institutions and Human Rights, 2013. ISBN 978-92-9234-869-4.
- [2] De-identification guidelines for structured data, 2016. Available online: <https://www.ipc.on.ca/resource/de-identification-guidelines-for-structured-data>.
- [3] *Cyber threats to Canada's democratic process*. Canada. Communications Security Establishment (Canada), 2017. Available online: <http://publications.gc.ca/site/eng/9.838566/publication.html>.
- [4] *Alternative Voting Technologies Report*. Elections Ontario, 2019. ISSN 978-1-4606-2017-5.
- [5] *Modernizing Ontario's Electoral Process: Report on Ontario's 42nd General Election*. Elections Ontario, 2019. Available online: <https://www.elections.on.ca/en/resource-centre/reports-and-publications.html>.
- [6] D. Anderson, C. Dunn, A. Dempster, B. Labby, and A. Neveu. Fraudulent emails used to cast votes in ucp leadership race. *CBC News*, Published April 10th, 2019. Available online: <https://www.cbc.ca/news/canada/calgary/ucp-leadership-voter-fraud-membership-lists-data-1.5091952>.
- [7] N. Boisver. Dead dog registered to cast vote in upcoming mono, ont. election. *CBC News*, Published October 11th, 2018. Available online: <https://www.cbc.ca/news/canada/toronto/decease-dog-voting-pin-1.4859489>.
- [8] C. Butler. Ontario civic elections: the problem with online voting. *CBC News*, April 4th, 2018. Available online: <https://www.cbc.ca/news/canada/london/london-ontario-online-voting-1.4598787>.
- [9] N. Chang-Fong and A. Essex. The cloudier side of cryptographic end-to-end verifiable voting: A security analysis of helios. In *32nd Annual Computer Security Applications Conference (ACSAC '16)*, CA, 2016.

- [10] N. Goodman, J. H. Pammett, and J. DeBardeleben. Internet voting: The canadian municipal experience. 33(3), 2010.
- [11] R. Haenni. Swiss post public intrusion test: Undetectable attack against vote integrity and secrecy, 2019. Available online: <https://e-voting.bfh.ch/publications/2019/>.
- [12] J. Laucius. Election night glitch points to the 'wild west' of online voting, says cybersecurity expert. *Ottawa Citizen*, October 25rd, 2019. Available online: <https://ottawacitizen.com/news/local-news/election-night-glitch-points-to-the-wild-west-of-online-voting-says-cybersecurity-expert>.
- [13] S. J. Lewis, O. Pereira, and V. Teague. How not to prove your election outcome. 2019. Available online: <https://people.eng.unimelb.edu.au/vjteague/HowNotToProveElectionOutcome.pdf>.
- [14] M. Lindeman and P. B. Stark. A gentle introduction to risk-limiting audits. *IEEE Security & Privacy*, 10(5):42–49, 2012.
- [15] A. MacVicar. Alberta ndp calls for special prosecutor to oversee rcmp investigation of ucp leadership race. *Global News*, Published May 2nd, 2019. Available online: <https://globalnews.ca/news/5233913/notley-special-prosecutor-ucp-leadership-race/>, May 2019.
- [16] M. Nemeč, M. Sys, P. Svenda, D. Klinec, and V. Matyas. The return of coppersmith's attack: Practical factorization of widely used rsa moduli. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1631–1648. ACM, 2017.
- [17] A. Regenscheid and N. Hastings. *A Threat Analysis on UOCAVA Voting Systems*. Number NISTIR 7551. US National Institute of Standards and Technology, 2008.
- [18] F. Scarpaleggia et al. *Strengthening Democracy in Canada: Principles, Process and Public Engagement for Electoral Reform*. Canada. Parliament. House of Commons. Special Committee on Electoral Reform, 2016. Available online: <http://publications.gc.ca/site/eng/9.828533/publication.html>.
- [19] D. Springall, T. Finkenauer, Z. Durumeric, J. Kitcat, H. Hursti, M. MacAlpine, and J. A. Halderman. Security analysis of the estonian internet voting system. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 703–715. ACM, 2014.
- [20] M. Warren. Online voting causes headaches in 51 ontario cities and town. *Toronto Star*. Published October 23rd, 2019. Available online: <https://www.thestar.com/news/gta/2018/10/23/internet-voting-causes-headaches-in-51-ontario-cities-and-towns.html>.
- [21] S. Wolchok, E. Wustrow, D. Isabel, and J. A. Halderman. *Financial Cryptography*, chapter Attacking the Washington, D.C. Internet Voting System, pages 114–128. 2012.

Election Integrity and Electronic Voting Machines in 2018 Georgia, USA

Kellie Ottoboni¹[0000-0002-9107-3402] and Philip B. Stark¹[0000-0002-3771-9604]

Department of Statistics, University of California, Berkeley, CA, USA

Abstract. Direct recording electronic (DRE) voting systems have been shown time and time again to be vulnerable to hacking and malfunctioning. Despite mounting evidence that DREs are unfit for use, some states in the U.S. continue to use them for local, state, and federal elections. Georgia uses DREs exclusively, among many practices that have made its elections unfair and insecure. We give a brief history of election security and integrity in Georgia from the early 2000s to the 2018 election. Nonparametric permutation tests give strong evidence that something caused DREs not to record a substantial number of votes in this election. The undervote rate in the Lieutenant Governor’s race was far higher for voters who used DREs than for voters who used paper ballots. Undervote rates were strongly associated with ethnicity, with higher undervote rates in precincts where the percentage of Black voters was higher. There is specific evidence of DRE malfunction, too: one of the seven DREs in the Winterville Train Depot polling place had results that appear to be “flipped” along party lines. None of these associations or anomalies can reasonably be ascribed to chance.

Keywords: Permutation testing, anomaly detection, DREs

Acknowledgements. We are grateful to Marilyn Marks and Jordan Wilkie for helpful conversations and suggestions.

1 Introduction

The state of Georgia was a focal point in the civil rights movement of the twentieth century. It also has a history of election problems: systematic voter suppression, voting machines that are vulnerable to undetectable security breaches, and serious security breaches of their data systems.

The 2018 midterm election returned Georgia to the national spotlight. Civil rights groups alleged that then Secretary of State Brian Kemp—who was running for Governor against Stacey Abrams, a Black woman—closed polling places, deleted voters from the rolls, and challenged voter signatures—disproportionately in Black neighborhoods [9,18,22]. A federal lawsuit against the Secretary of State demanded that Georgia replace paperless direct recording electronic (DRE) voting machines with optically scanned voter-marked paper ballot (opscan) voting systems [34]. While the judge accepted the plaintiffs’ argument that DREs (and

Georgia’s election management) have serious security problems, defendants successfully argued that they were unable to replace their equipment in time for the election. Ultimately, in-person voting in Georgia’s 2018 election was on DREs.

The 2018 election produced anomalous results that could have been caused by malfunctioning, misprogrammed, or hacked election technology, including DREs. The accuracy of DRE results cannot be checked (for instance by a risk-limiting audit) because the DREs used in Georgia do not produce a voter-verifiable paper record. There has been no forensic investigation of the DREs used in the 2018 election, although the (continuing) suit seeks to conduct one.

This paper begins with a short history of recent election integrity issues in Georgia. We summarize known security flaws of DRE voting systems and what took place in the months leading up the 2018 election. We analyze public election results and poll tapes photographed by a volunteer, finding strong statistical evidence that DREs were the source of these anomalies: that something caused DREs to miss votes in the Lieutenant Governor’s contest and to “flip” votes for one party into votes for another.

2 DRE Voting Machines

Congress passed the Help America Vote Act (HAVA) in 2002 after the problems in Florida in the 2000 presidential election. HAVA requires states to allow provisional voting and to build statewide voter registration databases, and provided funds for states to upgrade voting systems for accessibility. To receive funding, states were required to replace punchcard and lever voting systems and to provide at least one accessible voting machine per polling place [14].

Two types of systems were on the market: optical scanners (opscan), which primarily used hand-marked paper ballots, and DREs. DREs eliminate the need to print and store paper ballots, can present ballots in multiple languages, and satisfied the accessibility requirement [14].¹ DREs and in-precinct opscan systems also make it easier to report results faster than central opscan systems. While HAVA only required one accessible machine per polling place, some states opted to use DREs exclusively [43]. In 2002, four voting machine manufacturers offered DREs: Diebold Election Systems, Election Systems and Software (ES&S), Hart InterCivic, and Sequoia Voting Systems. This paper focuses on Diebold (now Premier), the lone DRE provider in Georgia.

In the following year, newly-adopted DREs caused serious problems. In the 2002 Florida primaries, some machines in Miami-Dade county failed to turn on, creating long lines that prevented some would-be voters from voting. In New Mexico, faulty programming caused machines to drop a quarter of the votes. In Virginia, the software on 10 machines caused one vote to be subtracted for every 100 votes cast for a particular candidate [40].

In 2007, studies sponsored by the Secretaries of State of California (the Top-to-Bottom Review, TTBR) and Ohio (the EVEREST study) gave conclusive

¹ There is ample evidence that the systems are not very usable in practice by voters with disabilities [33], yet they satisfy the legal requirement.

evidence that the DREs on the market had fundamental security flaws. The TTBR found physical and technological security flaws with Premier Election Systems' (formerly Diebold) DREs, including vulnerabilities that would allow someone to install malicious software that records votes incorrectly or miscounts them; susceptibility to viruses that propagate from machine to machine; unprotected information linked to individual votes that could compromise ballot anonymity; access to the voting system server software, allowing an attacker to corrupt the election management system database; "root access" to the voting system, allowing attackers to change the settings of any device on the network; and numerous physical security holes that would allow an attacker to disable parts of the device using standard office tools [5]. EVEREST found that the software for Premier DREs was "unstable" and lacked "sound software and security engineering practices" [17]. California decertified DREs from Premier, Hart InterCivic, and Sequoia, and the EVEREST study prompted Ohio to move to optical scanners.

White-hat hackers have found even more security flaws. In 2005 and 2006, Finnish computer scientist Harri Hursti demonstrated that Diebold's optical scanners could be hacked to change vote totals, and uncovered security flaws with Diebold's AccuVote-TSx machines that render "the voting terminal incurably compromised" [12,13]. In 2017, the annual DEF CON hacker conference held a "Voting Village" and supplied participating hackers with over 25 pieces of election equipment used in the United States. While EVEREST restricted the types of hacks that could be deployed against the machines, there were no such restrictions at DEF CON. Within minutes, hackers with little prior knowledge of voting systems penetrated several DREs, including the Premier AccuVote-TSx used in Georgia. They uncovered serious hardware vulnerabilities, including chips installed in sockets instead of being soldered in place to prevent removal and tampering [2]. The Voting Village has become a regular part of DEF CON as voting system vulnerabilities persist: the organizers reported in 2018, "while, on average, it takes about six minutes to vote, machines in at least 15 states can be hacked with a pen in two minutes" [3].

Security experts recommend that jurisdictions using DREs conduct forensic audits both before and after every election. An examination of the software and machines done by an independent, neutral party might detect tampering, bugs, or hacking, and would help discourage malicious attacks [35]. (However, forensic investigation is not guaranteed to detect all hacking: for instance, malware can be programmed to erase itself after doing its damage.) But historically, it has been illegal to examine voting machine software because it is considered proprietary information [24]. Without a forensic audit or a reliable paper trail against which to check reported results, there is no way to know whether a DRE accurately captured and tallied votes.

To make DREs more secure, printers can be added to create a "voter-verifiable paper audit trail" (VVPAT) displayed behind glass, so the voter can check whether their vote was cast as intended. The paper record can be used in a post-election audit, and serves as a back-up in case the device's electronic mem-

ory fails. The NIST Auditability Working Group found that the only satisfactory way to audit DREs is with a trustworthy paper record such as a VVPAT [20].

However, VVPATs can be compromised. If the printer malfunctions, the paper record is incomplete. VVPATs are difficult to audit: they are typically printed on continuous, flimsy, uncut rolls of paper, which need to be unrolled and segmented to count votes. Most VVPATs are thermal paper, which degrades quickly when exposed to heat, light, human touch [20], or household chemicals [5].

Verifiable does not imply *verified*: voters might not check a VVPAT effectively or at all. Research has shown that voters don't review their selections effectively. Voters often walk away from DREs before an electronic review screen is displayed. Errors in votes occur at the same rate whether a review screen is shown or not. In experiments where the wrong candidate was marked on an electronic review screen, only 37% of study participants noticed the error on the review screen, though 95% reported that they had checked their ballot either somewhat or very carefully [7]. A report by the Pennsylvania State Department found that when voters were shown VVPATs displayed behind glass, the glare and edges of the glass cage obstructed their selections [33]. VVPATs may not reflect voter intent, even if voters claim to review them.

Many states have been phasing out paperless DREs. In 2006, nearly 40% of voters used DREs to cast their vote. In 2016, 28 states used DREs in some capacity, but most jurisdictions had some paper record, either opscan or an electronic method with a paper backup [40]. Only five states still use paperless DREs exclusively: Delaware, Georgia, Louisiana, New Jersey, and South Carolina.

3 Voter Suppression in Georgia

Georgia faced heightened scrutiny under the Voting Rights Act of 1965 due to a history of discrimination in elections. Sections 4(b) and 5 of the 1965 Voting Rights Act required jurisdictions with prior evidence of racial discrimination to get "preclearance" from the federal government before changing their election policies. In 2013, the Supreme Court ruled in *Shelby County v. Holder* that these sections were unconstitutional because they placed undue burden on some states based on outdated evidence of discrimination against minority voters [30].

The ruling revitalized efforts to disenfranchise minority voters: without federal oversight, some states that were previously subject to the preclearance rule of the Voting Rights Act reinstated some discriminatory policies. States began to close polling places and create stricter voter registration laws. Previously, counties and states would have had to show that these changes would not differentially disenfranchise minority voters. After *Shelby County v. Holder*, Arizona, Louisiana, and Texas made changes that affect a large number of registered voters, disproportionately Black and Latino [36].

Strategically closing polling places can reduce voter turnout for specific demographic groups. It can force voters to travel farther to vote and create long lines in remaining polling places. Since the ruling, nearly a thousand polling places in the United States have been closed, many which served African Amer-

ican communities [36]. Since 2012, election officials in Georgia have closed 214 precincts—nearly 8% of the state’s polling places [22]. Officials claim that consolidating low-turnout polling places is purely a cost-saving measure [38]. However, 39 of the 159 counties in Georgia where polling places were closed have poverty rates above the state average and 30 of them served significant African American populations [22]. These closures would not have been permissible prior to 2013 under the Voting Rights Act’s preclearance rule.

Under Secretary of State Kemp, over 1.4 million voter registrations were cancelled in “routine maintenance” of the voter rolls, eliminating those marked inactive according to the law. Kemp implemented the first “exact match” law in 2010 with preclearance from the federal government, requiring a name on a registration application to exactly match the voter’s legal name. The law made it harder for voters whose registrations were removed to get back on the voter rolls. The law was dismantled after it was found unconstitutional in 2016 [32]. It was replaced in 2017 by a new exact match law. Any discrepancy between the name on the application and legal name—as innocuous as a missing hyphen—renders the registration “pending.” Civil rights groups argue that, though they are eligible, having a pending application discourages people from voting. Over 53,000 voter registration applications were pending leading up to November 2018. Nearly 70% of pending applications were from Black voters, more than double the 32% Black population percentage in the state [18].

Kemp denies he has attempted to suppress minority voting, claiming that the decision to close a precinct is up to county election officials. However, in 2015 his office provided a document giving county officials guidance on why and how to close polling places [22]. Kemp blames the racial disparity in pending voter registration applications is on sloppy voter registration efforts and poorly trained canvassers, in particular the New Georgia Project, a voter registration group (founded by Kemp’s gubernatorial opponent Stacey Abrams) that targeted African American voters and used primarily paper registration forms [18].

4 Georgia After HAVA

Georgia was the first state to adopt DREs statewide in the wake of HAVA: in November 2002, just days after HAVA was passed, the state signed a \$54 million contract with Diebold Election Systems to use the AccuVote-TS/TSx DREs [43].

During the summer of 2002, Diebold began preparing more than 20,000 DREs to be used in Georgia for the November election. A former Diebold employee alleged that during this time, before the machines had been delivered to counties, employees were asked to install three software patches on all of the DREs that would be used statewide that year. These patches did not undergo the federal certification process for voting equipment [41]. Another former Diebold employee reported that the president of Diebold’s election unit, Bob Urosevich, came to the warehouse himself to order the installation of uncertified software patches on about 5,000 machines used in DeKalb and Fulton, two historically Democratic counties [15].

This raised eyebrows when key contests in Georgia's 2002 election defied poll predictions. Longtime Democratic Senator Max Cleland was predicted to beat Republican opponent Saxby Chambliss by 3%, but in fact lost his seat by a 7% margin. Democratic incumbent Governor Roy Barnes was predicted to win 51% to 40%, but in fact lost to Republican candidate Sonny Perdue by 6% [8,24]. These Republican victories were a surprise in a historically Democratic state: Perdue was Georgia's first Republican governor in 130 years. There is no way to tell whether the outcome resulted from faulty programming or hacking, because the DREs left no paper trail.

Diebold has used political connections to ensure they remained the sole voting machine provider in Georgia. Former Secretary of State Cathy Cox, who signed the 2002 contract with Diebold, had strong ties to the company. The election director she appointed, Kathy Rogers, helped kill house bills that would have required paper records. In 2006, she resigned and took a job as Government Liaison at Diebold [6]. Cox's successor as Secretary of State, Karen Handel, started as a vocal supporter of paper trails and acknowledged publicly that she would not interact with Rogers as Diebold's liaison due to the conflict of interest. Later, Handel reversed her position on paper ballots, and the media revealed that she had received \$25,000 in campaign contributions from employees connected with Diebold's lobbying firm, Massey & Bowers [37]. Members of the state government have ignored security experts who pointed out problems with Diebold's touchscreen machines.

Georgia's election security issues reach beyond voting machines. In 2016, a cybersecurity researcher at Oak Ridge National Laboratory, Logan Lamb, discovered that he could download files from the state's "secure" election server. Among these files were the entire voter registration database for the state of Georgia, including sensitive personal information, instructional PDFs with passwords for poll workers to sign into a central server on Election Day, and software files for the state's ExpressPoll pollbooks that are used to verify voters' eligibility [44]. This intrusion would have allowed Lamb to alter entries in the voter registration database or the pollbooks, preventing some voters from casting their ballots. Lamb's concern about malicious hacking was not a purely theoretical: an NSA investigation found that Russian hackers targeted 39 states in the summer and fall leading up to the 2016 presidential election [26].

These were not the only security concerns at the state's Center for Election Services (CES), housed at Kennesaw State University under a long-standing contract with the Secretary of State. For instance, CES was using an outdated version of their content management software, Drupal, which would allow hackers to seize control of their websites. A software patch had been available since 2014, but CES had not installed it. Lamb notified the executive director of CES, Merle King, of the problems; King agreed to fix them and allegedly pressed Lamb not to talk to the media or other officials about the security issues [42].

CES did not secure their server, nor did they inform anyone about the Logan's breach. In March 2017, another cybersecurity researcher found that CES still had not secured its files properly. The issue was elevated to authorities above

King, and it was the first time that the Secretary of State’s office heard about the breach. In response to this poor management, the Secretary of State office signed a new agreement with Kennesaw State University to transfer CES to its own offices [42].

In July 2017, state voters and The Coalition for Good Governance filed a lawsuit against Georgia Secretary of State Kemp, alleging that he had ignored evidence that the state’s electoral system is vulnerable to fraud and hacking. The plaintiffs demanded that the state use paper ballots in future elections to guard against interference [34,35]. They requested to examine the CES servers at Kennesaw State University for evidence. Four days after the group filed the lawsuit, IT employees at CES wiped their servers of all prior election data. They later degaussed two remaining servers: key evidence was permanently erased. There is no proof that CES deliberately destroyed evidence, and the Secretary of State’s office claims that the servers were wiped before they were officially served with the lawsuit in late July. However, Kemp’s office was alerted about the lawsuit and declined to comment in the days between when the suit was filed and when the CES wiped its servers [27].

4.1 The November 2018 Election

The lawsuit, *Curling v. Kemp*, continued into September 2018, just before the midterm elections (and is ongoing at the time of writing). Testimony from the plaintiffs centered on two issues: security issues with DREs and the state’s procedures and data handling before and after Election Day. The current director of CES testified that the server that each county uses to construct its ballots is “air-gapped” from the Internet, but that he uses thumb drives, email, and an online repository to store and move data—all of which expose voting systems to malware. A county official testified that they use analog phone lines to transmit results to the Secretary of State. Computer scientists have testified that these are all vulnerable channels [19].

The state’s rebuttal did not seriously address the security concerns, but argued that there was not enough time before the election to switch to paper ballots. Kemp had convened the Secure, Accessible, & Fair Elections (SAFE) Commission in 2017 to select a new voting system in time for the 2020 election. Ultimately, U.S. District Judge Amy Totenberg ruled that the trade-off between election integrity and the feasibility of making changes before the impending election tipped in favor of continued use of DREs for the 2018 election. Judge Totenberg ruled that the plaintiffs provided sufficient evidence that DRE voting has the potential to cause irreparable harm to voters, but that the burden of switching to paper ballots so close to the election could cause even more harm to voters by causing bureaucratic confusion.

Ultimately, any chaos or problems that arise in connection with a sudden rollout of a paper ballot system with accompanying scanning equipment may swamp the polls with work and voters—and result in voter frustration and disaffection from the voting process. There is nothing like

bureaucratic confusion and long lines to sour a citizen. And that description does not even touch on whether voters themselves, many of whom may never have cast a paper ballot before, will have been provided reasonable materials to prepare them for properly executing the paper ballots.

Judge Totenberg also noted that the evidence and testimony “indicated that the Defendants and State election officials had buried their heads in the sand” [34].

Secretary of State Kemp refused to recuse himself from overseeing the election in which he ran for Governor, a clear conflict of interest [39]. Election Day voting in November 2018 was conducted on paperless DREs. Machines in four polling places in Gwinnett County malfunctioned, forcing voters to use paper ballots, which caused some voters to wait four hours to cast their vote [16]. Reported vote totals were anomalous: the rate of undervotes in the Lieutenant Governor (LG) contest was unusually high compared to historical LG races and compared to other statewide contests on the ballot, and the undervote rate was far higher for DREs than for paper ballots. The Coalition for Good Governance brought another lawsuit against the Georgia Secretary of State, calling for a redo of the LG contest [29]. Statistical evidence of anomalies in this election, presented in that lawsuit, is discussed below in Section 5.

After the election, Kemp’s office planned to certify the election results six days before state law required it, omitting nearly 27,000 provisional ballots. Provisional ballots are cast by voters whose registration or identification is in question; deliberately omitting provisional ballots is one way to disenfranchise voters. It would have ensured that the margin between Kemp and his opponent Stacey Abrams remained large enough to avoid a runoff election [4]. A civil rights group sued to delay the certification, and Judge Totenberg ruled against Kemp, ordering election officials to review the provisional ballots.

The SAFE Commission was scheduled to recommend a new voting system in January, 2019. In early January, the Democratic Party of Georgia called on Kemp to delay any decision to purchase new voting systems as more misbehavior came to light: now-Governor Kemp appointed Charles “Chuck” Harper, chief lobbyist for ES&S (the voting machine company that eventually acquired Diebold), as Deputy Chief of Staff in the Governor’s office [23].

5 Evidence of Malfunctioning DREs in 2018

While the controversy surrounding the Governor’s race did not result in anomalous election results, the LG race did. Shortly after the November 2018 election, The Coalition for Good Governance filed another lawsuit against the new Secretary of State, demanding a redo of the LG vote. The plaintiffs blamed malfunctioning DREs for an unusually high number of undervotes in the LG race, but not in others [29]. The judge overseeing the case initially agreed to let the plaintiffs examine the memory, but not the programming, of machines in three counties. She eventually dissolved this agreement and dismissed the case [45].

The plaintiffs did not specify the cause of the malfunction—faulty programming, poor electronic ballot design, hacking, or something else [29]. Numerous voters reported irregularities when attempting to cast their vote for LG on DREs, including many who reported that the race did not appear on their ballot until they were shown the review screen. Without forensic evidence, it is impossible to determine exactly what happened.

This section gives three lines of statistical evidence that DREs did not record every vote properly in this election. First, in 101 of Georgia’s 159 counties, the rate of undervotes in the LG race was much higher among DRE votes (those cast on Election Day and advance in-person) than on (paper) absentee ballots. (For other statewide contests, the undervote rates are similar across modes of voting in nearly all counties.) Second, in Fulton County, higher differential undervote rates tended to occur in precincts where a larger percentage of registered voters were Black. Third, on six of seven machines in the Winterville Train Depot polling place in Clarke County, Democrats got the majority of votes in every statewide contest, matching the overall results at the polling place. On the seventh machine, Republican candidates got a majority in every statewide contest.

Permutation tests show that these three anomalies are implausible unless something went wrong. Permutation tests require a minimum of assumptions, which can make them appropriate and convincing in situations where standard parametric tests require unrealistic or counterfactual assumptions, for instance, assumptions that voter preferences follow a parametric model, such as multinomial logistic. In contrast, the permutation tests we use treat one characteristic, such as the mode by which a ballot containing an undervote was cast or the machine on which a ballot was cast, as an arbitrary label that might as well have been assigned at random. Software implementing the tests reported here can be found at <https://github.com/pbstark/EvoteID19-GA>.

5.1 Undervotes for Lieutenant Governor

Undervotes occur when a voter selects fewer candidates in a contest than the contest rules allow, for instance, not voting for any candidate in a winner-take-all contest. The rate of undervotes tends to increase for “down-ticket” contests compared to major contests such as presidential and gubernatorial contests. In Georgia in 2018, the LG race had a 4% undervote rate, while the next contest on the ballot had an undervote rate of 1.4%. Moreover, this pattern appeared only in votes cast on DREs—Election Day votes and advance in-person votes.

Data were downloaded from Clarity Elections, the private sector vendor that reports official election results on behalf of the Georgia Secretary of State.²³ Data included the total number of ballots cast in each county and the number cast by each mode of voting (e.g. by mail) for each candidate by county. The file

² The fact that this crucial election function is outsourced without oversight might give the reader pause.

³ <https://results.enr.clarityelections.com/GA/91639/222278/reports/detailxml.zip>, downloaded in January 2019.

did not report ballots cast in each county by mode of voting. In order to calculate the number of undervotes, we assumed that the total number of ballots cast by county and mode of voting equalled the maximum number of votes cast in *any* contest for that county and mode of voting.

While political preferences might differ systematically between voters who vote by mail (on paper) and those who vote in person (on DREs), there is no reason to think that interest in a *contest* should differ across those groups. The usability literature suggests that DREs ought to help people of disparate education and ethnicities vote correctly, in which case, the undervote rate on DREs should be *lower* than the rate for paper ballots [31]. If so, then it is reasonable to treat the mode of voting as a label assigned randomly to ballots in such a way that the number of ballots cast on DREs and the number cast on paper is fixed (conditioned to be equal to the actual numbers). The number of undervotes in a contest among DRE votes then has a hypergeometric distribution. Under the alternative that undervotes are more likely on DREs, we would expect to see more undervotes on DREs (and fewer on paper ballots) than the hypergeometric distribution predicts.

In 101 of 159 Georgia counties, the difference in undervote rates between mail votes and DRE votes in the LG race is statistically significant at level 0.01%. In contrast, in the 8 statewide contests further down the ballot, the difference is statistically significant in no more than 5 counties. Table 1 shows the counts.

Table 1. Counties with statistically significant ($p < 0.0001$) disparities in undervote rates between paper ballots and DREs.

Contest	Counties with significant undervote rate disparities
Lt. Governor	101
Secretary of State	4
Attorney General	4
Commissioner of Agriculture	5
Commissioner of Insurance	4
State School Superintendent	5
Commissioner of Labor	2
Public Service Commission District 3	4
Public Service Commission District 5	4

5.2 Undervotes and Race in Fulton County

Undervote rates on touchscreen voting machines were reported to be higher in predominantly Black precincts across the state [10]. If so, that is evidence that security and usability issues with DREs disparately impact historically disadvantaged groups. We investigated this issue in Fulton County, which includes most of the capital, Atlanta, and had over 424,000 voters in November 2018.

Precinct-level reported vote totals were downloaded from the Clarity Election site that reports official results for the Georgia Secretary of State.⁴ Data included total votes cast for each candidate by each mode of voting, in each precinct within Fulton County. As with the statewide data, we estimated the number of undervotes by subtracting the votes from the maximum number of votes in any contest, by mode of voting and precinct.

Voter turnout data were downloaded from the Secretary of State’s website.⁵ From these data, we computed the percentage of registered voters who were Black in each precinct.

A permutation test was used to assess the correlation between the difference in undervote rates between voters who used paper ballots and voters who voted electronically and the percentage of registered voters who were Black. Of the 373 precincts in Fulton County, we restricted analysis to the 302 precincts in which at least 10 people voted electronically and at least 10 voted on paper.

The undervote rate was substantially lower for voters who used paper ballots than for voters who voted electronically, by an amount that—on average—was larger in precincts with a larger percentage of Black registered voters. Table 2 shows the correlation between the difference in undervote rates and the percentage of registered voters who are Black. p -values are for randomized permutation tests with 10,000 replications, carried out using the Python `permute` package.⁶ Small p -values for multiple statewide contests could be explained by voter behavior; prior research suggests that Black voters may intentionally undervote at a higher rate than other voters, and may cast valid votes at a rate that is lower than the rate for the general electorate [11,31]. However, it is notable that the correlation for the Lieutenant Governor’s contest is more than twice what it is for any other contest.

5.3 Party Preferences in Winterville Train Depot Polling Place

A citizen photographed printed poll tapes from the seven DRE machines in the Winterville Train Depot polling place in Clarke County. The photographs were transcribed to CSV and double checked by a second person.⁷

The Winterville Train Depot polling place is just one polling place in Georgia where a member of the public photographed poll tapes posted at the precinct after the polls closed. It was not selected at random, but neither was there particular reason to suspect problems there. There is no reason to believe that problems are confined to this polling place—where then-Secretary of State Kemp himself voted—but even if they were, any anomaly is of concern.

The DREs in the precinct recorded comparable numbers of voters (117, 135, 131, 133, 135, 144, 135). In this polling place, Democratic candidates won a

⁴ <https://results.enr.clarityelections.com/GA/Fulton/91700/221530/reports/detailxml.zip>, downloaded in January 2019.

⁵ http://sos.ga.gov/admin/uploads/PRECINCT_Nov_2018.zip, downloaded in January 2019.

⁶ <http://statlab.github.io/permute>

⁷ The data were submitted as evidence in [29].

Table 2. Correlation between the difference in undervote rates and percentage of registered voters who are Black, for the 10 statewide contests in Georgia in November 2018, in Fulton County.

Contest	correlation	<i>p</i> -value
Governor	-0.134	0.9903
Lt. Governor	0.557	0.0001
Secretary of State	0.092	0.0582
Attorney General	0.078	0.0902
Commissioner of Agriculture	0.207	0.0003
Commissioner of Insurance	0.246	0.0001
State School Superintendent	0.154	0.0050
Commissioner of Labor	0.041	0.2376
Public Service Commission District 3	0.042	0.2329
Public Service Commission District 5	0.125	0.0145

majority in all ten statewide contests. Every DRE reported a majority of votes for the Democratic candidate in every statewide contest except machine 3, which reported a majority for the Republican candidate in every statewide contest.

If voters were directed to DREs as if at random, then the number of voters who used different machines should be roughly equal, as should the percentage of votes for each candidate. Conditional on the number of ballots on each machine and the total number of votes for each candidate across machines, all permutations of votes across machines are equally likely under the null hypothesis. We performed a two-sided permutation test using the difference between the expected and actual fraction of Republican votes in each contest as the test statistic. Permutations were done using the `cryptorandom` pseudo-random number generator for Python⁸. The *p*-values for different contests were combined using Fisher’s combination function to obtain a global *p*-value on the assumption that the distribution of Fisher’s combining function under the null hypothesis is chi-square. That would be true if votes in different contests were independent; however, voters tend to vote along party lines. If ballot-level data were available, a Fisher’s combining function could be calibrated to take that correlation into account. However, the poll tapes give only totals by contest. Hence, while *p*-values for individual contests are on a firm statistical footing, the global *p*-value should be viewed as suggestive rather than precise.

On the assumption that voters were directed to DREs as if at random, the chance any of the seven machines would show disparities as large as machine 3 did in individual contests ranges from less than 1% to approximately 15%. Seven of the ten values are significant at level 5% or below; see Table 3. The global *p*-value for the ten tests is 0.00009%.⁹

⁸ <http://statlab.github.io/cryptorandom>

⁹ As mentioned above, the assumptions under which Fisher’s combining function has a chi-square distribution may not hold, so the global *p*-value should be viewed as suggestive.

Table 3. Consistency of results across DREs in Winterville Train Station Polling Place and consistency of results if D and R were flipped on machine 3.

Contest	p -value	p -value if machine 3 were flipped
Governor	0.114	0.464
Lt. Governor	0.025	0.795
Secretary of State	0.018	0.450
Attorney General	0.151	0.543
Commissioner of Agriculture	0.026	0.734
Commissioner of Insurance	0.030	0.604
State School Superintendent	0.097	0.807
Commissioner of Labor	0.008	0.797
Public Service Commission District 3	0.046	0.280
Public Service Commission District 5	0.025	0.939

These results are entirely driven by the results on machine 3. If the Democratic and Republican party labels were flipped on that machine, the anomaly disappears, and the global p -value for the ten contests becomes 97%. For individual contests, no p -value is then below 0.280, compared with values as small as 0.008 (and seven values below 5%) for the actual poll tapes. See Table 3.

These tests strongly suggest that machine 3 had some software or hardware problem: misconfiguration, error, defect, hack, or malfunction. The most plausible explanation is that misconfiguration caused votes for Republican candidates to be recorded as votes for Democratic candidates, and vice versa.

6 Conclusion

The 2018 midterms demonstrated that election integrity in Georgia remains fraught. In the weeks leading up to the election and for weeks after, citizens challenged the Secretary of State’s treatment of provisional ballots and voter registrations, alleging that these practices were intended to disenfranchise minority voters. Touchscreen DRE voting machines were used statewide, even after security experts voiced their concerns and a nonprofit organization sued the state to replace DREs with hand-marked paper ballots. There is evidence that some DREs malfunctioned in the election; statistical anomalies suggest that DREs failed to record a large percentage of votes cast in the Lieutenant Governor’s race, and that “missing votes” were more frequent in jurisdictions with large African American populations [10]. The Secretary of State has refused to investigate these issues. Some particular anomalies (i.e., the Winterville Train Depot data) are most easily explained by “vote flipping,” in which the DRE recorded votes for one candidate as votes for the candidate’s opponent.

Lawmakers are poised to replace the state’s DREs with a new system: either hand-marked paper ballots with optical scanners, using touchscreen ballot-marking devices (BMDs) for accessibility, or BMDs for all voters. In February

2019, the state legislature voted to purchase BMDs statewide [21]. While BMDs do produce a paper record, they are more expensive than opscan systems,¹⁰ and they are neither as reliable nor as secure as hand-marked paper ballots and opscan systems. Among other issues, BMD malfunctions can prevent voting on Election Day; inadequate provisioning of equipment can produce long lines; there is evidence that voters cannot and do not reliably verify their BMD selections; and BMDs require the same trust in software as DREs, with no practical recourse if machines malfunction and little possibility that outcome-changing errors will be detected [1,28]. The SAFE Commission's only security expert, Prof. Wenke Lee, warned against BMDs.

House Minority Leader Bob Trammell expressed his stance on the evidence for hand-marked paper ballots [21]:

It's unequivocally clear that cybersecurity experts have expressed concerns about the ballot-marking devices. It comes down to whether you think the opinion of election officials . . . is more important than the issue of credentialed experts in the field talking about a material risk to the voting process.

References

1. Appel, A., DeMillo, R., Stark, P.B.: Ballot-marking devices (BMDs) cannot assure the will of the voters. Social Science Research Network (2019)
2. Blaze, M., Braun, J., Hursti, H., Hall, J.L., MacAlpine, M., Moss, J.: DEFCON 25 Voting Village report. Tech. rep., DEFCON (2017), <https://www.defcon.org/images/defcon-25/DEFCON%2025%20voting%20village%20report.pdf>
3. Blaze, M., Braun, J., Hursti, H., Jefferson, D., MacAlpine, M., Moss, J.: DEFCON 26 Voting Village report. Tech. rep., DEFCON (2018), <https://www.defcon.org/images/defcon-26/DEF%20CON%2026%20voting%20village%20report.pdf>
4. Blinder, A.: Federal Judge Delays Certification of Georgia Election Results. The New York Times (2018), <https://nyti.ms/2DiWDzx>
5. Bowen, D.: Top-to-Bottom Review of voting machines certified for use in California. Tech. rep., California Secretary of State (2007), <https://www.sos.ca.gov/elections/voting-systems/oversight/top-bottom-review/>
6. Chronicle, T.A.: Voting machine maker hires former state election chief. The Augusta Chronicle (2006), <https://www.augustachronicle.com/article/20061224/NEWS/312249946>
7. Everett, S.P.: The Usability of Electronic Voting Machines and How Votes Can Be Changed Without Detection. Ph.D. thesis, Rice University, Houston, Texas (2007), <https://scholarship.rice.edu/handle/1911/20601>
8. Freeman, S.F., Bleifuss, J.: Was the 2004 Presidential Election Stolen?: Exit Polls, Election Fraud, and the Official Count. Seven Stories Press (2006)
9. Harnik, A., Press, A.: Officials Scrap Plan To Cut Most Polling Places In Majority Black Ga. County. WABE (2018), <https://www.wabe.org/officials-scrap-plan-to-cut-most-polling-places-in-majority-black-ga-county/>

¹⁰ The State of Georgia has claimed otherwise, but their analysis was deeply flawed, omitting costs associated with BMDs and overstating the cost of printing ballots, among other things. See [25].

10. Harriot, M.: Thousands of Black Votes in Georgia Disappeared. The Root (2019), <https://www.theroot.com/exclusive-thousands-of-black-votes-in-georgia-disappea-1832472558>
11. Herron, M.C., Sekhon, J.S.: Overvoting and representation: An examination of overvoted presidential ballots in broward and miami-dade counties. Electoral Studies **22**(1), 21–47 (2003)
12. Hursti, H.: Critical security issues with Diebold optical scan design. Tech. rep., Black Box Voting (2005)
13. Hursti, H.: Critical security issues with Diebold TSx. Tech. rep., Black Box Voting (2006)
14. Jones, D.W., Simons, B.: Broken Ballots: Will Your Vote Count? CSLI Publications (2012)
15. Kennedy, R.J.: Will the Next Election Be Hacked? Electronic Voting Machines Can't be Trusted. Rolling Stone (2006), <https://www.organicconsumers.org/news/robert-kennedy-jr-will-next-election-be-hacked-electronic-voting-machines-cant-be-trusted>
16. Lockhart, P.R.: Voting hours in parts of Georgia extended after technical errors create long lines. Vox (2018), <https://www.vox.com/policy-and-politics/2018/11/6/18068492/georgia-voting-gwinnett-fulton-county-machine-problems-midterm-election-extension>
17. McDaniel, P., Blaze, M., Vigna, G.: EVEREST: Evaluation and validation of election-related equipment, standards and testing. Tech. rep., Ohio Secretary of State (2007), <http://siis.cse.psu.edu/everest.html>
18. Nadler, B.: Voting rights become a flashpoint in Georgia governor's race. AP News (2018), <https://apnews.com/fb011f39af3b40518b572c8cce6e906c>
19. Nakashima, E.: In Georgia, a legal battle over electronic vs. paper voting. The Washington Post (2018), <https://wapo.st/2QADbm8>
20. National Institute of Standards and Technology Auditability Working Group: Report of the auditability working group. Tech. rep., NIST (2015), https://www.eac.gov/assets/1/28/AuditabilityReport_final_January_2011.pdf
21. Niese, M.: Bill to buy new Georgia voting machines clears committees. The Atlanta Journal-Constitution (2019), <https://www.myajc.com/news/state--regional-govt--politics/new-georgia-voting-machines-approved-house-committee/avz21tiapWPwM1Qx4bq3AP/>
22. Niese, M., Prabhu, M.T., Elias, J.: Voting precincts closed across Georgia since election oversight lifted. The Atlanta Journal-Constitution (2018), <https://www.myajc.com/news/state--regional-govt--politics/voting-precincts-closed-across-georgia-since-election-oversight-lifted/bBkHxptlim0Gp9pKu7dfn/>
23. Party, G.D.: Breaking: Democratic Party of Georgia Calls on SAFE Commission to Delay Vote Following News That Voting Machine Lobbyist Is Longtime Kemp Crony (2019), <https://www.georgiademocrat.org/2019/01/breaking-democratic-party-of-georgia-calls-on-safe-commission-to-delay-vote-following-news-that-voting-machine-lobbyist-is-longtime-kemp-crony/>
24. Peha, J.: Touch-and-Go Elections: The perils of electronic voting. The Nation (2006), <https://www.thenation.com/article/touch-and-go-elections-perils-electronic-voting/>
25. Perez, E., Miller, G.A.: Georgia State Election Technology Acquisition: A Reality Check. Tech. rep., OSET Institute (2018)
26. Riley, M., Robertson, J.: Russian Hacks on U.S. Voting System Wider Than Previously Known. Bloomberg (2017), <https://www.bloomberg.com/news/articles/2017-06-13/russian-breach-of-39-states-threatens-future-u-s-elections>

27. Stahl, J.: Georgia Destroyed Election Data Right After a Lawsuit Alleged Its Voting System Was a Mess. Why? Slate Magazine (2017), <https://slate.com/technology/2017/10/georgia-destroyed-election-data-right-after-a-lawsuit-alleged-the-system-was-vulnerable.html>
28. Stark, P.B.: Ballot-marking devices (BMDs) are not secure election technology (2019), <https://www.stat.berkeley.edu/~stark/Preprints/bmd19.pdf>
29. Superior Court of Fulton County, State of Georgia: Coalition for Good Governance, Martin, Duval, and Dufort v. Crittenden (2019), 2018-CV-3134-18
30. Supreme Court of the United States: Shelby County v. Holder (2013), 12-96
31. Tomz, M., van Houweling, R.P.: How does voting equipment affect the racial gap in voided ballots? American Journal of Political Science **47**(1), 46–60 (2003)
32. Torres, K.: Federal lawsuit alleges Georgia blocked thousands of minority voters. The Atlanta Journal-Constitution (2016), <https://www.myajc.com/news/state--regional-govt--politics/federal-lawsuit-alleges-georgia-blocked-thousands-minority-voters/EKb979oRoBe4yJ3Uo1nDfP/>
33. Torres, R.: Report Concerning the Examination Results of Election Systems and Software EVS 6021 with DS200 Precinct Scanner, DS450 and DS850 Central Scanners, ExpressVote HW 2.1 Marker and Tabulator, ExpressVote XL Tabulator and Electionware EMS. Tech. rep., Commonwealth of Pennsylvania Department of State (2018)
34. United States District Court for the Northern District of Georgia, Atlanta Division: Curling v. Kemp (2018), 1:17-CV-2989-AT (Order Denying Motion to Dismiss)
35. United States District Court for the Northern District of Georgia, Atlanta Division: Curling v. Kemp (2018)
36. Vasilogambros, M.: Polling Places Remain a Target Ahead of November Elections. Stateline (2018), <https://pew.org/2MCsiBT>
37. Voters Organized for Trusted Election Results in Georgia: Georgia Unverifiable Voting System Chronology. VOTER GA (2014), <https://voterga.org/history/>
38. Whitesides, J.: Polling places become battleground in U.S. voting rights fight. Reuters (2016), <http://reut.rs/2cKzOaZ>
39. Williams, V.: Georgia groups call on GOP gubernatorial nominee Brian Kemp to step down as the state’s elections chief. The Washington Post (2018), <https://wapo.st/2MrHZHP>
40. Wofford, B.: How to Hack an Election in 7 Minutes. POLITICO Magazine (2016), <https://politi.co/2K2OGOv>
41. Zetter, K.: Did E-Vote Firm Patch Election? Wired (2003), <https://www.wired.com/2003/10/did-e-vote-firm-patch-election/>
42. Zetter, K.: Will the Georgia Special Election Get Hacked? POLITICO Magazine (2017), <http://politi.co/2heBRW2>
43. Zetter, K.: The Crisis of Election Security. The New York Times (2018), <https://nyti.ms/2N3hoAh>
44. Zetter, K.: Was Georgia’s Election System Hacked in 2016? POLITICO Magazine (2018), <https://politi.co/2moAWUS>
45. Zetter, K.: Georgia voting irregularities raise more troubling questions about the state’s elections. POLITICO (2019), <https://politi.co/2SOlvas>

E-Voting, practical approaches

Cybersecurity and Electoral Integrity: The Case of Ukraine, 2014-present

Beata Martin-Rozumilowicz and Thomas Chanussot

International Foundation for Electoral Systems (IFES)
{[bmartinrozumilowicz](mailto:bmartinrozumilowicz@ifes.org),[tchanussot](mailto:tchanussot@ifes.org)}@ifes.org

Abstract. This article examines the historical context and framework of cybersecurity efforts and their impact on electoral integrity in the case of Ukraine since the Revolution of Dignity in 2014. It looks at the spectrum of incursions and attacks that the country has faced on the elections-aspect of its critical infrastructure, documenting historically the context that the country currently faces in this sphere. It examines the specific adversaries that have been involved in the Ukrainian cybersecurity space, the measures taken, and developments made to better protect the electoral process and buttress public confidence. Finally, it analyzes the current threats and challenges Ukraine faces and proposes measures to address them in the future. As such, this article presents a cohesive snapshot of a country that is at the frontlines of an iterative process in which lessons learned are then applied by adversaries to target other electoral democracies. Thus, it presents an important point of understanding as this field of research develops.

Keywords: Cybersecurity· Ukraine· Elections

1 Introduction

The challenges faced by Ukraine in the space of cybersecurity and elections are not unique. Numerous electoral democracies over the last decade and a half have been facing increased attacks by adversaries, both foreign and sometimes domestic. Ukraine, however, has been in a rather unique position to be a sort of ‘ground zero’ for weapons testing in this space. In what is clearly an iterative process, lessons from one game are analyzed, honed and re-deployed in further rounds and often in different jurisdictions. One can draw a relatively straight line from the advanced persistent threat (APT) attacks on Ukrainian critical infrastructure (CI) and its election administration in 2014 and 2015, further to the French and German incursions in April and May 2015 and then to the attacks in the 2016 U.S. presidential race, and beyond.

With this in mind, the authors believe that a thoughtful and analytical examination of the history of cybersecurity incursion and response is timely in a global perspective where such threats are becoming increasingly important to

the integrity of elections. Where once, focus was on the vulnerabilities of voting machines and electronic voting systems, increasingly ancillary systems (such as voter registration systems, political parties, media or local authorities) have proven to be equally vulnerable to such cyber-related threats. In fact, in the latest iteration of cybersecurity intrusions, it is precisely the ‘weak spots’ and soft targets that are being sought out by adversaries as states place more focus on shoring up their cyber defenses in the electoral sphere.

This paper will seek to analyze and understand the dynamics of the cyber threat to elections during three time periods: 1) pre-2008, where most attacks were DDoS-oriented and scattershot; 2) 2008-2014, where such cyberattacks were linked to military incursions in a new type of hybrid warfare, and 3) post-2014, where elections were added as a focal point of attacks in an overall CI-oriented strategy on the part of adversaries (Ukraine and U.S., as specific examples). From these exercises, lessons-learned are being drawn and ostensibly re-focused before being deployed in key upcoming elections such as the European Parliamentary elections in May, the Ukrainian parliamentary elections in the autumn (which are likely to be much a focal point of potential attack), and the U.S. presidential elections in 2020.

Thus, the paper examines the cybersecurity challenges in Ukraine from a historical perspective, before turning to an in-depth analysis of the cybersecurity challenges presented during the Ukrainian presidential elections held in March and April 2019. The paper then turns to an examination of the current threats and challenges facing the country prior to the upcoming parliamentary elections to gain a perspective of how the threats may be mutating. Finally, it attempts to draw lessons for the future, both in the specific Ukraine-oriented scenario, as well as at a more global level.

As such, it is hoped that this paper presents a unique perspective, which is at once specific to Ukraine, while also having global implications. It represents a singular examination of the Ukrainian and related cases in the region, which has been sorely lacking for some time and with unique, first-hand evidence and data. As such, the authors believe that represents a significant advance in this and related fields of a holistic and integral approach to the question of cybersecurity in an electoral context.

2 Ukraine and its Cybersecurity Challenges in Historical Perspective

It was election night of the Presidential race in the immediate aftermath of the February 2014 Revolution of Dignity (or EuroMaidan in common parlance). It was the first experience that Ukraine was having of holding a poll since the ouster of Viktor Yanukovich and the end of ever-increasing deterioration the country had faced since his taking power in 2010/11. It would be a key test of

Ukrainian democracy and its ability to reflect the will of the people through a legitimate nationwide vote.

Yet, as certain members within the Central Election Commission (CEC) watched the results being reflected on their official website, they knew there was a problem. The commission IT department had been under a massive Distributed Denial of Service (DDoS) attack over the last 24-hours.¹ Later, on election day, Commissioners and CEC staff looked in horror as Dmytro Yarosh was announced as the winner of the election on Russian TV from the CEC official website; an announcement clearly conflicting with the results of the electronic tabulation system.² [1]

The CEC deployed all resources at their disposal, as the results were isolated and located in a separate website, and the commission was able to identify and block the incursion and remove the fake results. Without surprises, Petro Poroshenko, who had been leading the results in the polls led the tabulated votes and was eventually identified as the winner in the official results, which are still based by law on paper protocols. The crisis had been averted.

Yet, the damage had been done. The narrative that the foreign adversary identified behind the attacks had been promoting was that Ukraine was on the brink of chaos, a proto-fascist state which could not even manage to hold credible elections. The false results fed into this narrative and attempted to undermine public trust and confidence in the results. It would prove to be a strategy that APT 28 and 29 were linked to in future elections, including the U.S. 2016 presidential election. Yet, Ukraine in 2014 was "ground zero". It was the inception point from which future strategies to challenge elections on the cyber side would develop.³

This was not the first challenge European states faced from foreign adversaries. Incidents of such "weapons-system" testing occurred in Estonia in 2007 (and even earlier), where DDoS attacks nearly shutdown the country's internet infrastructure (including websites of parliament, banks, ministries, newspapers,

¹ DDoS is a flood of fake internet traffic coming from different sources over the internet, with the intent to overwhelm the network and make the services, the CEC result website, unavailable to the public.

² If it had not been discovered and removed, the fake results would have portrayed ultra-nationalist Right Sector party leader Dmytro Yarosh as the winner with 37 percent of the vote (instead of the 1 percent he actually received) and Petro Poroshenko (the actually winner with a majority of the vote) with just 29 percent, Ukraine officials told reporters the next morning.

³ This would be replicated in the parliamentary elections in the autumn of 2014, which saw another attack launched against the election infrastructure, with hackers attacked Ukraine's CEC website on the eve of elections. Minimal damage was done in these attacks, which Ukrainian security officials blamed on DDoS attempts. The Ukrainian side, however, had learned from its previous presidential hack experience and had put measures in place.

broadcasters). For a country moving whole-heartedly and at a fast pace to a philosophy of e-governance and e-elections, this was a serious shock.⁴ A similar tactic was employed in Lithuania in June 2008 following the adoption of a law prohibiting the use of Soviet symbols. Some 300 websites were subject to DDoS attacks and vandalized, including those of key government agencies.⁵ This type of strategy was stage one, in an ever-evolving game plan that was further developed in the next iteration.

This attack plan was replicated in Georgia in 2008 during the first utilization of hybrid warfare tactics (where cyberattacks were used in concert with traditional military operations) when key government sites in the country were brought down on the eve of the Russian invasion. This strategy was also replicated in the invasion of Crimea and support to Donbas in 2014. It has also become the hallmark of this wider-scope attack, with the aim of bringing down a maximal of critical infrastructure to achieve maximum chaos.

This new style of attack was supplemented by additional tests targeting CI sectors. In 2015 and 2016, the Ukrainian energy grid as well as its train services were brought down by malware with a clear signature pattern). [2] These types of malware would later be discovered as far away as the U.S. electricity grid (Sandworm, BlackEnergy [3]).⁶ Such malware targeted bad cyber hygiene practices through spearphishing to gain and exploit administrator access to the CI systems and shut actual administrators out. It took nearly four months to properly address the incursion of the energy grid in Ukraine.

In Ukraine 2014, for the first time, attacks were directed at electoral systems and administration, specifically. This time, the tactics of the previous two iterations were added to, with a specific target focus. Rather than just causing disruption and sowing chaos, this latest iteration tried to project different electoral results to an entire population.⁷ This would mark a third generation of cyber incursion in an attempt at foreign interference. This is presented in the table below, showing the additive properties with each iteration.

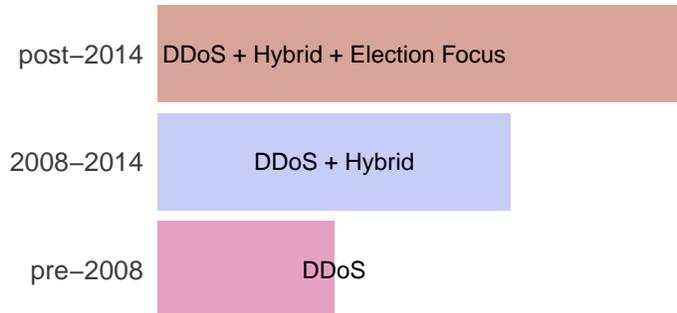
⁴ Most of the attacks were distributed denial of service (DDoS) attacks ranging from single individuals using various methods like ping floods to expensive rentals of botnets used for spam distribution. 20-year-old Dmitri Galushkevich was eventually charged in the attack.

⁵ Government, political party and business webpages were defaced by hackers with hammer-and-sickles and five-pointed stars.

⁶ Sandworm has shown a special interest in power grids. FireEye has tied the group to a series of intrusions on American energy utilities discovered in 2014, which were infected with the same Black Energy malware Sandworm would later use in its Ukraine attacks.

⁷ An earlier DDoS attack in March, which was 32 times larger than the previously largest attack when Russia invaded Georgia was coupled with Russian-armed, pro-Russian rebels seizing control of Crimea.

Table 1: Cyber Interference in Elections



The innovation of combining targeted CI incursions with the exploitation of soft targets due to bad cyber hygiene was ultimately exploited in the 2016 U.S. presidential election, where one soft target (namely, John Podesta [4]) was used to effectively bring down the Democratic National Committee (DNC) and sow doubt amongst American citizens as to the efficacy of their democracy. This had its seeds in 2014 Ukraine and prior. To prevent making the same mistakes in the future, we need to learn from this history.

3 The 2019 Ukrainian Presidential Election and Cybersecurity Challenges

In September 2018, six months before the March 2019 presidential elections in Ukraine, the CEC had yet to finalize preparations for its cyber-defense or take necessary measures that would protect its infrastructure against another potential cyber-attack. While the election commission did take important measures in establishing cybersecurity best practices to protect its network, these efforts only scratched the surface of the extensive actions the organization would need to take, in co-operation with Ukrainian cyber agencies, to improve the security of the electoral process. Moreover, the CEC was in the final months of preparation for an election that was widely seen in the media and among experts [5] [6] as the next testing ground for new types of attacks that would eventually be unleashed in other countries.

3.1 Pre-election preparation

The CEC Secretariat is a permanent organization with over 250 employees working in 15 departments. It has a dedicated IT team working on operational and cybersecurity needs. As part of the modernization and cybersecurity efforts beginning in 2015, the IT department segmented the CEC office network and the election information systems, disconnecting these systems from the regular day-to-day emails and internet communication of the CEC. Notably, election information systems in Ukraine ensure 2 critical functions [7]:

- The Result Management System (RMS) is used by clients in the District Election Commissions (DEC) on a dedicated line isolated from the internet. Its purpose is to capture electronically the polling station results (polling is paper based in Ukraine).
- The State Voter Register (SVR) database is used by 761 local administrative offices to update the voter list monthly.

In addition to isolating these critical systems from the internet and the rest of the organization’s network, the CEC installed modern and comprehensive network monitoring systems, partly owing to its collaboration with the government security agencies. Most importantly, the CEC, including the newly appointed Chairperson and commissioners, was significantly more aware of the risk cyber-attacks could represent to the integrity of the presidential elections.

Despite these efforts and recognition of the importance that cybersecurity would play by key figures, the infrastructure and the election information systems used by the CEC remained widely unaudited by the end of 2018. Little progress had been made to replace outdated equipment or to review and establish necessary protective measures to improve the CEC’s cybersecurity posture. With an unpredictable Russia on the border, and many cybersecurity vulnerabilities still unaddressed, Ukrainian cybersecurity resilience and protection efforts remained an urgent priority as 2019 - and the presidential election - dawned.

3.2 Involvement of the cybersecurity agencies, a coordinated effort

Ukraine’s CEC is a formally independent body and is responsible for ensuring the cybersecurity of its own systems. It is assisted during the election period by the State Service of Special Communications and Information Protection (SSSCIP) and the Security Service of Ukraine (SBU, Ukrainian acronym) and the National Police.

Like in most democracies, election infrastructure in Ukraine is not considered critical infrastructure according to the Law on Cybersecurity in Ukraine (in force since May 2018). This legal reality creates a different structure of cooperation between the CEC and cybersecurity agencies. Several cybersecurity think-tanks and international agencies [8] have advocated in favor of the classification of election systems, processes and infrastructures as critical infrastructure to ensure the necessary cybersecurity measures are put in place. This model could be useful in Ukraine, although the current bi-lateral Memorandums of Understanding regulating cooperation between the CEC and cybersecurity agencies have so far been sufficient.

Furthermore, the SSSCIP and the SBU have different mandates with regards to the protection of the CEC infrastructure. While the SSSCIP is responsible for implementing technical and organizational measures to prevent, detect and

respond to cyber-attacks, the SBU as a law-enforcement authority and government security agency, oversees system penetration testing (pen-testing) and provides support to the CEC with regards to detection of criminal intrusions on the election systems.

Both the SSSCIP and SBU cooperated with the CEC before, notably in mitigating and investigating the 2014 cyber-incidents. However, the preparation for the 2019 election required more coordinated and strategic approaches. This support came in the form of a working group, established several months prior to the election, composed of technical experts from the CEC, SSSCIP, Security Bureau of Ukraine and National Police [9]. A similar working group had been created before, and Ukraine managed to hold elections in difficult circumstances as the war unfolded [10]. The working group initially met monthly, but at the request of the CEC gathered more often during the weeks leading up to the elections.

During the pre-election period, the CEC and the cybersecurity agencies also participated in tabletop crisis simulation [11] [12]. During these activities, both managerial and technical crisis response capacity was tested and adapted based on the outcome of simulations. In addition to increasing the technical skills, these exercises also create a sense of common responsibility among stakeholders in ensuring cybersecurity.

3.3 The challenge of replacing outdated equipment

Outdated equipment was identified as a risk for the CEC infrastructure immediately in the aftermath of the 2014 elections. Although never confirmed by the manufacturer and not reported by any official, a 0-day vulnerability [13] was allegedly reported being used by the hacking group CyberBerkut on the old Cisco ASA software used by the CEC. This class of device was no longer supported by Cisco. Devices reaching “end-of-support” can put an organization at a higher risk of a security breach, as newly discovered vulnerabilities such as the one discovered by CyberBerkut are not being patched. By mid-2018, although several plans for replacement had been made, procurement difficulties were still getting in the way, leaving gaping holes in the security of the election commission network, in particular, the Result Management and the State Registry of Voters systems.

With the assistance of international organizations, the CEC received replacements for the outdated equipment in the 11th hour of their election preparation. While the new equipment did resolve the issue of un-patchable vulnerable network devices, the late delivery introduced a new kind of risk. Despite the short timeframe, this new equipment needed to be set up, configured, tested and audited to ensure functionality prior to the presidential elections. Procurement of the needed software and hardware shifted the risk, rather than eliminating it, as the new equipment had to be integrated into the existing infrastructure without creating performance bottlenecks, and more importantly, creating new or additional security vulnerabilities.

The SSSCIP support to the CEC following 2014 cyberattacks [13] was limited to forensic investigation and incident response management. In February 2019 and upon the delivery of the equipment, both the SSSCIP and the SBU provided support to the CEC and ensured that the equipment was installed, configured and audited ahead of the first round of the presidential election on March 31, 2019.

3.4 Cybersecurity awareness

Humans are often the weakest link when it comes to defending cybersecurity, while ideally, they should be the first line of defense. Studies report that 93 percent of reported cyber intrusion incidents could have been avoided with basic cyber-hygiene best practices [14]. Elections are no exception to this statistic and the CEC secretariat staff is a prime target for adversaries. Spear-phishing and social engineering campaigns can start months before an election, and the available tools and methods of cyberattacks are diverse and proven. These tactics can be used by adversaries to influence or blackmail users or exploit stolen information to compromise the integrity of the electoral process. With support from international organizations [15] the CEC developed a cyber-hygiene course to introduce stakeholders to the current risks linked to weak online security and the measures that can be taken to protect them and their workplace from adversaries who seek to harm the democratic process. By the end of 2018, all the CEC secretariat staff had received the training.

Members of DECAs are appointed the between 40 and 65 days prior to the election, they are hence more difficult to identify and target via social engineering techniques. Conversely, DECAs are less trained and more prone to mistakes as they work under a highly stressful environment with various levels of experience and understanding of the tasks they must perform. In consideration of this, members of DECAs also participated in introductory cyber-hygiene trainings. The training of more than 500 members of the electoral commission at the central and district level was a major milestone in improving the security posture of the election staff.

3.5 Incidents reported before the election

During the months leading to the presidential elections, the intensity, frequency, and types of cyber-attacks increased substantially. In January 2019, several government offices including the CEC, were the target of a wide phishing campaign [16] using greeting cards (Christmas cards), shopping invitations, offers for software updates and other malicious phishing material intended to steal passwords and personal information. As these attacks were using some of the same mechanisms used in 2014, they were suspected by the cyber police to originate from the same adversary, namely organizations under the control of Russian special agencies.

On 24 and 25 February, a DDoS attack targeted the CEC infrastructure [17] [18]. While the attack was mitigated, based on the timing of the attack experts interpreted that the assailants were testing the system in preparation of a future attack. During a press conference the next day, the President reported the attack as originating from the Russian Federation.

On 26 February, the SBU announced it had uncovered a plot by a Ukrainian telecom organization contractor (and a resident of Russia) [19]. The Russian contact was interested in collecting data on the networks of strategically important mobile operators, the location of telecommunication nodes and the periods of time necessary to restore them after damage. His objective was to block the communication capacity and disrupt activities of the SVR and disrupt the prior to elections.

While none of these attacks ended up seriously damaging the CEC systems or destabilizing the preparation activities of the cybersecurity agencies, it supported the general trend of growth of attacks both in intensity and sophistication [20], and validated the public's concern for the preparedness of the CEC with regards to its cybersecurity protection.

3.6 Elections days

The first round of the presidential elections witnessed several coordinated attacks on the CEC infrastructure. While the technical details of the attack have not yet been disclosed, it has been established that the systems, particularly the CEC web servers [21], were incessantly probed for vulnerabilities on election day. Several waves of DDoS attacks were also launched against the CEC's result website. As these attacks are frequent in Ukraine, notably mirroring the attacks during the 2014 elections the responsible agencies had the experience and capacity to mitigate them. The efficient collaboration between the SBU, the SSSCIP, and the CEC was vital in providing protection and preventing these attacks from having an impact on the electoral process.

In comparison to the first round, the second round of the presidential election on 21 April was relatively uneventful [22], with no major incident reported by media organizations or government agencies. This should however not be interpreted as an absence of attacks; lessons-learned and preparations made by stakeholders following the first round of the election had mitigated them.

4 Current Threats and Challenges

4.1 Interpretation of the cyberattacks of the presidential elections

As anticipated, the CEC infrastructure and the electoral process, as a whole, were the targets of several cyber-attacks prior to and during the presidential elections.

While the intensity of the attacks was not as high as many experts predicted, they were coordinated with a clear and consistent intent: *find an entryway into the CEC information systems*. There is no indication that the DDoS attacks launched against the CEC website and the RMS were meant to disrupt electoral operations. Instead, the level of intensity would suggest an attempt to cover other attacks to penetrate the systems to obtain control over the resources that could allow the modification of the results.

If the disruption of the electoral process was the main aim of the adversaries, critical infrastructures upon which the election information system heavily relies would likely have been targeted. Cybersecurity agencies did report an increase of attacks against critical infrastructure before the election [23] but they did not materialize during the election itself, which could be interpreted as a strategic decision rather than the lack of capacity for disruption.

Furthermore, a cost-benefit analysis by adversaries offers a probable explanation on **why** there were no indications of attacks attempting to change the presidential election results. The Ukrainian election is grounded in the fundamental use of paper ballots. The manual reconciliation of results and the multiple checks in place are important security features of the Ukrainian electoral system. While the potential to disrupt the electoral process by compromising the preliminary results published on the website exists, it would affect the public perception of the electoral process rather than the actual integrity of the election results⁸ By establishing important and necessary cybersecurity protections, including improved user cyber-hygiene, good information security practices and coordinated response capacity, the CEC has substantially increased the cost of an attack that would bring a limited return of investment.

The cost of any attack seeking to effectively alter the outcome of the election would be high, since it would require that the adversary alters votes from a significant number of polling stations. Moreover, there was also no clear pro-Russian candidate (the likely beneficiary of a “hacked” election) in the 2019 Ukrainian presidential election, making this kind of attack not very attractive. It is important to remember that the cost of this type of attack reduces with the size of the constituency. As such, an attack that could affect the election results of a parliamentary election is hence *cheaper* than it would be for a presidential election. This distinction is important to remember in preparations for the parliamentary elections this autumn.

4.2 Alternative attacks/secondary targets

Targeting ancillary systems (voter registration, candidate registration, etc.) is particularly attractive for adversaries, especially when the election infrastructure

⁸ That being said, the Ukrainian CEC has been in the process of considering moving to an automated results tabulation system, which would increase this risk. Thus, considering such questions and measures prior to any such implementation is considered good practice on their part.

does not offer enough attack surface (the different access points that an adversary can use to penetrate a system to be able to exploit it). The same is the case if there is not sufficient return on investment or the system has been hardened to make any kind of attack costly. Recent and visible attacks on ancillary systems include the attacks targeting the U.S. Democratic Party in 2016 [24] and the Macron leaks in 2017 [25].

Ukraine is not an exception, the website of then presidential candidate Volodymyr Zelensky's campaign was under attack multiple times before the election [26]. As political parties and candidates are increasingly using the internet and digital tools to communicate with their constituencies, they logically become the target of foreign or domestic groups who have an interest in changing the perception of certain candidates or parties. In several countries, national cybersecurity agencies provide instructions and in some cases advice to political agents (NCSC [27] in the UK, ANSSI [28] in France). While it is still unclear what type of support could be provided to political parties and candidates in the preparation of the parliamentary elections, cybersecurity agencies may consider a politically acceptable level of engagement to assist parties and candidates in ensuring the cybersecurity of the electoral process as a whole - even if this assistance is not comparable to the level of effort provided to the CEC. To note, the risk against these ancillary systems is particularly high in Ukraine during the campaigning period, as it comes at the crossroad of cybersecurity and disinformation, with a potentially high impact at a low cost from the part of foreign actors (including such as Russia).

4.3 Remaining Challenges for the CEC

The CEC has greatly improved its cybersecurity protection in the wake of the presidential election. However, some challenges remain and need to be addressed before the 2019 parliamentary elections. Addressing these remaining challenges is extremely important as the political environment and the potential gain that could be obtained from a successful foreign or domestic attack will make the parliamentary election an attractive target.

The CEC has understandably prioritized the implementation of cybersecurity measures to ensure the integrity of the presidential elections. With support from the cybersecurity agencies, the CEC has concentrated efforts on strengthening the accuracy of the results management and transmission systems for the presidential election. Other systems such as the SVR or the CEC's official website have not received the same level of attention. This is understandable, and the CEC already has plans to increase the cybersecurity level of these systems - nonetheless, it is essential these plans are followed through. A carefully planned and long-term attack on the SVR should not be excluded as a possibility, with the objective being the fraudulent modification of the voter list would be sufficient to destabilize or delegitimize election day in some constituencies *without* triggering the CEC's alarms.

The result management verification and data entry are highly decentralized operations run by 225 DECs. DECs rely on staff appointed 45 days before the election, preventing the necessary training on cybersecurity awareness. This represents a major challenge for the parliamentary elections as these DEC officers, particularly those responsible for the IT operations, will be the first line of defense against localized and potentially rewarding hacks against the results system. The CEC and the cybersecurity agencies may mitigate this risk by establishing an advanced cyber-hygiene program for IT officers operating for the temporary district election commissions.

Maintaining a high level of cyber vigilance will be another challenge for the CEC. As it successfully navigated the presidential election, the organization needs to be fully aware that its defense was tested but its capacity to respond to incidents was not. The cybersecurity agencies involved in the incident simulation organized before the parliamentary election expressed the wish to repeat this type of exercise on a regular basis, particularly before an election.

4.4 Information warfare

Foreign interference in the electoral process does not stop at hacking electoral commission infrastructure. There is a wide range of possible attacks that government and all political stakeholders must prepare for, and from recent experiences in Ukraine and abroad, foreign interference often leads to information warfare.

An election in which there is doubt whether the will of voters has been fairly translated into elected mandates leaves a wound that is difficult to heal. In this context, the perception of cyber-attacks against the electoral system might be as important as the attack itself. While the presidential election was widely seen as genuine and democratic by Western observers, disinformation campaigns from Moscow have been distributing a counter-narrative using a wide range of topics and targets [29]. Fighting disinformation is not part of the mandate of the CEC, even when it comes to the electoral process. On this front, the SBU has been proactive in addressing these matters. However, an approach limited to law enforcement and legislation can become a challenge as it needs to balance freedom of expression whilst still lessening the impact disinformation can have on society.

Facebook is currently the largest social network operating in Ukraine with more than 13 million users [30], accounting for more than half of the total internet audience. Facebook's status in Ukraine is a relatively recent phenomenon as it replaced the Russian social network VK or *Vkontakte* [31]. As part of its global objectives, Facebook has launched a global campaign against election interference, particularly the establishment of political advertising transparency policies and the crackdown on Russian-linked group pages or accounts distributing fake information [32].

5 Lessons for the Future

Coordination among cybersecurity and election stakeholders was a key element to the success of the presidential elections and it represents an important milestone in establishing a sustainable program for the protection of the elections in Ukraine. While coordination during the most recent election was efficient and considerate, it calls for higher transparency. There is no international norm of cooperation between Computer Emergency Response Team (CERT) teams and cybersecurity agencies with regards to the protection of electoral infrastructure. As Ukraine is in the front-line of foreign interference in the electoral process, it can take a leading role into establishing publicly available standards defining the rules of engagement of national cybersecurity agencies in response to cyber-incident during an election.

Timely and agile procurement allowed the CEC to patch critical vulnerabilities in its infrastructure. It is most likely that the CEC will need new equipment to secure the next national election cycle (2024). It is very difficult to sufficiently estimate the needs in advance due to the long and rigid rules imposed by government bureaucracy. A revision of the procurement procedures for cybersecurity equipment in general, and cybersecurity equipment for elections specifically, would provide the means to maintain the necessary protection against attempts to threaten the integrity of the electoral process. Additionally, an iterative and planned approach is preferable to an all-or-nothing, last minute solution.

Staff training and cyber awareness at all levels is proved to be effective way in reducing the amount of successful attacks on information systems. Of course, one can only know about the failed attempts, but the fact that Ukraine government agencies experienced several waves of cyber-attacks during the months before the elections, and that none of them provided tools to adversaries to compromise any information system for the presidential election is a good indicator of success. Thanks to users' awareness and cyber-hygiene trainings, targets knew what phishing would look like and thanks to coordination efforts, cyber agencies knew how to monitor and stop campaigns. It should be noted that analysts usually consider that users in Ukraine are very aware of the cyber risk, but do not really know how to behave in order to reduce it. Cyber hygiene programs are hence more adapted than cyber awareness.

The tools of adversaries will very often depend on the return on investment. Hacking an electoral database is a highly spectacular but very often sophisticated operation that represent a substantial cost, particularly if the system is well protected. Hacking the email account or developing a deepfake video of a candidate that have views opposed or not in line with those of adversaries, is cheaper and can provide more immediate benefits.

A nexus of political, technical and regional strategy made the cost of an attack high in comparison to the expected benefit. This will not likely be the case during the upcoming parliamentary elections. In the months prior to the Ukrainian

presidential election, the EU and the rest of the world closely watched the situation unfolding in Ukraine, learning from the CEC and Ukrainian society about the potential risks to the upcoming EU parliamentary elections. Now, Ukrainian officials both within and operating around the CEC should be looking at the EU parliamentary elections in the preparation for its next election in the autumn.

6 Conclusions

In conclusion, the current article has presented a comprehensive analysis of the background and history of the specific Ukrainian case and its significance to the global perspective on current cybersecurity threats to election processes. It has presented a typology of adversary incursion in this space within the former Warsaw Pact countries and sought to show how this impacts strategies likely to be implemented in other parts of the world, now focusing on a composite of DDoS attacks, hybrid warfare, and electoral focus.

The article has also analyzed the recent 2019 presidential election in Ukraine and the cybersecurity threats that were faced, including how they were mitigated. It has also examined the current threats and challenges that the country faces in specific areas such as alternative attacks / secondary targets, information warfare, and remaining challenges that the country's CEC must face.

Finally, it looks at lessons learned from this analysis and its applicability to elections that will take place in the coming period (2018-19), highlighting some of the issues that those concerned with cybersecurity in elections should be focused on. It also presents important issues and questions in need of further investigation.

As such, the authors feel that it makes an important contribution to the field. Both in documenting what has happened in a singular and analytical way. But also, in drawing the threads together to present a common narrative of an increasingly globalized threat to electoral security and integrity. It is hope that it will prompt further research and investigation into this increasingly vital field.

References

1. <https://www.nytimes.com/2017/08/16/world/europe/russia-ukraine-malware-hacking-witness.html>
2. <https://www.reuters.com/article/us-russia-cyber/hackers-accused-of-ties-to-russia-hit-three-east-european-companies-cybersecurity-firm-idUSKCN1MR1BO>
3. <https://www.wired.com/story/russian-hacking-teams-infrastructure/>
4. https://www.washingtonpost.com/world/national-security/how-the-russians-hacked-the-dnc-and-passed-its-emails-to-wikileaks/2018/07/13/af19a828-86c3-11e8-8553-a3ce89036c78_story.html
5. <https://www.politico.eu/article/ukraine-cyber-war-frontline-russia-malware-attacks/>

6. <https://securityintelligence.com/posts/how-ibm-x-force-iris-prepared-for-the-ukraine-election/>
7. <http://ifesukraine.org/cybersecurity-playbook-for-the-elections-in-ukraine/?lang=en>
8. <https://www.enisa.europa.eu/publications/enisa-position-papers-and-opinions/election-cybersecurity-challenges-and-opportunities/>
9. <http://www.rnbo.gov.ua/en/news/3213.html>
10. <https://m.tyzhden.ua/publication/211266>
11. <https://www.usukraine.org/elections-cybersecurity-threats-how-vulnerable-is-ukraine/>
12. <https://www.france24.com/en/20190317-ukraine-ready-to-take-russian-election-hackers>
13. https://ccdcoe.org/uploads/2018/10/Ch06_CyberWarinPerspective_Koval.pdf
14. https://otalliance.org/system/files/files/initiative/documents/ota_cyber_incident_trends_report_jan2018.pdf
15. <https://docs.house.gov/meetings/AP/AP04/20190312/108976/HHRG-116-AP04-Wstate-BanburyA-20190312.pdf>
16. <https://www.reuters.com/article/us-ukraine-cyber-exclusive/exclusive-ukraine-says-it-sees-surge-in-cyber-attacks-targeting-election-idUSKCN1PJ1KX>
17. <https://www.cyberscoop.com/ukraines-president-accuses-russia-launching-cyberattack-election-commission/>
18. <https://www.ukrinform.net/rubric-society/2649609-sbu-blocks-largescale-cyberattack-on-cec-website.html>
19. <https://ssu.gov.ua/en/news/1/category/2/view/5775#.NqSWhPUv.dpbs>
20. <https://www.seattletimes.com/business/technology/ukrainian-official-hacking-intensifies-as-election-nears/>
21. <https://www.kyivpost.com/ukraine-politics/disrupt-and-discredit-russia-still-has-ukrainian-elections-in-sights.html>
22. <https://www.ukrinform.net/rubric-elections/2688206-national-police-no-cyberattacks-on-cec-systems-recorded-during-second-round-of-elections.html>
23. https://news.vice.com/en_us/article/bjqe8m/inside-the-massive-cyber-war-between-russia-and-ukraine
24. https://www.wsj.com/articles/russian-hackers-evolve-to-serve-the-kremlin-1476907214?mod=article_inline&mod=article_inline
25. <https://www.bbc.com/news/blogs-trending-39845105>
26. https://news.vice.com/en_us/article/bjqe8m/inside-the-massive-cyber-war-between-russia-and-ukraine
27. <https://www.ncsc.gov.uk/guidance/guidance-for-political-parties>
28. <https://www.france24.com/en/20170114-france-vulnerable-cyber-attacks-hacking-presidential-elections>
29. <https://ukrainelects.org/live-updates/>
30. <http://www.uadn.net/2019/02/25/with-13-million-users-facebook-is-now-ukraines-leading-social-network-in-ukraine/>
31. <http://euromaidanpress.com/2019/03/24/ukraines-vkontakte-ban-led-to-drop-in-users-but-die-hards-now-more-radicalized/>
32. <https://www.kyivpost.com/ukraine-politics/facebook-rolls-out-new-political-ads-policy-for-ukraine-two-weeks-before-the-vote.html>

GI Elections with POLYAS: a Road to End-to-End Verifiable Elections

Bernhard Beckert³, Achim Brelle⁴, Rüdiger Grimm¹, Nicolas Huber⁵, Michael Kirsten³, Ralf Küsters⁵, Jörn Müller-Quade³, Maximilian Noppel³, Kai Reinhard⁴, Jonas Schwab⁵, Rebecca Schwerdt³, Tomasz Truderung⁴, Melanie Volkamer³, and Cornelia Winter²

¹ University of Koblenz

² Gesellschaft für Informatik e.V. (GI)

³ Karlsruhe Institute of Technology (KIT)

⁴ POLYAS GmbH

⁵ University of Stuttgart

Abstract. Starting from 2019, the annual elections of the GI (German Society for Computer Scientists) will be carried out using a new online voting system developed by POLYAS, aiming at providing high, state-of-the-art security guarantees. We describe the steps that POLYAS plans to take together with the GI and academic partners in order to achieve the level of transparency and trust that is expected from modern online voting. The participation of the academic partners is the key factor to make the verification process both practical and meaningful.

Online voting has been used by the GI — German Society for Computer Scientists (*Gesellschaft für Informatik e.V.*) — for its annual elections since 2004. These elections have been so far carried out using the POLYAS 2.3 voting system. Starting from 2019, the GI elections will be carried out using the new POLYAS online voting system, POLYAS 3.0, based on currently available cryptographic methods.

The current state of the e-voting research offers variety of techniques that address the fundamental security concerns of e-voting: *privacy* and *end-to-end verifiability*. However, to fully utilize the potential of those methods, one has to take into account some practical aspects and challenges, such as usability and constraints imposed by the organisation which carries out the elections. Also, as commonly agreed, the desired degree of confidence and trust cannot be achieved without a high level of transparency and without participation of independent experts. In this paper, we discuss the steps already taken and those planned by POLYAS, the GI, and the academic partners, which are designed to result in an election process which satisfies the pragmatics of the GI elections and, at the same time, provides practical and meaningful security guarantees.

Independent verification tools. POLYAS 3.0 implements standard mechanisms aiming at providing universal verifiability: all the important steps of the tallying process (such as shuffling and decryption) produce appropriate zero-knowledge proofs. It is then, in principle, possible for everyone to check those zero knowledge proofs in order to make sure that the tallying process has been carried out correctly. However, in order to utilize this ability in a practical and meaningful way, we have established a cooperation

including academic partners in order to build independent verification tools. The detailed documentation of the tallying process, based on which verification tools can be implemented, has been provided by POLYAS to the GI and to the academic partners for review. POLYAS has also provided a reference implementation of the verification procedure. Both the documentation and the reference implementation will be made publicly available. As of now, one independent implementation of the verification procedure has already been built by Maximilian Noppel from KIT and implementation of two other verification tools have been started in the group of Prof. Küsters from University of Stuttgart, one of which is being funded by POLYAS.

In the election process, an auditor will have an option to choose (one or more of) the verification tools in order to verify the final election data. We have already used the two existing verification tools to verify the tallying process of a test election carried out in July 2019.

Generation of voters credentials and formal verification of some security goals.

In order to establish a mechanism preventing ballot stuffing (or to provide, so-called, eligibility verifiability), the process of generating voters' private credentials (used to digitally sign the ballots) is carried out in a controlled way by an entity designated by the Election Council (GI). Consequently, POLYAS does not know the voters' private credential upfront. The specification of this process, as well as the source code of the credential generation tool has been provided by POLYAS (with the option to build independent implementations in the future). The process of credential generation (carried out by GI using the provided tool) has been part of the mentioned above test election. We plan to make both the specification of this process and the source code of the corresponding tool publicly available.

The central security goals of the credential generation tool is that the randomly generated plaintext passwords are only saved in an encrypted file designated for the trusted distribution facility and they do not leak in any other way (in particular, they should not make it to the file designated for POLYAS which should only contain the corresponding derived public credentials). We plan to use program verification methods to formally prove this property on the implementation level. In this non-trivial task, which has already been started by Michael Kirsten from the group of Prof. Bernhard Beckert (KIT), the KeY tool will be used in combination with techniques from simulation-based security.

Individual verifiability. The known solutions for individual verifiability involve several trade-offs, including the usability aspect: the voters should understand the process and be able to carry out the prescribed steps. Moreover, the election council must establish well define procedures for handling voters' complaints. We plan to address this security requirements in the second step, that is in the GI Elections 2020. The solution which POLYAS offers is based on the optional use of a second device (such as a mobile phone) by the voter. The details of this process are being still discussed.

The **goal of our cooperation between POLYAS and the academic partners** described above is to make the election process offered by POLYAS transparent and auditable in a practical and meaningful way. We believe that opening this process is the best way to build up trust and to improve the offered e-voting solutions.

Pakistan’s Internet Voting Experiment

Hina Binte Haq¹, Ronan McDermott², and Syed Taha Ali¹

¹ School of Electrical Engineering and Computer Sciences (SEECS), National University of Sciences and Technology (NUST), Islamabad, Pakistan.

{hhaq.dphd18seecs, taha.ali}@seecs.edu.pk

² MCDIS ronan@mcdis.com

Abstract. Pakistan recently conducted small-scale trials of a remote Internet voting system for overseas citizens. In this contribution, we report on the experience: we document the unique combination of socio-political, legal, and institutional factors motivating this exercise. We describe the system and its reported vulnerabilities, and we also highlight new issues pertaining to materiality. If this system is deployed in the next general elections —as seems likely —this development would constitute the largest enfranchised diaspora in the world. Our goal in this paper, therefore, is to provide comprehensive insight into Pakistan’s experiment with Internet voting, emphasize outstanding challenges, and identify directions for future research.

Keywords: Internet Voting · Overseas Voters · Pakistan.

1 Introduction

Pakistan recently piloted a remote Internet voting system for overseas citizens. This system, *i-Voting*³, was indigenously developed and originally scheduled for full-scale deployment in the General Elections of July, 2018. However, these plans were deferred after a third-party technical audit of the system uncovered numerous vulnerabilities and security concerns. *i-Voting* was deployed shortly after on a trial basis: in bye-elections, first in October, 2018, spanning 35 constituencies, and next in December, 2018 in 1 constituency. Some 7,538 votes were cast (7,461 in October and 77 in December) using this system and these were declared binding and incorporated into the final results.

It is widely expected that this pilot is a prelude to full-scale deployment in the General Elections of 2023.⁴ Since Pakistan currently has over 8 million overseas citizens [1], this may well be the largest deployment of Internet voting in the world. It is therefore essential to document and study this experiment.

In this paper, we make the following contributions:

1. We report on the deployment: we document the public debate on Internet voting and the legislative and political process to facilitate it. We describe the *i-Voting* system and we report on the pilot exercise.

³ also referred to as *iVoting*, *iVOTE*, *IVoting*

⁴ Currently, the law restricts use of such systems to bye-elections.

2. We describe the various risks involved in this modality of voting and summarize key findings of the technical audit. We examine the materiality, and therefore the potential political significance of overseas voting.
3. We highlight the unorthodox combination of unique political and social factors in Pakistan that have resulted in this exercise and we discuss various particular challenges that may arise as a result.

This paper holds relevance for election stakeholders including governments, political parties, election administrators, political scientists, researchers, and technologists. Pakistan’s Internet voting experience may also prove instructive for other countries, particularly in the developing world, where governments are severely limited in terms of financial resources, technical expertise, and infrastructure to undertake such critical large-scale projects.

2 Background

Organization of Government Pakistan has a parliamentary form of government with bicameral legislature, comprising a Senate (upper house) with 104 members and a National Assembly (lower house) with 342 members. Each of the four large provinces have a unicameral legislature, consisting of a Provincial Assembly.⁵ The electoral system is the first-past-the-post system under universal adult suffrage. Members of the National Assembly and Provincial Assembly are elected by representation in electoral districts (referred to as seats or *constituencies*). The number of seats in each administrative division is listed in Table 2. General elections are conducted every five years and are overseen by the Election Commission of Pakistan (ECP), which is an independent and autonomous body as defined in the Constitution of Pakistan.

Body	Total Seats	Federal Capital Islamabad	Baluch-istan	Federally Administered Tribal Areas	Khyber Pakhtun-khwa	Punjab	Sindh
National Assembly	272	3	16	12	39	141	61
Provincial Assembly	577	-	51	-	99	297	130

Table 1. National and Provincial Assembly Seats [2]

Overseas Pakistanis and the Right to Vote Pakistan, has over 8 million overseas citizens [1] which comprises the sixth largest diaspora in the world [3]. Overseas citizens are actively engaged in the socioeconomic well-being of the country and every year send home remittances worth approximately US\$19 billion, which accounts for around 5% of Pakistan’s GDP [4].

Article 17 of the Constitution of Pakistan grants all adult citizens the fundamental right to vote [5]. This article has generally been interpreted to acknowledge that this right extends to all Pakistani citizens, irrespective of place

⁵ Federally Administered Tribal Areas (FATA) and Federal Capital Islamabad are administrative divisions in addition to the four provinces, included in the contested elections and comprise National Assembly seats only.

of residence. Overseas Pakistanis have raised calls for enfranchisement and facilitation of their voting rights since the first general elections of 1970 [6].

The earliest constitutional petition filed to facilitate overseas voters was in 1993 by a British-Pakistani student and the Supreme Court of Pakistan referred it to the government and the ECP for consideration [7]. After a hiatus of almost two decades, more petitions followed in quick succession: in 2011, Dr. Arif Alvi, Secretary General of Pakistan Tehreek-e-Insaf (PTI), a popular political party petitioned the Supreme court in this regard; in 2014, the Islamabad High Court was likewise petitioned by a concerned overseas citizen, and in June 2015, by the Chairman of PTI, Mr. Imran Khan.

The judiciary, while addressing this grievance, has upheld this fundamental right of overseas citizens on multiple occasions [8] [9] [10], ruling that this right cannot be denied on technical grounds, and it has repeatedly directed the ECP to make the necessary logistical arrangements. We discuss these attempts next.

Efforts by the Election Commission and Parliament In 2012, in one of the earliest statements on the subject, the ECP dismissed the possibility of overseas citizens' participation in the General Elections of 2013, citing logistics and budgeting issues [11]. However, by 2015, the ECP had established a Directorate for Overseas Voting in its Secretariat, which conducted mock overseas voting exercises using postal ballots and voting via telephone [12].

These trials were unsuccessful. The reasons were clarified in a study commissioned by the ECP: *"We find that any remote voting solution using currently available technology whether postal, internet, telephone, or proxy will lack the necessary electoral integrity checks to preserve the credibility of an election result."* Commenting on the feasibility of other modalities, the report stated: *"...given the size and dispersal of the Pakistani diaspora, coupled with the limited official resources available in-country and abroad, any significant in-person voting operation would be expensive and logistically challenging"* [13].

In July of 2014, the Parliamentary Committee for Electoral Reforms was constituted with Finance Minister Ishaq Dar as chair. A sub-committee was formed in January, 2016 by MP Dr. Arif Alvi (one of the petitioners for overseas voting mentioned earlier and Secretary General of PTI) to devise a mechanism for overseas voting [14] and in March, 2017, it proposed remote Internet voting as a potential solution. Consequently, the committee authored the Elections Act of 2017, which authorized the ECP *"to conduct pilot projects for voting by Overseas Pakistanis in bye-elections"* [7].

The ECP subsequently requested the National Database Registration Authority (NADRA) to build a system. NADRA is an independent and autonomous agency working under the Ministry of Interior and tasked with managing government databases and issuing national identity cards to citizens. However the system did not materialize: in June 2017, ECP again contacted NADRA, but NADRA expressed its regrets at not having a solution available [12].

The Supreme Court Intervenes Interest in overseas voting peaked again in the six months leading up to the General Elections of July, 2018. The Supreme Court of Pakistan consolidated 16 similar constitutional petitions and resumed hearings on the issue [15]. It sought reports from the ECP and NADRA over non-compliance of Section 94 of the Elections Act (regarding pilot projects for overseas voters) [16]. In an attempt to break the deadlock, the Supreme court directed NADRA to develop an Internet voting system. NADRA informed the court that it would require 4 months to build a system and would cost Rs.150 million (approximately US\$ 1.36 million) [17]. However, the Court ordered NADRA to present it in 10 weeks [18].

The new system, i-Voting, was unveiled on April 12, 2018 at a public session convened by the Supreme Court [19]. The audience included members of various political parties, IT experts from Pakistani universities, concerned citizens, and members of the media. It was here that IT experts aired serious security concerns regarding this system and pointed out that similar systems had been demonstrably attacked and were being phased out in developed countries. The Supreme Court concluded the session by forming an Internet Voting Task Force (IVTF) to audit the system and assess its suitability for deployment in the forthcoming general elections of July, 2018.

3 The i-Voting System

Here we describe NADRA's i-Voting system and summarize the IVTF findings.

System Architecture The i-Voting system conforms to the traditional design of Internet voting solutions, where a centralized database is used to store and tabulate votes, which voters access using a Web portal.

More specifically, as depicted in Fig. 1, a central datacenter hosts the overseas votes database and an application and email Server. These servers interface with a webserver hosting the i-Voting Web portal (Fig. 2) and NADRA databases (for verification of voter information). The ECP can monitor the system using an administrative portal. Load balancing and backup arrangements are deployed as well as standard security solutions including firewalls, intrusion detection mechanisms, and mitigation of distributed denial of service (DDoS) attacks.

Voter Registration The registration process is depicted in Fig. 3 [20]. To enrol, a user must possess his/her passport, a National Identity Card for Overseas Pakistanis (NICOP), and a valid e-mail address. The ECP announces a public Registration Phase during which prospective voters enter their basic details into the system. A confirmation email with a PIN is then sent to the voter. To confirm the account, the user enters the PIN and solves a CAPTCHA.

Now the user logs in to the system and provides further details of his NICOP and passport. He also answers two randomly chosen questions pertaining to his identity after which he is successfully registered. The system allows a maximum of 3 answer attempts, failing which the NICOP number is restricted.

Vote Casting and Preparation of Results Prior to polling day, each registered voter is emailed containing a unique passcode (which acts as a one-time password), enabling him to log on to his i-Voting account and cast a vote for his respective National Assembly and/or Provincial Assembly seat (Fig. 3).

When polling concludes, the ECP tabulates the votes via the Reporting Portal and dispatches the tally to concerned officials for consolidation of results.

Internet Voting Task Force The Internet Voting Task Force (IVTF) was given a time window of 3 weeks to assess the security of the i-Voting system. The team comprised of IT and security specialists and academic researchers [21]. They conducted a high level security analysis of the system, examined the code, and mounted some typical attacks. The results were written up in a report and submitted to the Supreme Court. Their key findings are as follows:

1. i-Voting does not provide ballot secrecy, a fundamental right defined in the Constitution of Pakistan. This further opens up the possibility of vote buying and coercion of overseas voters.
2. A key security vulnerability allowed overseas voters to bypass their native constituencies and cast votes for any two seats of their choice in the country.
3. The IVTF successfully launched impersonation attacks, enabling them to send fake emails purportedly from the ECP to direct voters to fake websites.
4. i-Voting avails the services of a leading DDoS mitigation solution, a measure which researchers have recently demonstrated can potentially compromise ballot secrecy and election integrity [22].
5. The system employs certain third-party security components (such as text-based CAPTCHAs) which are obsolete and demonstrably insecure.

The IVTF also raised other critical non-security concerns: i-Voting lacked verifiability, fail-safe, or redundancy mechanisms. There were no security policies or procedural controls defined to protect critical security processes from insider attacks. No usability studies or trials had been conducted for the system. Furthermore, the system was built in an ad-hoc manner with key documentation missing. For instance, there was no documented Solution Requirements Specification (SRS) or documentation pertaining to key operational processes (such as administration, hosting, responsibility of critical components), which limited assessment for certain important security attacks.

The IVTF therefore strongly argued against the deployment of i-Voting in the upcoming General Elections of 2018. Their report stated that this would be “*a hasty step with grave consequences*” The report also emphasized that “*many*

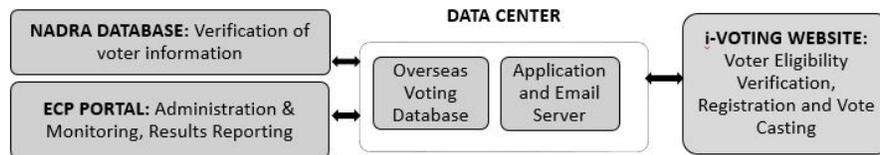


Fig. 1. i-Voting: System Architecture

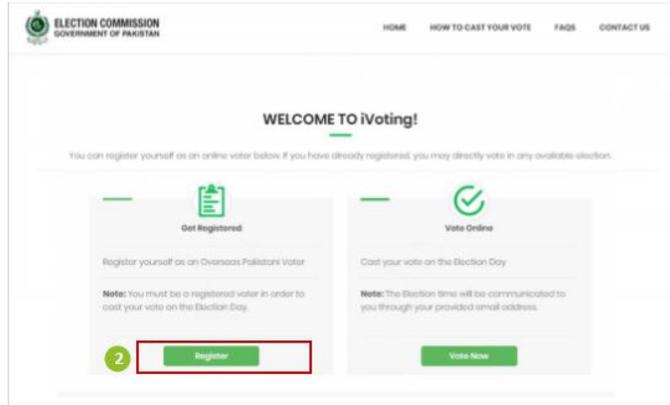


Fig. 2. iVoting: Interface

of these security vulnerabilities are not specific to iVOTE [sic] but are inherent to this particular model of Internet voting systems” [23].

The report also made various recommendations to facilitate overseas voters. We discuss these in Sec. 6 and 7.

4 Deployment

The Supreme Court of Pakistan revisited the matter of Internet voting after the general elections in August, 2018, and ruled: “Based on these representations we prima facie find the mechanism of I-Voting (sic) to be safe, reliable and effective for being utilized in a pilot project. We are sanguine that the aforesaid proposed rules shall be incorporated in the Election Rules, 2017 to enable overseas Pakistanis to exercise their right of vote in the forthcoming bye-elections.” The court further stipulated that votes cast using i-Voting not be added to the final tally until the ECP is satisfied with regards to their “technical efficacy, secrecy, and security”. In case of any dispute the ECP was authorized to exclude these overseas votes from the official results [10].

The ECP consequently amended the Election Rules to accommodate the requirements of Internet voting. NADRA implemented certain technical recom-

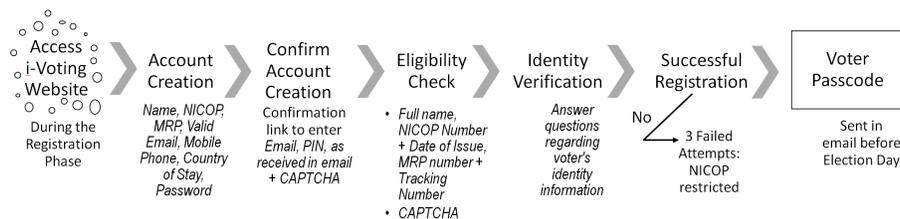


Fig. 3. Voter Registration

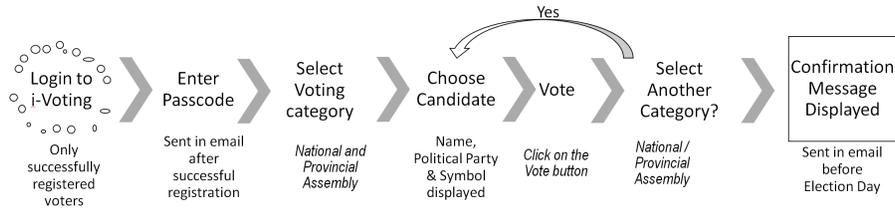


Fig. 4. Vote Casting

mentations of the IVTF⁶ and trained ECP officials to administer the system. The ECP launched a media campaign for voter awareness and published detailed guides and video tutorials for the i-Voting system. A dedicated support center was also set up to provide telephone and email assistance [12].

First Pilot (Bye-Elections - 14 October, 2018) Bye-elections were held for 35 constituencies (11 National Assembly and 24 Provincial Assembly seats). The total overseas Pakistanis eligible to participate in these polls numbered a significant 631,909. However, out of these only 7,419 citizens (1.17%) actually registered to vote using the new system. On the day of the elections, a total of 6,146 voters of these citizens cast their votes [12].

ECP later reported that on the day of the polls the system successfully withstood 7,476 DDoS attempts.⁷ The top 5 countries by voter count were the United Arab Emirates (1,654), Saudi Arabia (1,451), the United Kingdom (752), Canada (328), and the United States (298). The pilot project cost approximately Rs. 95 million (0.67 million USD approximately) [12].

The trial was smooth and uneventful. In its own report, the ECP attributed the low turnout to the short time frame within which the system was deployed and advertised. The ECP also cited key issues which echoed the concerns of the Internet Voting Task Force (discussed in Sec. 3), in that the system violates ballot secrecy, enables voter coercion, lacks auditability, and may be vulnerable to state-level cyberattacks.

Second Pilot (Bye-Elections - 13 December, 2018) As many as 4,667 overseas Pakistanis from more than ten countries were eligible to vote for one Provincial Assembly seat [24]. However, only 77 overseas Pakistanis registered to vote. The ECP has not released any further details about this trial [25].

5 Materiality - Are Overseas Votes Decisive?

In this section we undertake a basic post-election analysis to examine the potential impact of overseas votes on final results.

⁶ No details have been published on what specific changes were made.

⁷ No details of these attacks have been released to the public.

The leading Pakistani citizen observation group, Free and Fair Election Network (FAFEN) has conducted an analysis of 2018 General Election results [26] and has determined that, in a significant number of constituencies, the Margin of Victory (MoV) is less than the number of invalid votes. A similar analysis, but comparing Margin of Victory with number of Eligible Overseas Voters would be useful to highlight the materiality⁸. There is no publicly available voter registration or population data which shows how many registered voters are actually Overseas Pakistanis for all constituencies. The only exception to this is for the constituencies where by elections were conducted using i-Voting in October and December 2018. The average percentage of eligible overseas voters, as a percentage of total voters in the October 2018 By-elections is 6.88% (Table 2). With this assumption, we estimate the number for eligible overseas voters for individual constituencies under scrutiny.

We then calculate an estimated Overseas Pakistani Voters value for each contested constituency in the October 2018 By elections that was also contested in the July 2018 General Elections. We now compare these Estimated Overseas Voters (EOV) value with the margin of victory and flag where the MoV is less. We do this both for the October 2018 By-elections and the July 2018 General Elections. We may describe the number of Overseas Pakistani Voters in these cases as material to the outcome of the election⁹.

As Table 3 shows, in ten of twenty-seven by-election races, Overseas Pakistani Voters had the potential to be material. This grows to thirteen of twenty-seven for General Election races. In five races, both Bye-Election and General Election saw Margin of Victory less than estimated Overseas Pakistani Voters. It is, we believe, reasonable to assume that, in competitive future General Elections at least one in five races may be decided by votes cast by overseas Pakistani voters. This places the integrity of and trust in any internet voting solutions deployed by the Elections Commission of Pakistan into very sharp focus. This is positive in the sense that overseas Pakistanis can feel their votes count. At the same time it necessitates election integrity checks so that this right is not misused.

6 Discussion

In this section we further examine outstanding issues arising from this experiment and we attempt to contextualize these by examining the unique political and institutional factors that motivated these pilot projects.

⁸ Materiality in this context refers to the theoretical scenario where all possible overseas votes are cast, and all are cast for the second place or losing candidate, the outcome of the election might have been different/might be different.

⁹ The public domain sources for Table 2 and 3 are no longer available on ECP Website. The documents will be made available at a URL which will be cited later (to retain anonymity).

Description	Number
Total Registered Voters in bye-election	9,185,705 []
Total Eligible NICOP	734,777 []
Non-Machine Readable Passports (14%)	102,868 []
Estimated Total Overseas Pakistani Voters	631,909 []
Eligible Overseas Pakistani Voters as % of Voters Registered	6.88% []

Table 2. Calculating Estimated Percentage of Overseas Pakistani Voters



Fig. 5. Twitter users Posting Screenshot of Vote

6.1 Ballot Secrecy and Voter Coercion

In delivering one right (the right of overseas Pakistanis to vote), the solution risks undermining another (the right to secrecy). The i-Voting system does not comply with Article 226 of the Constitution of Pakistan [27] and the Elections Act 2017, Section 81¹⁰ [28], that impose ballot secrecy. Being a remote voting modality, there is no mechanism to prevent an individual from revealing their vote to others. Similarly, certain event logging software (specially on a shared/public device) can secretly capture the choice of a voter.

Electoral offences were committed by voters unintentionally, by posting screenshots on social media, as shown in Fig 5. The tweeters seem to be unaware their actions are electoral offences, and being outside the jurisdiction of Pakistan it is unclear, how such offenders can be brought to justice¹¹.

In addition, to the lack of secrecy for the voter at the client end, the low levels of participation in pilots also mean that, in some cases, typically PA seats (PB-35, PP-165, PP-292 from October 2018 bye-elections [12]), the voter’s choice is revealed. The usual solution to this problem (mixing votes from multiple ballot

¹⁰ The exceptions in Section 81 are not to the secrecy requirement, rather to requirement of casting a vote by inserting paper ballot into a ballot box.

¹¹ Section 178 of the Elections Act 2017 elaborates the offences relating to ballot secrecy.

Constituency	Total Registered Voters	Estimated Overseas Voters (EOV)	General Elections July, 2018		Bye Elections October, 2018		Both Elections
			MoV	MoV <EOV	MoV	MoV <EOV	MoV <EOV
PB-40	76,173	5,240	13,345	-	9,141	-	-
NA-53	313,141	21,541	48,763	-	18,630	-	-
NA-35	582,785	40,091	7,001	Yes	23,455	-	-
PK-3	146,180	10,056	5,550	Yes	1,163	-	-
PK-7	155,719	10,712	5,825	Yes	334	-	-
PK-44	202,601	13,937	10,857	Yes	1,630	-	-
PK-53	153,352	10,549	6,729	Yes	61	-	-
PK-61	139,517	9,597	4,593	Yes	5,247	-	-
PK-64	160,728	11,056	18,579	-	13,215	-	-
PK-97	155,032	10,665	16,461	-	10,172	-	-
NA-56	640,133	44,036	64,490	-	41,593	Yes	-
NA-63	371,713	25,571	35,979	-	26,292	-	-
NA-65	553,289	38,062	51,963	-	68,591	-	-
NA-69	469,177	32,275	73,172	-	50,803	-	-
NA-124	535,172	36,815	65,287	-	47,533	-	-
NA-131	365,677	25,155	756	Yes	10,031	Yes	Yes
PP-3	224,755	15,461	37,008	-	227	Yes	-
PP-27	312,370	21,488	1,766	Yes	656	Yes	Yes
PP-118	222,190	15,285	548	Yes	5,189	Yes	Yes
PP-164	137,906	9,486	20,870	-	7,561	Yes	-
PP-165	132,077	9,085	20,372	-	5,742	Yes	-
PP-201	232,120	15,968	17,297	-	7,024	Yes	-
PP-222	196,858	13,542	11,446	Yes	6,083	Yes	Yes
PP-261	187,510	12,899	9,371	Yes	14,261	-	-
PP-272	167,467	11,520	5,390	Yes	8,899	Yes	Yes
PP-292	141,297	9,720	253	Yes	10,692	-	-
NA-243	402,731	27,704	67,291	-	21,601	-	-

Table 3. Materiality of Overseas Voters in Selected Constituencies, General and Bye-Elections 2018

boxes or polling stations) is not available in the i-Voting context or could only be implemented at the cost of further erosion of already minimal transparency.

As the IVTF report points out, some jurisdictions [29] allow a voter to waive their right to secrecy. This is not a solution to the problem, as any voter (or party or candidate) to assert their constitutional and legal rights to secrecy for the system to be challenged.

Almost half of the diaspora, over 4 million Pakistanis reside in the Middle East, and about a quarter (over 1.5 million) reside in Europe [30]. A bulk of the diaspora specifically in the Middle East are labourers. The ECP itself recognizes the risk of vote buying and coercion when it speaks of the "kafeel"¹² abusing custody of passports [12]. The ILO¹³ describes this system as placing migrant workers in "a position of vulnerability and have very little leverage to negotiate with employers, given the significant power imbalance embedded within the em-

¹² Sponsor for a migrant worker

¹³ International Labour Organization

ployment relationship. Common grievances expressed by migrant workers include restrictions on free movement, confiscation of passports, delayed or non-payment of salaries, long working hours, untreated medical needs, and violence all conditions that can give rise to situations of forced labour and human trafficking” [31]. It is reasonable to assume that anyone who will treat migrant workers in this manner will not hesitate to exploit their votes for political or financial benefit.

Migrant workers are bound to face difficulty to independently use the i-Voting System. This could pave the way for coercion, vote buying and compromise secrecy if vote casting is aided by a computer literate party. Thus, usability tests need to be conducted to receive direct input from real users. It might be argued that low usability, was a primary reason of the low registration turnout, where only 1.17 % [12] of the total eligible overseas voters successfully registered with the i-Voting system. Further, there seem to be no special accessibility features incorporated to address the needs of voters with disabilities.

6.2 Voter Authentication

The process of registration on the i-Voting platform [32] is entirely out of ECP’s control, relying as it does on a verification method conducted and adjudicated by a computer programme. Potential overseas voters are quizzed with questions, whose answers are considered “*secret*”. Common sense dictates that, despite the familial/personal nature of these questions, the answers will not be known exclusively to the voter. As a consequence, ECP cannot guarantee that the voter registered via the online platform [32] is indeed the eligible citizen, or an impostor. Other key mechanisms to protect the integrity of the electoral rolls - the public display of and claims/objections on the draft electoral rolls - are omitted from the online process. Political parties, observers, and voters themselves, are not given the access to these electoral rolls to allow for the scrutiny that would contribute to stakeholder confidence in the electoral rolls.

Furthermore, the mechanism within i-Voting to “*lock*” an identity following repeated incorrect answers or CAPTCHA verification may be used for voter suppression - a sort of denial-of-service attack, albeit on a vote-by-vote basis. Voters may not know the answers to all the questions they might be asked (where, for example, an 18 year old was registered and a parent provided the information). Corrupt or partisan Presiding Officers could merely strike out the names of legitimate voters saying that they had registered online for i-Voting.

6.3 Election Integrity and Dispute Resolution

On election day in polling stations across Pakistan, a long list of integrity mechanisms are in place, arising from the Constitution, the Elections Act 2017 and the Election Rules 2017, as amended¹⁴. In the i-Voting system these fourteen

¹⁴ The election is conducted in full view of polling staff, party/candidate agents and observers, who first-hand witness the integrity checks in place: verify ballot boxes are initially empty, identify voters on arrival, ink their fingers (to prevent multiple

separate mechanisms are missing with consequences for electoral integrity. The exclusion of party/candidate agents and citizen observers from the i-Voting process is compounded by the inherent absence of any verifiability mechanism or possibility to audit the i-Voting system - by design - "*In order to ensure that Voting is kept secret, all data was encrypted and no audit trail of voting was kept by the system*" [12].

ECP may exclude based on its "*opinion*" as to whether the "*technical efficacy, secrecy and security of the voting has not been maintained*" [12]. It is not known how ECP informs that opinion, or whether it has the required access to the i-Voting system. Given the 2018 recourse to establishing the IVTF, it seems likely that ECP lacks the technical capacity to properly arrive at an informed opinion. Given the likely materiality of votes cast by overseas Pakistanis in a significant proportion of contests, we may expect many electoral disputes to centre around the integrity of i-Voting system.

Specifically, in a developing country like Pakistan, where the democratic process is at an inflection point, and the mechanisms to investigate and resolve electoral disputes, are still very fragile, electoral improprieties or even the impression of such can potentially lead to political deadlock and turmoil. An indication of this is the PILDAT (Pakistan Institute of Legislative Development and Transparency) report on the perception of pre-poll fairness which notes that "*the internet-based OP voting may also be a major instrument of rigging in 2018 General Election*" [33]. Diverse stakeholders including the ECP itself [12], and prominent mainstream political parties expressed similar reservations [34] [35]. Lack of auditability features means there is no evidence if results are challenged through an election petition. These, coupled with questions over the capacity and willingness of the judiciary, raises concerns about the resolution of electoral disputes around internet voting.

6.4 Threat Model

Concerns have also been raised regarding the threat model on which i-Voting is based. The recent controversy of foreign interference in US elections hints that a developing country like Pakistan may also be at risk. ECP's report on the pilot deployment highlights this concern: "*[adversaries] did not materially interfere merely to put us off track. When the system is finalized and put into practice in the next elections, shall we be able to counter/control cyber attacks[?]*" [12].

Furthermore, whereas the high number of DDoS attacks (around 7476 on bye-elections day), posed an outage threat, the use of a DDoS mitigation service, as demonstrated recently by Culnane et al. [22], introduces a new attack vector. The mitigation service is in a position to decrypt incoming traffic, thereby able to compromise ballot secrecy and potentially even alter the content. The IVTF

voting), vote casting in secrecy behind a screen, placing the ballot paper into the transparent ballot box, the Presiding officer conducting the count and disseminating the results form to all stakeholders, and packaging all ballot papers (valid, invalid, challenged, spoiled) separately in tamper-evident envelopes.

audit highlighted this concern in their report and pointed out that the servers employed by the DDoS mitigation service were all based overseas and beyond control of Pakistani authorities.

6.5 The Curious Case of Pakistan

We see the Supreme Court at the forefront, driving the institutions to deliver a voting solution to overseas Pakistanis. Here we try to make sense of the unique predicament and examine the various factors that led to this situation. The judiciary has time and again reiterated the ECP to roll out a voting mechanism for overseas Pakistanis, but to no avail. A concrete step in this direction was long overdue and it had to take the Supreme Court to push it through, given the institutional inertia within the government.

The Supreme Court of Pakistan frequently takes Government ministries and other public bodies to task for not fulfilling their obligations [36]. Whether through judicial activism (using *suo moto* powers) or responding to petitions from interested parties, it often gets involved in the technical specifics of cases. Its jurisdiction “*is not limited to mere procedural technicalities as it enjoys certain inherent powers to do complete justice in any case*” [37]. A vivid example of the Supreme Court’s ambition beyond procedural technicalities is the fund established in July 2018 to raise money to build dams. This fund currently exceeds ten billion Pakistani rupees (approximately 71 million US \$) [38].

Cognizant of this deviation, the Honourable Chief Justice of Pakistan, inquired whether it was the job of the Supreme Court to give the right to vote to overseas Pakistanis? [16] In the matter of how best to enfranchise overseas Pakistanis the Supreme Court initially directed the elections management body to develop an internet voting system and then later mandated the use of this system in binding political bye-elections. In doing so, the Supreme Court dismissed unambiguous and dire warnings from the IVTF about the hazards of the proposed system as mere “*technical and security apprehensions*”, and that the report was “*generally positive and encouraging*” [10]. While the IVTF report clearly says “*Hopefully, this discussion thus far demonstrates to the reader why internet voting is recognized by security experts to be a controversial and risky undertaking*”, and it concludes by asserting “*We would, therefore, urge all stakeholders to exercise extreme caution in approaching the question of internet voting*” [23]. This disconnect has received media recognition [39].

The ECP itself was very reluctant to adopt this modality of overseas voting. Recently, when the matter was taken up in the Senate, in May 2019, Senator Javed Abbasi recollected that “*the ECP had convinced political parties that the system should not be introduced in Pakistan, but could not convince the Supreme Court*”, at which an ECP representative expressed his dismay that while the ECP tried to dissuade the Supreme Court, no political party supported the ECP in Supreme Court [40]. The absence of a broad political consensus on the use of i-Voting to enfranchise the diaspora does not bode well for the future.

Neither the Elections Commission of Pakistan, nor the developers of the i-Voting system challenged the Supreme Court’s interpretation of the IVTF find-

ings or recommendations. Since there is no higher court than the Supreme Court, no appeal is possible. If a future election is decided on votes cast by overseas Pakistanis, via the i-Voting system, and the result is challenged —which seems highly likely, given both the deficiencies and materiality described earlier in this paper —it will be interesting to see if the electoral dispute resolution process ends up in the same court.

7 Way Forward and Conclusion

Pakistan’s experiment for the October 2018 by elections was the largest deployment of Internet voting in a binding election, anywhere in the world. The recommendations of both the IVTF report, and ECP’s own report on the October 2018 pilot exercise are comprehensive and we endorse these. Going beyond these, and comparing the Pakistani experience with other countries who are further along the internet voting pathway, we would highlight two vital priorities. First, transparency: ECP and NADRA succeeded in delivering a working prototype system in the short time available, but the details of the process were, and remain, opaque. Stakeholder acceptance cannot be assured in future without meaningful transparency and greater consultation, having regard to voter secrecy. Second, capacity building across all stakeholders, starting with, and led by ECP (such as establishing a dedicated R&D cell within the ECP), to deliver competent national ownership and informed policymaking. It seems likely that escalation from pilots in bye-elections to full-scale use of internet voting for the enormous Pakistani diaspora will happen in 2023. The issues highlighted in this paper should receive urgent attention by all Pakistani stakeholders.

References

1. PRIO. Pakistan as a return migrant destination. <https://opf.org.pk/media/1410/pakistan-as-a-return-migration-destination.pdf>, 2015.
2. ECP. Final List of Constituencies (Final Delimitation 2018). <https://www.ecp.gov.pk/frmGenericPage.aspx?PageID=3100>.
3. Haq, R. India tops labour export, Pakistan ranks 6th. <https://en.dailypakistan.com.pk/opinion/blog/international-migrants-day-india-tops-labor-export-pakistan-ranks-6th/>, December 2016.
4. Rizvi, M. Pakistan remittances may hit \$22 billion in 2018-19. [https://www.khaleejtimes.com/business/economy/Pakistan-remittances-may-hit-\\\$22-billion-in-2018-19-](https://www.khaleejtimes.com/business/economy/Pakistan-remittances-may-hit-\$22-billion-in-2018-19-), October 2018.
5. Article: 17 Freedom of association — The Constitution of Pakistan, 1973. <https://pakistanconstitutionlaw.com/article-17-freedom-of-association/>, 1973.
6. Mehboob, A. B. Voting from abroad. <https://www.dawn.com/news/1335670>, May 2017.
7. Mehboob, A. B. Arranging Voting by Overseas Pakistanis. <https://pildat.org/blog/arranging-voting-by-overseas-pakistanis-when-and-whose-job-is-it-2>, April 2018.

8. Const.p.39and90of2011-dt-3-5-2013. http://www.supremecourt.gov.pk/web/user_files/File/Const.P.39and90of2011-dt-3-5-2013.pdf.
9. Rao, S. Ihc refers matter to ECP. <https://www.pakistanpressfoundation.org/ihc-refers-matter-ecp/>, December 2014.
10. Const.p..74.2015. http://www.supremecourt.gov.pk/web/user_files/File/ConstP._74_2015.pdf.
11. Butt, Q. Election preparations: Overseas Pakistanis unlikely to vote in polls. <https://tribune.com.pk/story/435109/election-preparations-overseas-pakistanis-unlikely-to-vote-in-polls/>, Sept 2012.
12. Report on I/-Voting Pilot Test. <https://ecp.gov.pk/documents/ivotingreport.pdf>, October 2018.
13. Mehboob, A. B. Voting from abroad. <https://www.dawn.com/news/1335670>, May 2017.
14. Ali, F. Apathy of electoral reform. <https://www.dawn.com/news/1258537>, May 2016.
15. Bhatti, H. SC accepts 16 petitions regarding overseas Pakistanis. <https://www.dawn.com/news/1380952>, January 2018.
16. Correspondent. CJP asks whether parliament will give voting rights to overseas Pakistanis. <http://dunyanews.tv/en/Pakistan/423769>, January 2018.
17. Monitoring Report. NADRA to develop Internet voting system for expats. <https://www.pakistantoday.com.pk/2018/01/25/nadra-to-develop-internet-voting-system-for-expats-report/>, January 2018.
18. Malik, H. Overseas Pakistanis suffrage: SC orders NADRA to develop voting software in 10-weeks. <https://tinyurl.com/y24gfhzh>, January 2018.
19. Supreme Court Press Release. SC hears constitution petitions regarding right of vote to overseas Pakistanis in general elections. http://www.supremecourt.gov.pk/web/user_files/File/Press_Release_47_2018.pdf.
20. iVoting- Registration Closed. <https://www.overseasvoting.gov.pk/i-voting/en-guide.pdf>.
21. TOR and ECP order. <https://ecp.gov.pk/TOR\%20and\%20ECP\%20order.pdf>, April 2018.
22. Culnane C., Eldridge M., Essex A., and Teague, V. Trust implications of DDoS protection in online elections. In *International Joint Conference on Electronic Voting*, pages 127–145. Springer, 2017.
23. IVTF Report Executive Version. <https://www.ecp.gov.pk/ivoting/IVTF\%20Report\%20Executive\%20Version\%201.\%20Final.pdf>, May 2018.
24. FAFEN. By-Election Report - PP-168 Lahore. <http://fafen.org/wp-content/uploads/2018/12/By-Election-Report-PP-168-Lahore-XXV.pdf>, Dec 2018.
25. Bilal, R. PTI defeats PML-N by close margin in Lahore's PP-168 by-poll. <https://www.dawn.com/news/1451213>, December 2018.
26. General Election Observation 2018: Key Findings and Analysis. <http://fafen.org/category/publications/featured-publications/>.
27. Constitution of Pakistan, part VIII: Chapter 2: Electoral Laws and Conduct of Elections. <http://www.pakistani.org/pakistan/constitution/part8.ch2.html>.
28. Elections act, 2017. http://www.na.gov.pk/uploads/documents/1506961151_781.pdf.
29. Orcutt, M. Internet Voting leaves out a Cornerstone of Democracy: The Secret Ballot. <https://www.technologyreview.com/s/602204/>, August 2016.

30. Tanoli, Q. 2.43 million Pakistanis working in Europe. <https://tribune.com.pk/story/1391730/overseas-workforce-2-43-million-pakistanis-working-europe/>, April 2017.
31. International Labour Organisation. Employer-Migrant Worker Relationships in the Middle East. https://www.ilo.org/beirut/publications/WCMS_552697/lang--en/index.htm, May 2017.
32. i-Voting - Registration Closed. <https://www.overseasvoting.gov.pk/i-voting/index.html>.
33. PILDAT. Scorecard on Perceptions of Pre-poll Fairness. https://pildat.org/wp-content/uploads/2018/05/PILDATScorecardonPerceptionofPre-PollFairness_May-2018.pdf, May 2018.
34. Khan, I. A. No haste in introducing i-voting for expats. https://epaper.dawn.com/DetailImage.php?StoryImage=12_09_2018_005_008, Sept 2018.
35. Staff Reporter. i-Voting being pushed to rig polls, says raza rabbani. <https://www.dawn.com/news/1431256/i-voting-being-pushed-to-rig-polls-says-raza-rabbani>, Sept 2018.
36. Junaidi, I. CDA promises proper waste disposal to Supreme Court. <https://www.dawn.com/news/1082080>, Jan 2014.
37. Shabbir, S. S. Judicial Activism Shaping the Future of Pakistan. <https://ssrn.com/abstract=2209067>, January 2013.
38. Supreme Court of Pakistan Diamer Basha and Mohmand Dam Fund. <http://www.supremecourt.gov.pk/DamFund/index.html>.
39. Younis, W. Risky voting. <https://www.dawn.com/news/1436323>, October 2018.
40. Junaidi, I. Electronic Voting not appropriate for Pakistan, senate body told. <https://www.dawn.com/news/1481153>, May 2019.

Implementing a public security scrutiny of an online voting system: the Swiss experience

Jordi Puiggali¹ [0000-0003-1472-415X]

¹ Scytl Secure Electronic Voting, S.A.
Research and Security Department
08008 Barcelona, Spain
jordi.puiggali@scytl.com

Abstract. During February and March 2019, Switzerland successfully conducted a Public Intrusion Test (PIT) event on the new Swiss Internet Voting system developed jointly with Scytl. The goal was to identify, categorize and correct vulnerabilities as early as possible, to improve the security of the voting system. This event drew the attention of most of the e-voting community, especially the attention of security researchers. The event took place in parallel with the publication of the source code of the voting system, that was also used to detect and design attacks for the PIT. A total number of 3187 researchers/teams have registered for the event and reported 173 findings. 16 of them were accepted under the category of “best practices”. Regarding source code findings, 84 were reported and 3 of them were deemed critical. The latest three findings, despite not having been exploited in the PIT, affected the main cryptographic properties for achieving complete verifiability: universal and individual verifiability. The attack related to individual verifiability also affected the version of the voting system used in production elections. However, the attack could be detected by auditors and it is excluded that past elections or votes have been manipulated. The production system was put on hold for scrutiny until the issues are solved and verified again by the experts.

Keywords: Switzerland, internet voting, regulation, verifiability, certification, source code publication, public intrusion test, online voting experiences

1 Introduction

Between February 25th and March 24th 2019, the security of the new Swiss Internet Voting system implemented jointly with Scytl (sVote) was publicly tested by researchers and security experts. The closest experience conducted before is the public test on the “D.C. Digital Vote-by-Mail Service” voting system piloted by the Washington, D.C. Board of Elections and Ethics (BOEE) in 2010 [17]. In this case, the voting system consisted of an open-source web-based platform for downloading and uploading PDF files through the internet.

Thus, sVote is the second internet-based voting system that has gone through a public intrusion test (PIT). The PIT conforms with the security and transparency require-

ments [1] set in August 2018 by the Electronic Voting Experts Group (EXVE) established by the Federal Chancellery. Prior to this experience, transparency was only achieved through the publication of the source code (excluding [17]). In 2011, within the context of the Norwegian Municipal Elections, the first case of source code publication took place, namely the Norwegian voting system, also developed by Scytal [2]. The source code of the Estonian voting system was partially published in 2013 [3] (the voting application was excluded). Finally, the Canton of Geneva published in 2016 part of the source code of its CHVote voting system (only the offline administration application), as a transparency effort [4].

The source code of the cryptographic protocol of the new sVote voting system was also published before starting the PIT. The publication is mandatory for any voting system compliant with the complete verifiability requirements of the Swiss Internet Voting legislation (VEleS) [5]. The new sVote voting system passed the audit certification processes for complete verifiability compliance in January 2019 and therefore, the source code was published under the previous registration before starting the authorization process from any Canton [22].

These two steps (source code publication and PIT) were important to detect and solve any issue that was not previously detected during the certification process. The new sVote is an evolution of the sVote voting system previously certified in September 2017 by the individual verifiability requirements of VEleS (Level 2 certification). Despite the differences between both voting systems (individual and complete verifiability requirements), the source code publication and PIT also served to detect issues that could affect the individual verifiability properties of the previous sVote voting system. In other words, the PIT was relevant since it could also affect the sVote voting system already in production.

In this paper, we analyse the source code publication and provide statistics arising from the PIT. Such event was open for a limited period (one month), but, as required by VEleS, the source code will remain publicly accessible. Interestingly, the analysis of the source code conducted by researchers has been concentrated during the PIT. This paper will cover all the aspects regarding the source code publication. We explain the requirements gathered in VEleS, the certification process, preparation of source code publication, and the results and impact of the source code publication during the PIT. The paper concludes with lessons learnt and conclusions that the author has taken from this experience.

The goal of this paper is therefore to share the acquired experience during the PIT related to source code publication. We hope this experience can be used by other governments or entities to establish transparency practices on electronic voting systems. Statements, analysis and conclusions are those of the author of this paper.

2 Complete verifiability in Switzerland

2.1 Online voting regulation (VEleS)

As described in [6], in 2014 a new online voting regulation (VEleS) [5] was established by the Swiss Federal Council regarding the deployment of electronic voting

systems in Switzerland. The new regulation introduces three levels of authorization for internet voting in the Cantons. Each level increases the security of the previous one. Systems are authorised to be used by up to 30% of the Canton electorate provided they meet a minimal set of functional and security requirements gathered by the Federal Chancellery (ChF). Systems authorised for up to the 50% of the Canton electorate must be compliant with the so-called individual verifiability requirements of VELeS. Finally, systems authorized for up to the 100% of the electorate must be compliant with the so-called complete verifiability requirements of VELeS. The latter requirements were enforced in 2018 by requesting publication of the source code of those software components in charge of providing complete verifiability. This enabled public scrutiny of these components (Art. 7a⁴ of VELeS). On top of the VELeS regulations, the Swiss Federal Council required, as good practice, that complete verifiable systems should be exposed to a Public Intrusion Test before its deployment [1].

The new sVote voting system is designed to be used by 100% of the electorate, and therefore needs to be compliant with the complete verifiability requirements of VELeS and PIT.

2.2 Complete verifiability requirements

According to the VELeS, a system with complete verifiability must achieve the following security requirements:

- Individual verifiability (VELeS Art. 5 para. 3): It must provide a proof to the voters that allows them to check that the vote received by the voting system contains their selected voting options (cast-as-intended and recorded-as-cast)
- Universal verifiability (VELeS Art. 5 para. 4-5): the voting system needs to provide proofs that the voting system did not modify any vote that has been individually verified by the voter. These proofs should be public auditable by the auditors or observers (universal) and should allow the verification of the voting (generation of the individual verifiable proofs) and counting (vote anonymization and decryption) phases.
- Use of a trustworthy environment (control components) for the generation of the individual and universal verifiable proofs (VELeS Art. 5 para. 6) under the trust assumption that at least one component must be honest for detecting any manipulation attempt (see Section 2.3.1 for more details).

The VELeS regulation comes with a technical annex (Tech. VELeS) [7] setting up the scope to evaluate voting systems that aim to be certified as individual and complete verifiable. The annex describes the trust model (abstract model) that is to be used for evaluation purposes. The abstract model defines the main components of the voting system, the trust assumptions made on these components, and the objectives that must be achieved under such trust assumptions.

The abstract model is at the core of the compliance process. A voting system must be accompanied with cryptographic (computational) and formal proofs (Tech. VELeS Ref. 5.1.1). Below, we provide an overview of the abstract model required for com-

plete verifiability. Also, the abstract model can be used to validate complete verifiability compliance of online voting systems.

2.3 Complete verifiability abstract model

The complete abstract model is defined in section 4.3 of the Tech. VEleS. The main system components according to the ordinance are:

- **Voters:** the actors that cast, verify and confirm their votes after previous authentication to the System component.
- **User platform:** the technical component used by the voters to interact with the System
- **Voter’s technical aids:** the technical aids independent from the User Platform components used by the voters to authenticate and cast their votes.
- **System (server-side):** the technical component used to authenticate voters, and to store, decrypt and count their votes.
- **Print Office:** the component used to print the authentication credentials and secret data required by voters to individually verify their votes (e.g., Return Codes)
- **Control Components:** the technical components interacting between them and the System component for proving the correct results.
- **Auditors:** the actors that check the proofs generated by the Control Components and System for confirming the correct results.
- **Auditor’s technical aids:** the technical aids used by the Auditors to universally verify the correctness of the proofs.

In addition to the above, the abstract model also defines the assumptions put on communication channels across system components to achieve the security requirements (see **Fig. 1**).

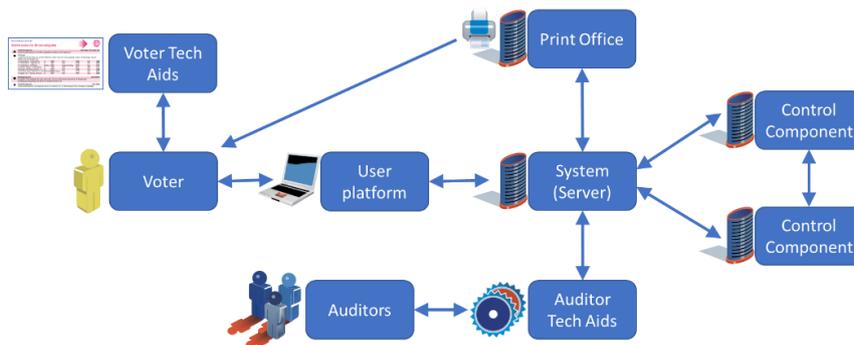


Fig. 1. Complete verifiability trust model components and interactions

2.3.1 Trust assumptions

The trust assumptions for individual verifiability (that is, for systems to be used by up to 50% of the electorate) are specified in the so-called reduced abstract model de-

scribed in section 4.1 of Tech. VEleS. The trust assumptions for complete verifiability (systems used by up to 100% of the electorate), extend the reduced model by removing the trustworthiness of the server System and trusting only one of the control components (without knowing which component exactly is trusted).

To ensure individual and universal verifiability, and as an exception to the above, the following components are considered trustworthy in the complete verifiability model:

- The Print Office
- The Voter Technical Aids
- The One Control Component
- One auditor
- The Auditor’s technical aids

To ensure vote secrecy the trust assumption is also extended to Voters and User Platforms. The abstract model lies the assumptions made in the cryptographic and formal proofs that must be provided to certify the system as completely verifiable. **Table 1** provides an example of how the new sVote voting system is aligned with the abstract model components and trust assumptions of VEleS [8].

VEleS	sVote	Trust assumption
Voters	Voters	Significant proportion of voters are non-trustworthy
User platform	Voting Client	Untrustworthy for individual and complete verifiability, trustworthy for privacy
Voter’s technical aids	Voting Card	Trustworthy
System (server-side)	Voting Server	Untrustworthy
Print Office	Print Office	Trustworthy
Control Components	Return Codes and Mixing Control Components	Trustworthy only as the whole. At least one is honest.
Auditors	Auditors	At least one is trustworthy
Auditor’s technical aids	Verifier	At least one honest auditor has a trustworthy aid

Table 1. - Mapping between components in the protocol and in the VEleS Swiss regulation

2.3.2 Objectives

Finally, the security objectives defined in the abstract model under the trust assumption mentioned above can be summarized as follows:

- To detect any attack against the votes processed by the system with a high probability (before, during or after being cast) through the verification proofs provided to the voters and/or auditors.

- To prevent attackers from compromising the secrecy of the votes cast by the system (under the trust assumption that the user platform is not controlled by an attacker).

Voting systems need to provide cryptographic and formal proofs for each of these objectives to demonstrate compliance with complete verifiability.

2.4 Verifying complete verifiability compliance

As already mentioned, the bases of the VELeS' certification process is to prove that the design of the voting system is compliant with the complete security abstract model and that the system is properly implemented in the software. This necessarily implies that the design, deemed compliant with the security model, is not broken when the voting system is implemented in practice (source code).

To this end, we can define a three-layer hierarchy, where each layer depends on the previous one. These layers are:

- **The regulatory layer:** sets-up the abstract model requirements. In the case of the Swiss regulation, we are talking about the VELeS regulation (an ordinance passed by the Swiss Federal Chancellery). The nexus between this layer and the next one (abstract layer) is the VELeS' technical annex.
- **The abstract layer:** in this layer, voting systems need to prove, in a mathematical sense, that they are compliant with the abstract model defined in the regulatory layer. The basis of this process is the VELeS' technical annex, and the way to certify compliance is by providing cryptographic and formal proofs compliant with the abstract model. In the case of complete verifiability, proofs must be provided to guarantee verifiability and vote secrecy.
- **The implementation layer:** this final layer contains the software development of a system based on the cryptographic protocol described in the abstract layer. The implementation of this layer is mainly based on two deliverables: the protocol specification (used for the development process) and the source code of the voting system and related documentation (deployment guides, audit documentation, etc.).

Based on the information managed at each layer we can identify the following main documentation and relationships (See **Fig. 2**):

- **VELeS:** Defines the security requirements and abstract model for achieving individual and complete verifiability, and how compliance with these requirements are to be verified.
- **Security and formal Proofs**
 - Security proof: Computational security proof of the proposed cryptographic protocol implemented according to the VELeS abstract model.
 - Formal proof: Formal security proof of the protocol defined in the security proof.

- **Protocol Specs:** Defines how to implement the protocol proven by the security formal proofs and how to deploy the system to ensure compliance with the security assumptions.
- **Source Code:** Source code that implements the processes described in the protocol specification.

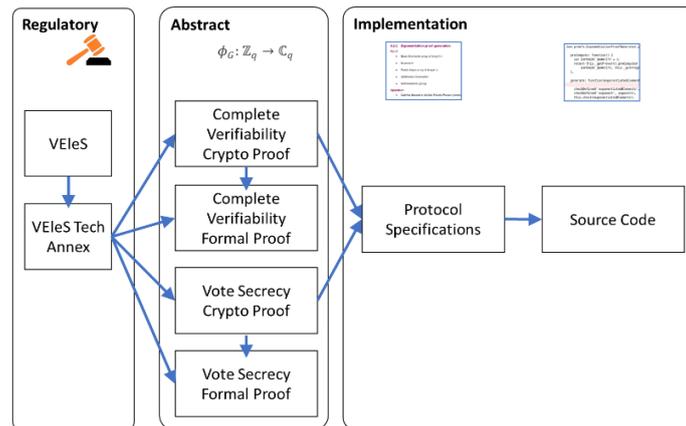


Fig. 2. Documentation relationship on the different compliance analysis layers

Ensuring that the documentation provided in the different layers is aligned is of paramount importance to guarantee that VELeS compliance covers the whole voting system design and implementation. As we will explain later, this has been the main cause of the findings notified during the PIT and source code publication.

3 Preparing public scrutiny of voting system

Federal Chancellery introduced the requirement of publication of the source code in a revision of the VELeS in August 2018. Article 7.b of the regulation details the requirements for publishing the source code. These requirements can be summarized as:

- **Quality:** The source code and related documentation must follow best practices.
- **Accessibility:** The source code and related documentation must be publicly available on the Internet and free of charge.
- **Completeness:** The published source code and documentation must cover all the relevant security aspects of the voting system.
- **Openness:** The source code and documentation terms of use must allow to analyse, modify, compile and execute the source code for study purposes. These terms will not ban the publication of the results of the analysis.

Therefore, the conduct of a PIT is not a requirement in the VELeS regulation. However, the internet voting Expert Group created by the Federal Council required it to be provided before first-time use of the system with a view to enhance transparency (Transparency section of [1]). This recommendation was considered relevant by Feder-

al Chancellery and Cantons, and a PIT was organized in addition to the publication of the source code to enforce the transparency of the process [20].

3.1 Source code publication

Considering the requirements from the VELeS, a public repository was configured to allow researchers to access the source code and related documentation, as well as a reporting environment to allow researchers to report their findings.

From an accessibility point of view, the source code was published in a GitLab repository [9], accessible upon request. Access to the repository requires a previous registration process where any interested party should provide basic contact information. The required information was name, surname, email address, reason for application and a github address [10]. The email and github address were the only ones verified (the first to ensure a valid contact and the second to ensure access to the source code). In this registration process, security researchers should accept the Conditions of Use [11] that include a Responsible Disclosure policy that ensures a grace period of up to 45 days since the last communication exchanged with the communication manager, before any finding can be made public by the participants (see clause 9 of [11]¹). The 45 days period was setup just in case a different period is not agreed between both parts. Currently, all the findings reported under this responsible disclosure were published in a shorter period. Despite responsible disclosure is a common accepted practice in software security testing, some potential participants did not accept its terms (arguing that disclosure clauses could be interpreted in a way that could allow indefinitely block any publication), or they just decided not to follow them. So, it was a matter of time until a replica of the official repository appeared on the Internet. Some of the main findings were reported by researchers that used an unofficial repository for making their research. However, these findings were accepted after researchers contacted the Federal Chancellery.

Another measure implemented for openness was the shared code license under which the source code was made public. Since the source code was not published as free open source software, it was important to publish it in a way that was compliant with the public scrutiny requirements of the regulation while keeping the IP rights from the authors. In this case, no issues were reported.

Finally, some issues were reported regarding the quality of the source code. Quality in this context does not cover security aspects, but software complexity, presence of test code (containing hardcoded passwords), unused code, or missing references to third party software. Most of these issues are related to the fact that the voting system per se implemented a complex cryptographic protocol that includes the distribution of operations in several components (control components). However, the feedback given is used to improve the current voting system quality and to solve some of the issues.

¹ “No Vulnerability shall be published within a period of forty five (45) days since the last communication exchanged with the Owners with regards to such potential Vulnerability, unless the Owners have agreed to a shorter period or defined a longer period.”

3.2 Public intrusion testing

The goal of the PIT was to identify vulnerabilities as early as possible and to correct them in order to strengthen and improve the security of the system. Furthermore, the PIT is a requirement of the Swiss government and the cantons and is therefore co-organized and supervised by the Swiss government and a specific group of Swiss cantons.

The scope of the PIT was based on the requirements from the Electronic Voting Experts Group [21] and focused on attacks from external actors with voter rights. Campaigns for incentivizing the participation started in August 2018 at the DEF CON® Hacking Conference in Las Vegas (USA). The Public Intrusion Test (PIT) has been executed from February 25th and until March 24th, 2019.

Category	Vulnerabilities	Compensation
Undetectable vote manipulation	<ul style="list-style-type: none"> - Manipulation of individual votes that is undetectable by voters and trusted auditors; - Scalable manipulation of votes that is undetectable by voters and trusted auditors; 	Between 30'000.- and 50'000.-
Vote manipulation	Manipulation of individual votes while maintaining universal verifiability mechanism (manipulation detectable by a trusted auditor) - e.g. the vote is modified after being cast;	20'000.-
Vote privacy (server-side)	<ul style="list-style-type: none"> - The privacy of a voter is broken (who voted) on the server; - The privacy of a vote is broken (what did they vote) on the server; 	10'000.-
Vote corruption	<ul style="list-style-type: none"> - A vote is stored in the ballot box and that vote cannot be decrypted; - A vote is stored in the ballot box in a way that gives the voter an unfair advantage; - Destruction of the electronic ballot box; 	5'000.-
Intrusion	<ul style="list-style-type: none"> - Intrusion into one of the servers (shell access); - Ability to execute arbitrary code on one or multiple servers; - Ability to execute arbitrary code on one or multiple control components; 	1'000.-
Best Practices	The configuration of a server or a service does not follow best practices of the security industry;	100.-

Table 2. - vulnerability categories and compensation (expressed in Swiss Francs CHF)

The PIT event was set-up similar to a bug bounty program and, therefore, additional considerations were made for managing the findings reported through this event: detected vulnerabilities were classified according to the impact of the finding with the prize that will be given to the researcher/team that reported it (see **Table 2**).

4 Statistics of the public intrusion testing

4.1 Global statistics

In total 3,186 people registered for the PIT. The total number of received findings on the PIT-Platform is 173. After careful inspection 16 were accepted and received a compensation.

System findings	Number
Total submitted	173
Of which confirmed by ChF, Cantons and Swiss Post	16
Of which critical	0
Non-critical optimization proposals	16

Table 3. – PIT accepted findings classification

4.2 Findings

From all those 173 submitted findings, 16 were accepted belonging to the category “best-practice” (see **Table 4**). There were no accepted findings falling into a higher category.

Vulnerability	Category	ID
Crafted X-Forwarded-For HTTP header injection	Best-practice	153
Missing HTTP to HTTPS redirection on 'pit-admin.evoting-test.ch'	Best-practice	166
Outdated version of Bootstrap Web Framework	Best-practice	168
Vulnerable TLS cipher-suites (LUCKY13)	Best-practice	175
Missing 'Expect-CT' HTTP header	Best-practice	179
Missing 'base-uri' in Content Security Policy	Best-practice	183
Incorrect 'HTTP-Strict-Transport-Security' header on 'pit-admin.evoting-test.ch'	Best-practice	188
Use of 'unsafe-eval' and 'unsafe-inline' in Content Security Policy	Best-practice	232
Multiple occurrences of 'X-XSS-Protection' HTTP header	Best-practice	234
Use of outdated version of AngularJS	Best-practice	257
Strict Transport Security Misconfiguration	Best-practice	272
Use of cipher suites without forward secrecy support	Best-practice	285
Missing charset declaration in some response's Content-Type header	Best-practice	294
Missing CSP header in redirect responses	Best-practice	295
Cross Origin Request possible on specific endpoint	Best-practice	296
Missing CSP header on http://pit-admin.evoting-test.ch/	Best-practice	318

Table 4. – PIT accepted vulnerabilities

Reported findings cover those that have been exploited in-scope of the PIT and those that constitute a misconfiguration according to industry best-practice. Findings which could not be exploited remotely, or which are only visible in the source code,

were rejected for the PIT but redirected to the source code program where they were evaluated again.

5 Results of the source code publication

5.1 Global statistics

System findings	Number
Total submitted	84
Of which accepted for revision	28
Of which optimization proposals	25
Critical findings	3

Table 5. – Source code participation statistics (May 2019)

Up to the moment of writing this paper (source code access is still open), there are about 1,600 registered users with access to the source code repository. **Table 5** shows the participation statistics on the source code review. The first critical finding was reported by 3 different sources, however we considered it as one unique finding.

5.2 Findings

As shown in **Table 5** the statistics from the 28 submissions accepted for revision, 25 were considered as improvements and only 3 reported as security findings. The improvements range from proposals of using larger key sizes to mechanisms for improving the verifiability of write-ins. All these improvements will be considered in the future improvements of the new sVote voting system, but were not considered as critical for implementing them, since they do not imply any security vulnerability. Registered users could see the proposals of the others in the environment that had been setup for reporting them.

The main critical findings were the three that affected the core cryptographic verifiability features of the voting system: universal and individual verifiability. The first finding affected the universal verifiability of the Mixing process implemented in the first control component (the one that anonymizes the votes before decryption). The second affected the universal verifiability of the partial decryption process implemented also in the first control component. Finally, the last one affected individual verifiability of the Return Codes generated in the voter device. While all these findings were reported with examples of theoretical attacks, none of these attacks were carried out in practice in the PIT platform. It can be excluded that past votes or elections have been manipulated because of these findings, since the attack always generates invalid votes and such votes have never been reported in previous elections.

A more detailed information of the three findings follows.

5.2.1 Universal verifiability – Mixing proof

This finding was detected in the protocol implementation but was not present in the protocol design (i.e., it was not present in the cryptographic and formal proofs). The issue was related to a missing verification in the independent generation of the commitment parameters of the Mixing proofs. The Mixing process is used to anonymize the votes before decryption, shuffling and re-encrypting the votes using the election public key. Since the process prevents to correlate the output shuffled and re-encrypted votes from the input ones, an attacker could use this property to substitute the output votes by other encrypted ones. To prevent this, the voting sVote system implemented a universal verifiable proof that shows that the Mixnet did not modify any content of the votes during the process. This proof is based on the proposal made by Bayer-Groth [18].

The issue was specifically related to the way that the commitment parameters used to generate the Bayer-Groth proofs were generated: these parameters were generated in a random way, but this generation cannot be verified as independent. This opens the door to an attacker controlling the Mixing node to generate these commitment parameters in a specific way (mathematical relationship). This would allow them to modify the output of the mixing and generate a fake proof that would be accepted as valid and therefore, modify vote contents without being detected.

To exploit this issue, it is necessary to perform some mathematical calculations based on the cyphertext and the internal structure of the vote. Therefore, the exploitation is not straightforward (e.g., cannot use the sVote software but another one developed by the attacker). In fact, the main theoretical attack identified by the researchers requires the attacker to control also the vote casting process in the voter device to learn the randomness used to encrypt the vote. Therefore, the attacker needs to control the voter device used to cast each vote that they want to manipulate. Another constraint of the reported attack was that it requires to control the first control component, otherwise it is not feasible. A second theoretical attack was also proposed, but the internal structure of the vote made it not feasible.

In any case, the importance of the detected finding was not related to the practical feasibility of the attacks proposed but the fact that the universal verifiability of the Mixing proof cannot be guaranteed. Therefore, it is necessary to correct the issue to restore universal verifiability property. In fact, the solution was easy to implement and already designed and present in the source code (the use of a FIPS 186-4 verifiable generation of a random group member), but it was not mentioned in the protocol specifications to use this verifiable generation instead of the standard one. This remarks the importance of ensuring that the protocol design and implementation are properly aligned.

This finding was independently reported by two individual researchers, Rolf Haenni (BfH) [13] and Thomas Haines (NTNU), through the official reporting environment, and a research team, composed by Sarah Jamie Lewis (Open Privacy Research Society), Olivier Pereira (UCLouvain), and Vanessa Teague (University of Melbourne) [14], by other means.

5.2.2 Universal verifiability – Decryption proof

This second finding was reported by the same research team of the first finding (Lewis-Pereira-Teague) [15] and affected the universal verifiability of the decryption proof. In this case, the root of the issue was that the Non-Interactive Zero Knowledge Proof (NIZKP) of the decryption process was not using the statement proved for the generation of the non-interactive heuristic (Fiat-Shamir) [19]. This does not guarantee the soundness of the proof in an adaptative scenario (when the attacker can manipulate the information that needs to be proven) and therefore, the universal verifiability of the decryption process. In a non-adaptative scenario, information of the statement that is not under the control of the attacker (e.g., public information) is not needed to be used to generate the non-interactive heuristic (weak Fiat-Shamir). Otherwise, the statement needs to be used (strong Fiat-Shamir).

In the new sVote voting system, control components implement a Mixing process and then a partial decryption of the mixed votes during the vote counting phase. Therefore, this configuration opens the door to an adaptative scenario if the attacker is fully controlling a control component: the attacker can play with the re-encryption transformation made by the Mixing to generate a ciphertext that can help him to generate a fake proof that will pass any validation. The research team notified again about this issue and reported a theoretical attack. The attack requires also a similar scenario as in the first attack: full control of the first control component and each voter device used for each voter whose vote will be manipulated. However, this attack has a constraint since it can only substitute the vote by random data. Therefore, this attack will always be detected when the vote is decrypted in the last control component: the vote contains non-sense data and therefore, is excluded from the count and isolated as an auditable vote. This situation is not possible in the cryptographic protocol, because invalid votes cannot be accepted by the voting system: cannot generate a valid return code.

Therefore, the main impact of the finding was that universal verifiability of the decryption process cannot be guaranteed and therefore, it is not possible to prove that the decryption is correct. Therefore, despite the attack is detectable (generates a vote with non-sense data), it requires to make an investigation to find the source of the issue.

Again, the issue was detected in the protocol implementation (protocol specification) but not in the protocol design (cryptographic proof).

5.2.3 Individual verifiability – Exponentiation proof

This is the last flaw detected by the same research team [16], and this time the impact was larger than the impact of the issues mentioned above, since it also affected the individual verifiability property of the sVote voting system already used in previous elections (i.e., production).

The finding was similar to the second one, since it was related to the information used in the non-interactive transform (Fiat-Shamir). However, this time it affected the exponentiation proof made in the voting client to prove that the information used to generate the Return Codes contains the same options than the vote cast by the voter. The issue did not happen because the statement was not used by the non-interactive transform, but since the transform only used part of the statement information: the

ElGamal encryption has two components (c_0, c_1) and in this case the protocol specification failed to require that both be included in the statement. Once again, the cryptographic proof (protocol design) was correct. Individual verifiability was based on the combination of two proof systems: exponentiation proof and a proof of plaintext equivalence. The plaintext equivalence was properly implemented in the design and implementation documents, so it was sound and not vulnerable to any attack.

The research team reported the finding again and provided a theoretical attack. The attack shows how to generate the Return Codes expected by the voter while casting an encrypted vote with different contents. The main limitation of the attack is that the attacker only can mathematically find a vote with random data. Therefore, despite the possibility that the attack can cheat the voter, it will be always detected by the auditors when the votes are decrypted since the tampered vote will contain non-sense data. As mentioned in relation to the second finding, votes with non-sense data are not possible and are detected as an attack. This property has been used to demonstrate that, despite the vulnerability has been present in the sVote voting system in production, no election has been manipulated since votes with non-sense data were never reported (i.e., Cantons never notified in the official counting report the presence of auditable votes).

In any case, the finding was critical since it impacts the individual property of the voting system and therefore, cannot provide guarantees to the voter that their votes are cast-as-intended. For this reason, it was decided to put the sVote voting system in production on hold until the issue was solved, and further scrutiny is done to ensure that the issue has been properly tested and solved.

5.3 Lessons learnt and future improvements

One of the main lessons learnt of this experience is the benefits of opening the source code to public scrutiny. It allows to find issues that have been undetected by other reviewers but could have an impact on the cryptographic properties. In fact, all the issues detected were not related to an improper protocol design, but to a missing step or information when describing the protocol at implementation level (protocol specifications). Therefore, they were not detected by experts just reviewing cryptographic and formal proofs. This made us revisit and improve the processes and tools used to keep the protocol design and implementation documentation aligned. That way, if the protocol has been computationally and formally proven, this assurance can be also applied to the implementation.

Another lesson learnt is that there is still room for improvement on how to manage the source code publication of voting systems. Specially, how to manage situations in which a vulnerability in the voting system is found near or in the middle of an election. One of the advantages of electronic voting systems is that issues can be corrected very fast, reducing the impact of availability. However, the main issue is the impact on the confidence of the voting system, that is usually more sensible to these issues than in other more traditional voting channels.

From the point of view of participation in the source code review, it is still too early to make some analysis. However, based on Scytl participation in previous source code publication experiences (e.g., Norway 2011 and 2013), the amount of people that participated in the source code review greatly exceeded expectations. Maybe this hap-

pened in part because the proximity of the PIT experience, but it is important to follow up the activity during the next events (specially in periods near election processes).

6 Conclusions²

The first conclusion from the experience is that its contribution has been relevant for improving the security of internet voting systems.

Another conclusion from the experience is the importance of having strong measures to ensure that what has been proven as secure from a design point of view (cryptographic and formal proofs), is then implemented following the same design principles. This is especially relevant considering the dynamic nature of the software and technology, since changes on the architecture can imply changes in the implementation that generate misalignments. In this case, in addition to better traceability measures between design and implementation, external reviews are important to detect issues that remained hidden to internal reviews.

A third conclusion, that we also detected in other experiences in other countries, is the lack of a common consensus of which mechanism must be used for reporting issues to the election managers. This is especially relevant when opening the source code, since the probability of reporting findings is higher than when the voting system is only accessible during the election period. Responsible disclosure agreements are usually the way to manage this on the security industry, to provide a balance between the protection of the reporter before making their findings public and the time needed by the provider for solving the issue before it is made public. However, not all the participants agreed on the responsible disclosure terms or followed them. For this reason, we propose to make an open debate on this aspect for reaching a common way for managing public disclosure and reporting of findings.

To conclude, it is crucial to identify and correct the issues raised by this experience (e.g., the terms and conditions required to researchers) in order to make it easier to replicate such experiences in the future (e.g., in any other countries).

References

1. Swiss Federal Chancellery. Rapport final du groupe d'experts Vote électronique (GE VE) (26.06.2018).
2. Barrat, Jordi; Goldsmith, Ben and Turner, John: International Experience with E-Voting - Norwegian E-Vote Project. International Foundation for Electoral Processes (IFES). (2012).
3. IVXV online voting system. <https://github.com/vvk-ehk/ivxv>
4. République et Canton de Geneve – CHVote 1.0. <https://github.com/republique-et-canton-de-geneve/chvote-1-0>
5. Swiss Federal Chancellery. Federal Chancellery Ordinance on Electronic Voting (VEleS) of 13 December 2013 (OEV, SR 161.116). (Status as of 1 July 2018).

² **Acknowledgement:** We would like to congratulate to all the researchers/teams that participated in the experience for the valuable feedback, and especially we would like to thank those that found the main critical issues in the implementation of cryptographic parts of the voting system.

6. Puiggalí, Jordi and Rodríguez-Pérez, Adrià. Defining a national framework for online voting and meeting its requirements: the Swiss experience. In: Proceedings E-Vote-ID 2018, TUT Press. pp. 82-97 (2018).
7. Swiss Federal Chancellery. Technical and administrative requirements for electronic vote casting. Annex to the FCh Ordinance of 13, December 2013 on Electronic Voting (OEV, SR 161.116). (Status as of 1 July 2018).
8. Scytl Secure Electronic. Scytl sVote. Protocol Specifications. Software version 2.1, document version 5.1. 2018.
9. Swiss Post. sVote 2.1 source code public access repository. <https://gitlab.com/swisspost/evoting-solution>
10. Swiss Post. Source code access registration form. <https://www.evoting.ch/sourcecode/ui/home?lang=en>
11. Swiss Post. Electronic Voting Solution Source Code Access Agreement. January 2019. <https://www.post.ch/-/media/post/evoting/dokumente/nutzungsbedingungen-quellcode.pdf?la=en&vs=2>
12. Swiss Post. Public Intrusion Test – Code of Conduct. <https://www.onlinevote-pit.ch/conduct/>
13. Haenni, Rolf. Swiss Post Public Intrusion Test: Undetectable attack against vote integrity and secrecy. 2019. <https://e-voting.bfh.ch/publications/2019/>
14. Lewis, Sarah Jamie; Pereira, Olivier and Teague, Vanessa. Ceci n'est pas une preuve: The use of trapdoor commitments in bayer-groth proofs and the implications for the verifiability of the Scytl-SwissPost internet voting system, 2019.
15. Lewis, Sarah Jamie; Pereira, Olivier and Teague, Vanessa. How not to prove your election outcome: The use of non-adaptive zero knowledge proofs in the Scytl-SwissPost Internet voting system, and its implications for decryption proof soundness, 2019.
16. Lewis, Sarah Jamie; Pereira, Olivier and Teague, Vanessa. How not to prove the validity of your ballot the use of non-adaptive zero knowledge proofs in the Scytl-SwissPost Internet voting system, and its implications for individual verifiability, 2019
17. Wolchok, Scott; Wustrow, Eric; Isabel, Dawn; Halderman, J.Alex. Attacking the Washington, D.C. Internet Voting System. In: Keromytis A.D. (eds) Financial Cryptography and Data Security. FC 2012. Lecture Notes in Computer Science, vol 7397. Springer, Berlin, Heidelberg
18. Bayer, Stephanie and Groth, Jens. Efficient zero-knowledge argument for correctness of a shuffle, Advances in Cryptology - Eurocrypt, 2012.
19. Fiat, Amos and Shamir, Adi. How to prove yourself: Practical solutions to identification and signature problems, Conference on the Theory and Application of Cryptographic Techniques, 1986.
20. Swiss Federal Chancellery. Fact sheet produced by the Management Committee of the Confederation and the Cantons. February 2019. https://www.bk.admin.ch/bk/en/home/politische-rechte/e-voting/oeffentlicher_intrusionstest.html
21. Swiss Federal Chancellery. Federal and cantonal requirements regarding public intrusion tests. https://www.bk.admin.ch/bk/en/home/politische-rechte/e-voting/oeffentlicher_intrusionstest.html
22. Swiss Federal Chancellery. Public intrusion test for e-voting to take place in February and March. <https://www.bk.admin.ch/bk/en/home/dokumentation/medienmitteilungen.msg-id-73898.html>

Attacks and Security Requirements in Practice

UnclearBallot: Automated Ballot Image Manipulation

Matthew Bernhard, Kartikeya Kandula, Jeremy Wink, and J. Alex Halderman

Department of Computer Science and Engineering, University of Michigan
{matber, kartkand, jr Jeremy, jhalderm}@umich.edu

Abstract. As paper ballots and post-election audits gain increased adoption in the United States, election technology vendors are offering products that allow jurisdictions to review ballot images—digital scans produced by optical-scan voting machines—in their post-election audit procedures. Jurisdictions including the state of Maryland rely on such image audits as an alternative to inspecting the physical paper ballots. We show that image audits can be reliably defeated by an attacker who can run malicious code on the voting machines or election management system. Using computer vision techniques, we develop an algorithm that automatically and seamlessly manipulates ballot images, moving voters’ marks so that they appear to be votes for the attacker’s preferred candidate. Our implementation is compatible with many widely used ballot styles, and we show that it is effective using a large corpus of ballot images from a real election. We also show that the attack can be delivered in the form of a malicious Windows scanner driver, which we test with a scanner that has been certified for use in vote tabulation by the U.S. Election Assistance Commission. These results demonstrate that post-election audits must inspect physical ballots, not merely ballot images, if they are to strongly defend against computer-based attacks on widely used voting systems.

Keywords: optical scan, paper ballots, image manipulation, drivers, image processing

1 Introduction

Elections that cannot provide sufficient evidence of their results may fail to adequately gain public confidence in their outcomes. Numerous solutions have been posited to this problem [9], but none has been as elegant, efficient, and immediately practical as post-election audits [21, 25, 39]. These audits—in particular, ones that seek to limit the risk of confirming an outcome that resulted from undue manipulation—are one of the most important layers of defense for election security [32].

Risk-limiting audits (RLAs) rely on sampling robust, independent evidence trails created by voter-verified paper ballots. However, other types of post-election

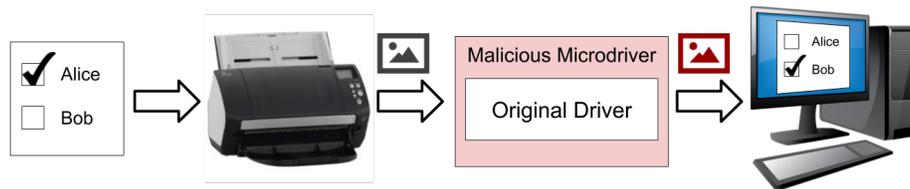


Fig. 1. Attack overview—A voter’s paper ballot is scanned by a ballot tabulator, producing a digital image. Malware in the tabulator—in our proof-of-concept, a micro-driver that wraps the scanner device driver—alters the ballot image before it is counted or stored. A digital audit shows only the manipulated image.

audits are gaining popularity in the marketplace. In particular, Clear Ballot, an election technology vendor in the United States, pioneered audit software designed to perform audits of *images* of ballots which have been scanned and tabulated, which we shall refer to as “image audits”. Other vendors have adopted support for this kind of audit, and one U.S. state, Maryland, relies on image audits to provide assurances of its election results [33].

While image audits can help detect human error and aid in adjudicating mismarked ballots, we show that they cannot provide the same level of security assurance as audits of physical ballots. Since ballot images are disconnected from the actual source of truth—physical paper ballots—they do not necessarily provide reliable evidence of the outcome of an election under adversarial conditions.

In this paper, we present UnclearBallot, an attack that defeats image audits by automatically manipulating ballot images as they are scanned. Our attack leverages the same computer vision approaches used by ballot scanners to detect voter selections, but adds the ability to move marks from one target area to another. Our method is robust to inconsistent or invalid marks, and can be adapted to many ballot styles.

We validate our attack against a corpus of over 180,000 ballot images from the 2018 election in Clackamas County, Oregon, and find that UnclearBallot can move marks on 34% of the ballots while leaving no visible anomalies. We also test our attack’s flexibility using six widely used styles of paper ballots, and its robustness to invalid votes using an established taxonomy of voter marks. As a proof-of-concept, we implement the attack in the form of a malicious Windows scanner driver, which we test using a commercial-off-the-shelf scanner certified for use in elections by the U.S. Election Assistance Commission.

UnclearBallot illustrates that post-election audits in traditional voting systems must involve rigorous examination of *physical ballots*, rather than ballot images, if they are to provide a strong security guarantee. Without an examination of the physical evidence, it will be difficult if not impossible to assure that computer-based tampering has not occurred.

The remainder of this paper is organized as follows: Section 2 provides background on image audits, ballot scanners, and image processing techniques we use to implement our attack. Section 3 describes the attack scenarios against

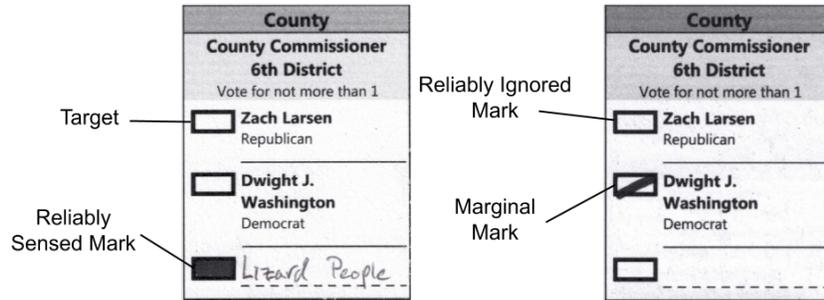


Fig. 2. Terms for parts of a marked ballot, following Jones [23].

optical scanners and image audits. Section 4 explains the methodology of our attack. In Section 5 we present data indicating that our attack can be robust to various ballot styles and voter marks. Section 6 contextualizes our attacks and discusses mitigations. We conclude in Section 7.

2 Background

Our attack takes advantage of two aspects of optical scanner image audits: the scanning and image processing techniques used by scanners, and the reliance on scanned images by image audits. Here we provide a brief discussion of both.

2.1 Ballot Images

Jones [23] put forth an analysis of the way that ballot scanners work, particularly the mark-sense variety that is most common today. All optical scanners currently sold to jurisdictions, as well as the vast majority of scanners used in practice in the U.S., rely on mark-sense technology [44]. Scanners first create a high-resolution image of a ballot as it is fed past a scan head. Software then analyzes the image to identify dark areas where marks have been made by the voter.¹ Once marks have been detected, systems may use template matching to translate marks into votes for specific candidates, typically relying on a barcode or other identifier on the ballot that specifies a ballot style to match to the scanned image.

Detecting and interpreting voter marks can be a difficult process, as voters exhibit a wide range of marking and non-marking behavior, including not filling in targets all the way, resting their pens inside targets, or marking outside the target. The terms Jones developed to refer to the ballot and marks are illustrated in Figure 2. Marks that adequately fill the target and are unambiguously interpreted as votes by the scanner are called *reliably sensed* marks, and targets that are unambiguously not filled and therefore not counted are *reliably ignored* marks.

¹ The details of how marks are identified vary by hardware and scanning algorithm. See [13] for an example.

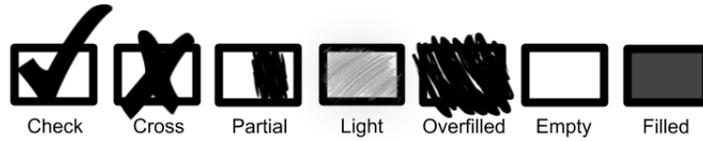


Fig. 3. Taxonomy of voter marks adapted from Bajcsy [2], including the five leftmost marks that may be considered marginal marks.

Marks of other types are deemed *marginal*, as a scanner may read or ignore them. Moreover, whether a mark should be counted as a vote is frequently governed by local election statute, so some marginal marks may be unambiguously counted or ignored under the law, even if not by the scanner.

Bajcsy et al. [2] further develops a systematization of marginal marks and develops some improvements on mark-detection algorithms to better account for them. An illustration of Bajcsy et al.’s taxonomy is shown in Figure 3. Ji et al. [22] discuss different types of voter marks as applied to write-in votes, as well as developing an automated process for detecting and tabulating write-in selections.

2.2 Image Audits

Risk-limiting post-election audits rely on physical examination of a statistical sample of voter-marked ballots [24, 26, 39, 40]. However, this can create logistical challenges for election officials, which has prompted some to propose relaxations to traditional audit requirements. To reduce workload, canvass audits and recounts in many states rely on retabulation of ballots through optical scanners (see the 2016 Wisconsin recount, for example [31]).

Some election vendors take retabulation audits a step further: rather than physically rescan the ballots, the voting system makes available images of all the ballots for independent evaluation after the election [15, 16, 42].² While the exact properties of these kinds of image audits vary by vendor, they typically rely on automatically retabulating all or some images of cast ballots, as well as electronic adjudication for ballots with marginal marks. These “audits” never examine the physical paper trail of ballots, which our attack exploits.

Several jurisdictions have relied on these image audits, including Cambridge, Ontario, which used Dominion’s AuditMark [17], and the U.S. state of Maryland, which uses Clear Ballot’s ClearAudit [28]. Maryland has also codified image audits into its election code, requiring that an image audit be performed after every election [27].

² While the review is made available to the public, the actual images themselves are seldom published in full out of concern for voter anonymity.

3 Attack Scenarios

Elections in which voters make their selections on a physical ballot are frequently held as the gold standard for conducting a secure election [32]. However, the property that contributes most to their security, software independence [34], only exists if records computed by software are checked against records that cannot be altered by software without detection. Image audits enable election officials to view images of ballots and compare them with the election systems' representation of the particular ballot they are viewing (called a cast vote record or CVR). While these two trails of evidence may be independent from each other (for example, Clear Ballot's ClearAudit [15] technology can be used to audit a tabulation performed by a different election system altogether), they are not software independent. A clever attacker can exploit the reliance on software by both evidence trails to defeat detection.

To surreptitiously change the outcome of the election in the presence of an image audit, the attacker must alter both the tabulation result as well as the ballot images themselves. Researchers have documented numerous vulnerabilities that would allow an attacker to infect voting equipment and change tabulation results (see [10, 20, 30] among others), so we focus on the feasibility of manipulating ballot images once an attacker has successfully infected a machine where they are stored or processed.

The most straightforward attack scenario occurs when the ballot images are created by the same equipment that produces the CVR. In this case, the attacker can simply infect the scanner or tabulator with malware that corrupts both the CVR and the images at the same time. The attack could change the image before the tabulator processes it to generate the CVR, or directly alter both sets of records.

In some jurisdictions, the ballot images that are audited are collected in a separate process from tabulation—that is, by scanning the ballots again, as in Maryland's use of ClearAudit from 2016 [28]. In this case, the adversary has to separately attack both processes, and has to coordinate the cheating to avoid mismatches between the initial tally and the altered ballot images.

Depending on the timing of the audit, manipulation of ballot images need not be done on the fly. For example, if the ballot images are created during tabulation but the image audit does not occur until well after the election, an attacker could modify the ballot images while they are in storage.

For ease of explication, the discussion that follows assumes that ballot images are created at the time of tabulation, in a single scan. The attack we develop targets a tabulation machine and manipulates each ballot online as it is scanned.

4 Methodology

To automatically modify ballot images, an attacker can take a few approaches. One approach would be to completely replace the ballot images with ballots filled in by the attacker. However, this risks being detected if many ballots have

the same handwriting, and requires sneaking these relatively large data files into the election system without being detected. For these reasons, we investigate an alternative approach: automatically and selectively doctoring the ballot scans to change the vote selections they depict.

For the attack to work successfully, we need to move voter marks to other targets without creating visible artifacts or inconsistencies. We must be able to dynamically detect target areas and marks, alter marks in a way that is consistent with the voter’s other marks, and do so in a way that is undetectable to the human eye. However, there is a key insight that works in the adversary’s favor: an attacker seeking to alter election results does not have to be able to change *all* ballots undetectably, only sufficiently many to swing the result. This means that the attacker’s manipulation strategy is not required to be able to change *every* mark—it merely has to reliably detect *which* marks it can safely alter and change enough of them to decide the election result.

4.1 Reading the ballot

To interpret ballot information, we rely on the same techniques that ballot scanners use to convert paper ballots into digital representations. Attackers have access to the ballot templates, as jurisdictions publish sample ballots well ahead of scheduled elections. Using template matching, an attacker does not have to perform any kind of sophisticated character recognition, they simply have to find target areas and then detect which of the targets are filled.

Our procedure to read a ballot is illustrated in Figure 4. First, we perform template matching to extract each individual race within a ballot. Next, we use OpenCV’s [11] implementation of the Hough transform to detect straight lines that separate candidates and break the race into individual panes for each candidate. Notably, the first candidate in each race may have the race title and extra information in it (see Figure 4c), which is cropped out based on white space.

Target areas are typically printed on the ballot as either ovals or rectangles. To detect them, we construct a bounding box around the target by scanning horizontally from the left of the race and then vertically from the bottom up, and compute pixel density values. The bounds are set to the coordinates where the density values first increase and last decrease. Once we have detected all the target areas, we compute the average pixel density of the area within the bounding box to determine whether or not a target area is marked. We then use our template to convert marks into votes for candidates.

4.2 Changing marks

Once we have identified which candidate was marked by the voter, we can move the mark to one of the other target locations we identified. If the vote is for a candidate the attacker would like to receive fewer votes—or if it is not a vote for a candidate they would like to win—the attacker can simply swap the pixels within the bounding boxes of the voter’s marked candidate and an unmarked candidate.

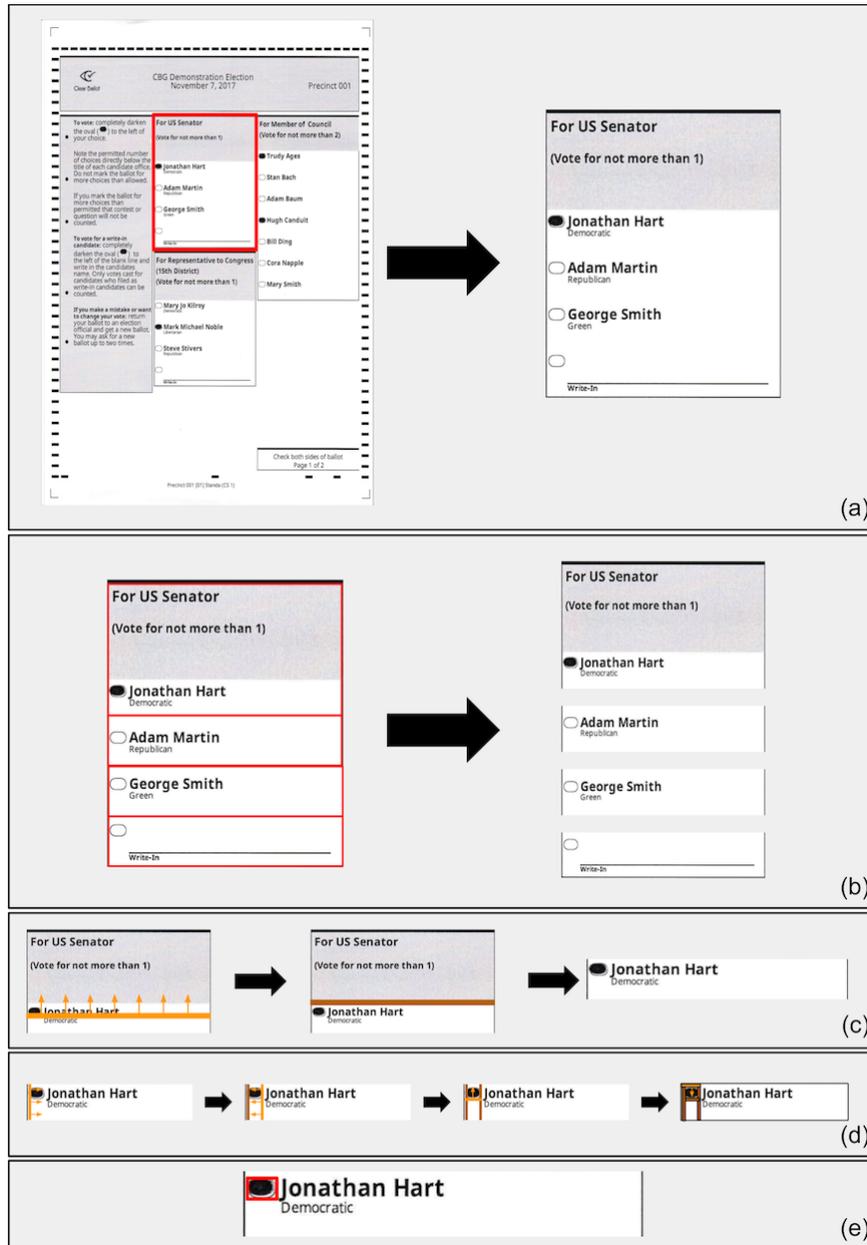


Fig. 4. Ballot manipulation algorithm—First, (a) we apply template matching to extract the race we intend to alter. Then, (b) we use Hough line transforms to separate each candidate. If the first candidate has a race title box, (c) we remove it by computing the pixel intensity differences across a straight line swept vertically from the bottom. For each candidate, (d) we identify the target and mark (if present) by doing four linear sweeps and taking pixel intensity. Finally, (e) we identify and move the mark. At each step we apply tests to detect and skip ballots where the algorithm might leave artifacts.

Original		Manipulated	
County		County	
Supervisor, District 1		Supervisor, District 1	
Vote for One		Vote for One	
Alfred Hitchcock		Alfred Hitchcock	<input type="radio"/>
Vincent Price	<input type="radio"/>	Vincent Price	
Write In	<input type="radio"/>	Write In	<input type="radio"/>
State		State	
Governor		Governor	
Vote for One		Vote for One	
Amelia Earhart	<input type="radio"/>	Amelia Earhart	
Howard Hughes		Howard Hughes	<input type="radio"/>
Charles Lindbergh	<input type="radio"/>	Charles Lindbergh	<input type="radio"/>
Write In	<input type="radio"/>	Write In	<input type="radio"/>

Fig. 5. Automatically moving voter marks—UnclearBallot seamlessly moves marks to the attacker’s preferred candidate while preserving the voter’s marking style. It is effective for a wide variety of marks and ballot designs. In the examples above, original ballot scans are shown on the left and manipulated images on the right.

By moving marks on each ballot separately, we ensure that the voter’s particular style of filling in an oval is preserved and consistent across the ballot. Figure 5 shows some marks swapped by our algorithm, and how the voters original mark is completely preserved in the process.

4.3 UnclearBallot

To illustrate the attack, we created UnclearBallot, a proof-of-concept implementation packaged as a malicious Windows scanner driver, which consists of 398 lines of C++ and Python. We tested it with a Fujitsu fi-7180 scanner (shown in Figure 6), which is federally certified for use in U.S. elections as part of Clear Ballot’s ClearVote system [43]. These scanners are typically used to handle small volumes of absentee ballots, and must be attached to a Windows workstation that runs the tabulation software.

The UnclearBallot driver wraps the stock scanner driver and alters images from the scanner before they reach the election management application. We chose this approach for simplicity, as the Windows driver stack is relatively easy



Fig. 6. The **Fujitsu fi-7180 scanner** we used to test our attack has been certified by the U.S. Election Assistance Commission for use in voting systems. Our proof-of-concept implementation is a malicious scanner driver that alters ballots on the fly.

to work with, but the attack could also be implemented at other layers of the computing stack. For instance, it could be even harder to detect if implemented as a malicious change to the scanner’s embedded firmware. Alternatively, it could be engineered as a modification to the tabulation software itself.

Once a ballot is scanned, the resulting bitmap is sent to our image processing software, which manipulates the ballot in the way described in Section 4.1. Prior to the election, the attacker specifies the ballot template, which race they would like to affect, and by how much. While ballots are being scanned, the software keeps a running tally of the actual ballot results, and changes ballot images on the fly to achieve the desired election outcome. To avoid detection, attackers can specify just enough manipulated images so that the race outcome is changed.

5 Evaluation

We evaluated the performance and effectiveness of UnclearBallot using two sets of experiments. In the first set of experiments, we marked different ballot styles by hand using types of marks taxonomized by Bajcsy et al. [2]. In the second set of experiments, we processed 181,541 ballots from the 2018 election in Clackamas County, Oregon.

5.1 Testing Across Ballot Styles

In order for our application to succeed at its goal (surreptitiously changing enough scanned ballots to achieve a chosen election outcome), it must be able to detect marks that constitute valid votes as well as distinguish marks which would be noticeable if moved. The marks in the latter case represent a larger set than just marginal marks, as they may indeed be completely valid votes, but considered invalid by our mark-moving algorithm. For example, if we were to swap the targets on a ballot where the user put a check through their target, we may leave a significant percentage of the check around the original target when swapping. The same applies for marked ballots where the filled in area extends into the candidate’s name, which could lead our algorithm to swap over parts of the candidate’s name when manipulating the image.

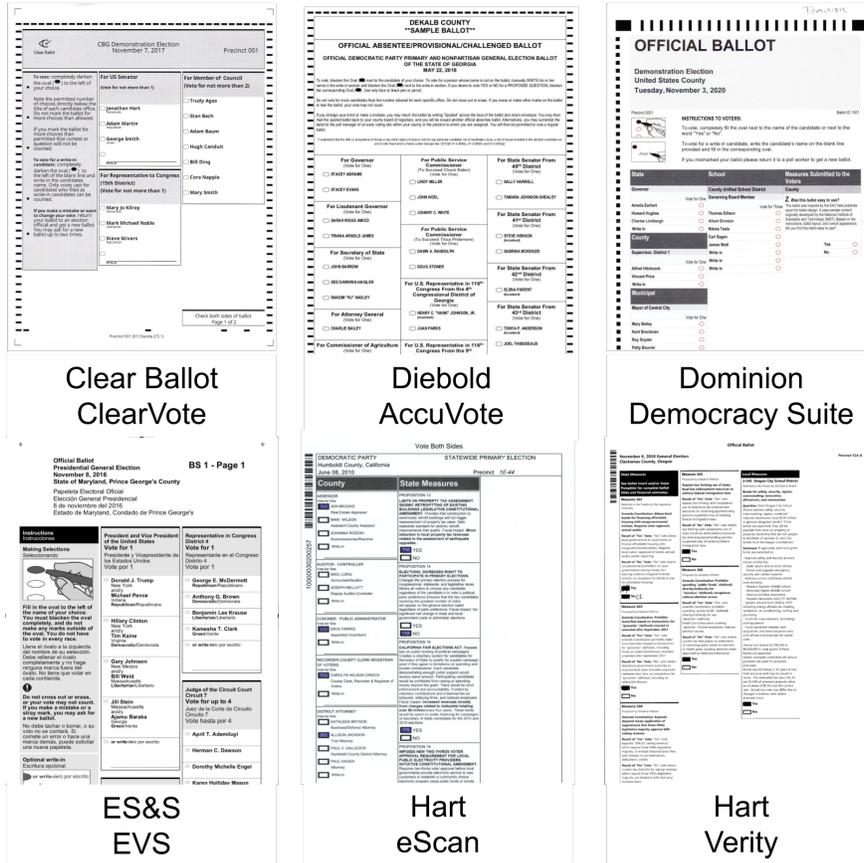


Fig. 7. Ballots Styles— We tested ballot designs from five U.S. voting system vendors: Clear Ballot, Diebold, Dominion, ES&S, and Hart (two styles, eScan and Verity).

To detect anomalies for invalid ballots, we leverage the same intensity checking algorithm that first found the marked areas. The program checks if the width or height is abnormally large, which would indicate an overfilled target, as well as if there are too few or too many areas of high intensity, which would indicate no target or too many targets are filled out. If the program detects an invalid ballot, it will not be modified by the program.

To show our attack is replicable on a variety of different ballot styles, we modified our program to work on six different sample ballot styles, shown in Figure 7. The ballots we tested come from the four largest election vendors in the U.S. (ES&S, Hart InterCivic, Dominion, and Clear Ballot), as well as two older styles of ballots from Hart and Diebold.

Our first experiment was designed to characterize the technique’s effectiveness across a range of ballot styles and with both regular and marginal marks. We

Ballot Style	Invalid Marks			Valid Marks			Time/Success
	Skipped	Success	Failure	Skipped	Success	Failure	
Clear Ballot	55	5	0	26	34	0	25 ms
Diebold	60	0	0	6	54	0	11 ms
Dominion	38	22	0	7	53	0	30 ms
ES&S	52	8	0	29	31	0	54 ms
Hart (eScan)	60	0	0	38	22	0	46 ms
Hart (Verity)	60	0	0	27	33	0	21 ms

Table 1. Performance of UnclearBallot— We tested how accurately our software could manipulate voter marks for a variety of ballot styles using equal numbers of invalid and valid marks. The table shows how often the system skipped a mark, successfully altered one, or erroneously created artifacts we deemed to be visible upon manual inspection. We also report the mean processing time for successfully manipulated races, excluding template matching.

prepared 720 marked contests, split evenly among the six ballot styles shown in Figure 7. For each style, we marked 60 contests with what Bajcsy [2] calls “Filled” marks, i.e. reliably detected marks that should be moved by our attack. We marked another 60 ballots in each ballot style with marginal marks, ten each for the five kinds of marginal marks shown in Figure 2 and ten empty marks.

Because the runtime of the template matching step of our algorithm is highly dependent on customization for the particular races on a ballot, we opted to skip it for this experiment. Rather than marking full ballots, we marked cropped races from each ballot style and then ran them through our program. We then manually checked to ensure that the races the program moved were not detectable by inspection. Results for these experiments are shown in Table 1.

Despite rejecting some valid ballots, our program is still able to confidently swap a majority of valid votes. In a real attack, only a small percentage of votes would need to actually be modified, a task easily accomplished by our program. Our program also correctly catches all votes that we have deemed invalid for swapping. This would make it unlikely to be detected in an image audit.

Dominion ballots saw a much higher rate of invalid mark moving, and Diebold and Dominion ballots saw a much higher rate of valid mark moving. This is likely due to the placement of targets: on the Dominion ballots, the mark is right justified, separating it significantly from candidate label information, as can be seen in Figure 7. Similarly, the Diebold ballot provides more space around the target and less candidate information that can be intercepted by marks, which would cause Unclear Ballot to skip moving the mark.

In an online attack scenario (such as if a human is waiting to see the output from the scanner), the attacker needs to be able to modify ballot scans quickly enough not to be noticed. Factors which might affect how quickly our program can process and manipulate ballots include ballot style, layout, and type of mark. During the accuracy experiment just described, we collected timing data for

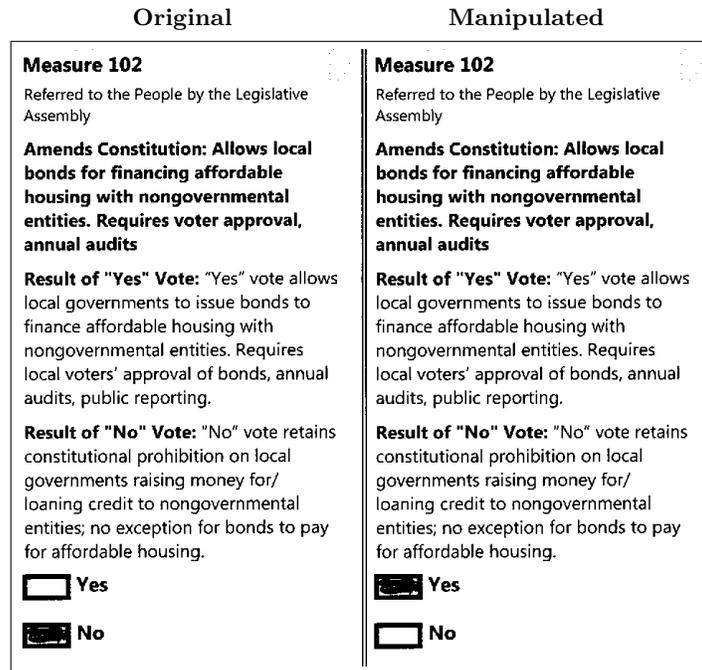


Fig. 8. Attacking Real Ballots—Using 181,541 images of voted ballots from Clackamas County, Oregon, we attempted to change voters' selections for the ballot measure shown above. UnclearBallot determined that it could safely alter 34% of the ballots. For reference, Measure 102 passed by a margin of 5%, well within range of manipulation [14]. We inspected 1,000 of them to verify that the manipulation left no obvious artifacts.

successfully manipulated ballot, and report the results in Table 1. The results show that after the target race has been extracted, the algorithm completes extremely quickly for all tested ballot styles. We present additional timing data at the end of the following section.

5.2 Testing with Real Voted Ballots

To assess the effectiveness of UnclearBallot in a real election, we used a corpus of scans of 181,541 real ballots from the November 6, 2018, General Election in Clackamas County, Oregon, which were made available by Election Integrity Oregon [18]. Like all of Oregon, Clackamas County uses vote-by-mail as its primary voting method, and votes are centrally counted using optical scanners. All images were Hart Verity-style ballots, as shown in Figure 7.

We selected a ballot measure that appeared on all the ballots (Figure 8) and attempted to change each voter's selection. UnclearBallot rejected 20,117 (11%) of the ballots because it could not locate the target contest. We examined a subset of the rejected ballots and found that they contained glitches introduced

during scanning (such as vertical lines running the length of the ballot), which interfered with the Hough transform.

To simulate a real attacker, we configured UnclearBallot with conservative parameters, so that it would only modify marks when there was high confidence that the alteration would not be noticeable. As a result, it would only manipulate marks that were nearly perfectly filled in. In most cases, marks that were skipped extended well beyond the target, but the program also skipped undervotes, overvotes, or mislabeled scans. Under these parameters, the program altered the target contest in 62,400 (34%) of the ballot images.

Two authors independently inspected a random sample of 1,000 altered ballots to check whether any contained artifacts that would be noticeable to an attentive observer. Such artifacts might include marks which were unnaturally cut off, visible discontinuities in pixel darkness (i.e. dark lines around moved marks), and so on. If these artifacts were seen during an audit, officials might recheck all of the physical ballots and reverse the effects of the attack. None of the altered ballots we inspected contained noticeable evidence of manipulation.

We also collected timing data while processing Clackamas County ballots. Running on a system with a 4-core Intel E3-1230 CPU running at 3.40 GHz with 64 GB of RAM, UnclearBallot took an average of 279 ms to process each ballot. For reference, Hart’s fastest central scanner’s maximum scan rate is one ballot per 352 ms [37], well above the time needed to carry out our attack.

These results show that UnclearBallot can successfully and efficiently manipulate ballot images to change real voters’ marks. Moreover, the alterations likely would be undetectable to human auditors who examined only the ballot images.

6 Discussion and Mitigations

UnclearBallot demonstrates the need for a software-independent evidence trail against which election results can be checked. It shows that audits based on software which is independent from the rest of the election system is still not software independent. To date, the only robust and secure election technology that is widely used is optical-scan paper ballots with risk-limiting audits based on a robust, well-maintained, *physical* audit trail. However, image audits are not useless, and here we discuss uses for them as well as potential mitigations for our attack.

Uses for image audits. So long as image audits are not the sole mechanism for verifying election results, they do provide substantial benefits to election officials. Using an image audit vastly simplifies some functions of election administration, like ballot adjudication in cases where marks cannot be interpreted by scanners or are otherwise ambiguous. Image audits can be used to efficiently identify and document election discrepancies, as has occurred in Maryland where nearly 2,000 ballots were discovered missing from the audit trail in 2016 [28]. Image audits also identified a flaw in the ES&S DS850 high speed scanner, where it was causing some ballots to stick together and feed two at a time [29].

Another way to utilize image audits is a transitive audit. Methods like SOBA [8] seek to construct an audit trail using all available means of election evidence, rooting the audit in some verification of physical record. By using physical records to verify other records, like CVRs or ballot images, confidence in election outcomes can be transitively passed on to non-physical audit trails. The drawback with this kind of audit is that it usually requires the same level of work as an RLA, plus whatever work is needed to validate the other forms of evidence. However, since ballot image audits already require a low amount of effort, they may augment RLAs and provide better transparency into the auditing process.

Image audits are an augmentation and a convenience for election administration, however, and should not be viewed as a security tool. Only physical examination of paper ballots, as in a risk-limiting audit, can provide a necessary level of mitigation to manipulated election results.

End-to-end (E2E) systems. Voting systems with rigorous integrity properties and tamper resistance such as Scantegrity [12] and Prêt à Voter [35] provide a defense to UnclearBallot. In Scantegrity, when individuals mark their ballots, a confirmation code is revealed that is tied to the selected candidate. This enables a voter to verify that their ballot collected-as-cast and counted-as-collected, as they can look up their ballot on a public bulletin board. Since each mark reveals a unique code, moving the mark would match the code with the wrong candidate, so voters would be unable to verify their ballots. If enough voters complain, this might result in our attack being detected.

Prêt à Voter randomizes the candidate order on each ballot, which creates a slightly higher barrier for our attack, as an additional template matching step would be needed to ascertain candidate order. More importantly, the candidate list is physically separated from the voter’s marks upon casting the ballot, so malware which could not keep track of the correct candidate order could not successfully move marks to a predetermined candidate. Since the candidate order is deciphered via a key-sharing scheme, malicious software would have to infect a significant portion of the election system and act in a highly coordinated way to reconstruct candidate ordering. Moreover, as with Scantegrity, votes are published to a public bulletin board, so any voter could discover if their vote had not been correctly recorded.

Other E2E systems which make use of optical scanning and a bulletin board, like STAR-Vote [6], Scratch and Vote [1], and VeriScan [7], are similarly protected from attacks like UnclearBallot.

Other mitigations. Outside of E2E, there may be other heuristic mitigations that can be easily implemented even in deployed voting systems to make our attack somewhat more difficult. As mentioned above, randomizing candidate order on each ballot increases the computation required to perform our attack. Voters drawing outside the bubbles can also defeat our attack, though this might also result in their votes not counting and may be circumvented by replacing the whole race on the ballot image with a substituted one. Collecting ballot images

from a different source than the tabulator makes our attack more difficult, as votes now have to be changed in two places. Other standard computer security technologies, like secure file systems, could be used to force the attacker to alter ballot images in a way that also circumvents protections like encryption and permissions.

Detection. Technologies that detect image manipulation may also provide some mitigation. Techniques like those discussed in [3–5, 38], among others, could be adapted to try to automatically detect moved marks on ballots. However, as noted by Farid [19], image manipulation detection is a kind of arms race: given a fixed detection algorithm, adversaries can very likely find a way to defeat it. In our context, an attacker with sufficient access to the voting system to implant a manipulation algorithm would likely also be able to steal the detector code. The attacker could improve the manipulation algorithm or simply use the detector as part of their mark-moving calculus: if moving a mark will trip the detector, an attacker can simply opt not to move the mark.

While a fixed and automatic procedure for detecting manipulation can provide little assurance, it remains possible that an adaptive approach to detection could be a useful part of a post-election forensics investigation. However, staying one step ahead of sophisticated adversaries would require an ongoing research program to advance the state of the art in detection methods.

A less costly and more dependable way to detect ballot manipulation detection would be to use a software independent audit trail to confirm election outcomes. This can be accomplished with risk-limiting audits, and the software independence enabled by RLAs provides other robust security properties to elections, including defending against other potential attacks on tabulation equipment and servers.

Future work. We have only focused on simple-majority elections here, because those are the kinds of elections used by jurisdictions that do image audits. Audits of more complex election methods, like instant-runoff voting or D’Hondt, have been examined to some extent [36, 41], but future work is needed into audits of these kinds of elections altogether. Because the marks made in these elections are different than the kind we’ve discussed here, manipulating these ballot images may not be able to employ the same image processing techniques we have used. Additionally it may be difficult for malware to know how many marks it needs to move, since margins in complex elections are difficult to compute. We leave exploration of image manipulation of these elections to future work.

7 Conclusion

In this paper, we demonstrated an attack that defeats ballot image audits of the type performed in some jurisdictions. We presented an implementation using a real scanner, and evaluated our implementation against a set of real ballots and a set of systematically marked ballots from a variety of ballot styles. Our

attack shows that image audits cannot be relied upon to verify that elections are free from computer-based interference. Indeed, the only currently known way to verify an election outcome is with direct examination of physical ballots.

Acknowledgements

The authors thank Vaibhav Bafna and Jonathan Yan for assisting in the initial version of this project. They also thank Josh Franklin, Joe Hall, Maurice Turner, Kevin Skoglund, Jared Marcotte, and Tony Adams for their invaluable feedback. We also thank our anonymous reviewers and our shepherd, Roland Wen. This material is based upon work supported by the National Science Foundation under grants CNS-1518888.

References

1. Adida, B., Rivest, R.L.: Scratch and Vote: Self-contained paper-based cryptographic voting. In: ACM Workshop on Privacy in the Electronic Society. pp. 29–40 (2006)
2. Bajcsy, A., Li-Baboud, Y.S., Brady, M.: Systematic measurement of marginal mark types on voting ballots. Tech. rep., National Institute for Standards and Technology (2015)
3. Bayar, B., Stamm, M.C.: A deep learning approach to universal image manipulation detection using a new convolutional layer. In: Proceedings of the 4th ACM Workshop on Information Hiding and Multimedia Security. pp. 5–10. ACM (2016)
4. Bayram, S., Avcibas, I., Sankur, B., Memon, N.: Image manipulation detection with binary similarity measures. In: 2005 13th European Signal Processing Conference. pp. 1–4. IEEE (2005)
5. Bayram, S., Avcibas, I., Sankur, B., Memon, N.D.: Image manipulation detection. *Journal of Electronic Imaging* **15**(4), 041102 (2006)
6. Bell, S., Benaloh, J., Byrne, M.D., DeBeauvoir, D., Eakin, B., Fisher, G., Kortum, P., McBurnett, N., Montoya, J., Parker, M., Pereira, O., Stark, P.B., Wallach, D.S., Winn, M.: STAR-vote: A secure, transparent, auditable, and reliable voting system. *USENIX Journal of Election Technology and Systems* **1**(1) (Aug 2013)
7. Benaloh, J.: Administrative and public verifiability: Can we have both? In: USENIX/ACCURATE Electronic Voting Technology Workshop. EVT '08 (Aug 2008)
8. Benaloh, J., Jones, D., Lazarus, E., Lindeman, M., Stark, P.B.: SOBA: Secrecy-preserving observable ballot-level audit. In: proc. Proc. USENIXAccurate Electronic Voting Technology Workshop (2011)
9. Bernhard, M., Benaloh, J., Halderman, J.A., Rivest, R.L., Ryan, P.Y., Stark, P.B., Teague, V., Vora, P.L., Wallach, D.S.: Public evidence from secret ballots. In: International Joint Conference on Electronic Voting. pp. 84–109. Springer (2017)
10. Bowen, D.: Top-to-Bottom Review of voting machines certified for use in California. Tech. rep., California Secretary of State (2007), <https://www.sos.ca.gov/elections/voting-systems/oversight/top-bottom-review/>
11. Bradski, G.: The OpenCV Library. *Dr. Dobb's Journal of Software Tools* (2000)
12. Carback, R., Chaum, D., Clark, J., Conway, J., Essex, A., Herrnson, P.S., Mayberry, T., Popoveniuc, S., Rivest, R.L., Shen, E., Sherman, A.T., Vora, P.L.: Scantegrity II municipal election at Takoma Park: The first E2E binding governmental election with ballot privacy. In: 18th USENIX Security Symposium (Aug 2010)

13. Chung, K.K.t., Dong, V.J., Shi, X.: Electronic voting method for optically scanned ballot (Jul 18 2006), US Patent 7,077,313
14. November 6, 2018 general election. <https://dochub.clackamas.us/documents/drupal/f4e7f0fb-250a-4992-918d-26c5f726de3c>
15. Clear Ballot: ClearAudit, <https://clearballot.com/products/clear-audit>
16. Dominion Voting: Auditmark. <https://www.dominionvoting.com/pdf/DD%20Digital%20Ballot%20AuditMark.pdf>
17. Dominion Voting: Cambridge Case Study. <https://www.dominionvoting.com/field/cambridge>
18. Election Integrity Oregon, <https://www.electionintegrityoregon.org>
19. Farid, H.: Digital forensics in a post-truth age. *Forensic science international* **289**, 268–269 (2018)
20. Feldman, A.J., Halderman, J.A., Felten, E.W.: Security analysis of the Diebold AccuVote-TS voting machine. In: USENIX/ACCURATE Electronic Voting Technology Workshop. EVT '07 (Aug 2007)
21. Hall, J., Miratrix, L., Stark, P., Briones, M., Ginnold, E., Oakley, F., Peaden, M., Pellerin, G., Stanionis, T., Webber, T.: Implementing risk-limiting post-election audits in California. In: 2009 Workshop on Electronic Voting Technology/Workshop on Trustworthy Elections. pp. 19–19. USENIX Association (2009)
22. Ji, T., Kim, E., Srikantan, R., Tsai, A., Cordero, A., Wagner, D.A.: An analysis of write-in marks on optical scan ballots. In: EVT/WOTE (2011)
23. Jones, D.W.: On optical mark-sense scanning. In: *Towards Trustworthy Elections*, pp. 175–190. Springer (2010)
24. Lindeman, M., Halvorson, M., Smith, P., Garland, L., Addona, V., McCrea, D.: Principles and best practices for post-election audits (Sep 2008), <http://electionaudits.org/files/bestpracticesfinal.0.pdf>
25. Lindeman, M., Stark, P.: A gentle introduction to risk-limiting audits. *IEEE Security and Privacy* **10**, 42–49 (2012)
26. Lindeman, M., Stark, P., Yates, V.: BRAVO: Ballot-polling risk-limiting audits to verify outcomes. In: 2011 Electronic Voting Technology Workshop / Workshop on Trustworthy Elections (EVT/WOTE '12). USENIX (2012)
27. Maryland House of Delegates: House Bill 1278: An act concerning election law – postelection tabulation audit. <http://mgaleg.maryland.gov/2018RS/bills/hb/hb1278E.pdf>
28. Maryland State Board of Elections: 2016 post-election audit report. http://dlslibrary.state.md.us/publications/JCR/2016/2016_22-23.pdf (12 2016)
29. Maryland State Board of Elections: December 15, 2016 meeting minutes. https://elections.maryland.gov/pdf/minutes/2016_12.pdf (Dec 2016)
30. McDaniel, P., Blaze, M., Vigna, G.: EVEREST: Evaluation and validation of election-related equipment, standards and testing. Tech. rep., Ohio Secretary of State (2007), <http://siis.cse.psu.edu/everest.html>
31. Mebane, W., Bernhard, M.: Voting technologies, recount methods and votes in Wisconsin and Michigan in 2016. 3rd Workshop on Advances in Secure Electronic Voting 2018 (Voting '18) **3** (2018)
32. National Academies of Sciences, Engineering, and Medicine: *Securing the Vote: Protecting American Democracy*. The National Academies Press, Washington, DC (2018), <https://www.nap.edu/catalog/25120/securing-the-vote-protecting-american-democracy>
33. National Conference of State Legislatures: Post-election audits (January 2019), <http://www.ncsl.org/research/elections-and-campaigns/post-election-audits635926066.aspx>

34. Rivest, R.: On the notion of ‘software independence’ in voting systems. *Phil. Trans. R. Soc. A* **366**(1881), 3759–3767 (October 2008)
35. Ryan, P.Y.A., Bismark, D., Heather, J., Schneider, S., Xia, Z.: Prêt à Voter: A voter-verifiable voting system. *IEEE Transactions on Information Forensics and Security* **4**(4), 662–673 (2009)
36. Sarwate, A.D., Checkoway, S., Shacham, H.: Risk-limiting audits and the margin of victory in nonplurality elections. *Statistics, Politics and Policy* **4**(1), 29–64 (2013)
37. ScannerOne: Kodak i5600. <http://www.scannerone.com/product/KOD-i5600.html>
38. Stamm, M.C., Liu, K.R.: Forensic detection of image manipulation using statistical intrinsic fingerprints. *IEEE Transactions on Information Forensics and Security* **5**(3), 492–506 (2010)
39. Stark, P.: Conservative statistical post-election audits. *Ann. Appl. Stat.* **2**(2), 550–581 (2008)
40. Stark, P.: Super-simple simultaneous single-ballot risk-limiting audits. In: 2010 Electronic Voting Technology Workshop / Workshop on Trustworthy Elections (EVT/WOTE ’10). USENIX (2010)
41. Stark, P.B., Teague, V., Essex, A.: Verifiable European elections: Risk-limiting audits for D’Hondt and its relatives. *{USENIX} Journal of Election Technology and Systems ({JETS})* **1**, 18–39 (2014)
42. Unisyn Voting Solutions: OpenElect OCS Auditor. <https://unisynvoting.com/openelect-ocs/>
43. U.S. Election Assistance Commission: Certificate of conformance: ClearVote 1.5. <https://www.eac.gov/file.aspx?A=zgte4IhsHz%2bswC%2bW4LO6PxIVssxXBebhvZiSd5BGbbs%3d> (3 2019)
44. Verified Voting Foundation: The Verifier: Polling place equipment (2019), <https://www.verifiedvoting.org/verifier/>

Election Manipulation with Partial Information

Michelle Blom¹, Peter J. Stuckey², and Vanessa J. Teague¹

¹[michelle.blom,vjteague]@unimelb.edu.au
School of Computing and Information Systems
The University of Melbourne

²peter.stuckey@monash.edu
Faculty of Information Technology
Monash University

Abstract. We consider the case of manipulating the results of Instant Runoff Voting (IRV) elections. Previous work in this area looked at *posthoc* manipulation with complete information, where the manipulator may alter ballots after reading the whole election profile. In this paper we examine the much more realistic, but challenging, problem of manipulating ballots during the election process, having observed only some ballots. The aim of the manipulator is to modify as few ballots as possible to ensure their candidate’s victory with high probability. We show that this is quite feasible in practice to generate efficient manipulations with a high probability of success. We also add some extra conditions on the manipulations to it less likely they will be detected by naive methods.

1 Introduction

Instant Runoff Voting (IRV), also known as Alternative Vote (AV), is a system of preferential voting in which voters rank candidates in order of preference. Tallying proceeds by eliminating the least popular candidate and redistributing their votes to the next-preferred candidate—see Section 2 for a precise algorithm. It is used in presidential, parliamentary and local government elections in countries such as Australia, Fiji, Papua New Guinea, Ireland, Bosnia/Herzegovina, the UK and United States [6].

IRV elections can be complicated—an early alteration in the elimination order can cascade into an entirely different election result. It is NP-hard even to compute the true margin of victory [7]. But the hard examples are somewhat contrived, and real elections follow patterns which make it relatively easy in practice to compute both margins [5, 3] and successful manipulations [2]. However, prior work has required complete information about all the votes cast in the election. In contrast, the related problem of manipulating one’s own vote to increase the likelihood of a desired outcome has been studied in the context of both full and partial knowledge of the preferences of other voters [4].

In this work, we examine for the first time whether an attacker with only partial information can devise a successful manipulation with a high probability of success. This models a realistic attacker (such as a corrupt scanner or voting machine) who must generate manipulated electronic records on the fly after reading only a few of them. The problem we consider is known as ‘election control by replacing voters’ in the computational social choice literature. We find that it remains easy to generate successful

manipulations, though of course success is not guaranteed. Our emphasis is on practical attacks in a realistic setting. We do not prove optimality nor a bound on the probability of success, but we demonstrate successful practical attacks by simulating the algorithm on state election data from Australia. We also examine a number of ways the attacker could refine the manipulation to be less likely to raise suspicion.

We assume the adversary has sufficient access to the election system to read ballots as they are received, and to modify each ballot before it is recorded. Unlike in prior work [2], our adversary cannot modify ballot recordings once made, but must perform manipulations based on the partial sample seen so far. This models a man-in-the-middle attack where ballots are intercepted between the voter and the final counting, and altered by the adversary. It is precisely the attacker model of a corrupted scanning process in which scanned ballots must be immediately output, but may be misrecorded on the fly. It is also appropriate for an e-voting process with immutable logs, in which the manipulated votes must be written immediately onto the log and cannot later be altered.

We look at an adversary that computes some manipulation rules to apply at periodic intervals, after they have seen some ballots scanned through. The attacker succeeds if their desired candidate wins. We compare a number of ways of generating the rules, and how effective they are at actually achieving a desired winner change. We find that if the adversary sees the first quarter or half of the votes then, with a few seconds of computation, it can compute a manipulation close to the optimal, and then apply it to the next 3/4 or half of the votes with a success probability usually higher than 90%.

Avoiding suspicion The last round margin (LRM) of an IRV election – the difference in tallies of the final two remaining candidates, divided by two and rounded up – is commonly used as an indicator of how close the election was. Blom *et al.* [3] have shown that the true margin of victory (MOV) of the election – the smallest number of votes one would have to alter to change who won the election – is generally equal to the last round margin, but not always. In some cases, the MOV can be much smaller than the last round margin. The Australian Electoral Commission (AEC) use the “margin between the two leading candidates” after all remaining candidates have been eliminated, and their preferences distributed, to determine whether an automatic recount of cast votes should be automatically performed.¹ In practice, the LRM plays a major role in determining whether further scrutiny of the outcome is performed.

Our first suspicion-avoiding extension is for the adversary to generate manipulations that result in a large last-round margin. The adversary wants to alter the smallest number of electronic records so that their desire of changing the outcome is realised, while at the same time ensuring that the last round margin of the manipulated election is larger than a given threshold. Another important way of avoiding suspicion is to minimize changes in first preferences. This is significant because, in preferential elections, first-preference tallies are often manually counted in a polling place independently of the scanning process. Modifying first preferences is therefore much more likely to be detected than other alterations. We find, surprisingly, that this constraint usually makes an insignificant difference to the number of ballots that need to be changed. A clever

¹ <https://www.aec.gov.au/Elections/candidates/files/hor-recount-policy.pdf>. The AEC definition of a “margin” is the difference in tallies of two candidates (not divided by two).

Initially, all candidates remain standing (are not eliminated)
While there is *more than one* candidate standing
 For every candidate c standing
 Tally (count) the votes in which c is the highest-ranked
 candidate of those standing
 Eliminate the candidate with the smallest tally
The winner is the one candidate not eliminated

Fig. 1: The IRV vote tallying process: the candidate with the smallest tally is repeatedly eliminated, with the ballots in their tally redistributed according to their next preference.

surreptitious manipulation is almost as easy as an obvious one. These attacks would be defeated if a rigorous risk-limiting audit was applied. The trick with LRMs would be defeated by a careful computation of the true MOV [3]. The purpose of this paper is to demonstrate that these rigorous methods are necessary, because an active adversary can defeat heuristic methods of assessing when to re-examine an election result.

Our contributions Using the Australian New South Wales (NSW) 2015 Legislative Assembly election as a case study, we simulate the attack by randomly determining the order in which ballots arrive at the scanner, computing a best-guess manipulation at periodic intervals, and then applying that manipulation to later ballots as they are scanned. We then compare the outcome of the modified election to that of the unmodified election to see whether the manipulation succeeded.

We report the number of ballots that this adversary needs to modify, and the number of first preferences on ballots they need to modify, in order to have a high chance of altering the election outcome, assuming that the proportion of the ballots they saw before computing the manipulation were a random sample of all ballots.

2 Preliminaries and prior work

Votes are tallied in an IRV election in a series of rounds (see Figure 1). In each round, the candidate with the smallest number of votes (their tally) is eliminated, with the last remaining candidate declared the winner. All votes in an eliminated candidate’s tally are distributed to the next most-preferred (remaining) candidate in their ranking.

We denote an IRV election, with a set of candidates \mathcal{C} , as \mathcal{B} . A sequence of candidates π is represented in list notation (e.g., $\pi = [c_1, c_2, c_3, c_4]$, which means that c_1 is the highest preference, c_2 the next-preferred, and so on). Such sequences represent both votes and the order in which candidates are eliminated. An election \mathcal{B} is defined as a multiset² of votes, each vote $b \in \mathcal{B}$ a sequence of candidates in \mathcal{C} , with no duplicates, listed in order of preference (most preferred to least). The first candidate appearing in a sequence π is denoted $first(\pi)$ (e.g., $first([c_2, c_3]) = c_2$). In each round of vote counting, there are a current set of eliminated candidates \mathcal{E} and a current set of candidates still standing $\mathcal{S} = \mathcal{C} \setminus \mathcal{E}$. The winner c_w of the election is the last standing candidate.

² A multiset allows for the inclusion of duplicate items.

Our definitions for a candidate's tally, the margin of victory (MOV), and last round margin (LRM) of an IRV election are replicated from [2].

Definition 1. Tally $t_{\mathcal{S}}(c)$ Given candidates $\mathcal{S} \subseteq \mathcal{C}$ are still standing in an election \mathcal{B} , the tally for candidate $c \in \mathcal{C}$, denoted $t_{\mathcal{S}}(c)$, is defined as the number of votes $b \in \mathcal{B}$ for which c is the most-preferred candidate of those remaining.³ Let $p_{\mathcal{S}}(b)$ denote the sequence of candidates mentioned in b that are also in \mathcal{S} .

$$t_{\mathcal{S}}(c) = |\{b \in \mathcal{B} \mid c = \text{first}(p_{\mathcal{S}}(b))\}| \quad (1)$$

Definition 2. Margin of Victory (MOV) The MOV in an election \mathcal{B} with candidates \mathcal{C} and winner $c_w \in \mathcal{C}$, is the smallest number of votes in \mathcal{B} whose ranking must be modified (by an adversary) so that a candidate $c' \in \mathcal{C} \setminus \{c_w\}$ is elected.

Definition 3. Last Round Margin (LRM) The LRM of an election \mathcal{B} , in which two candidates $\mathcal{S} = \{c, c'\}$ remain with $t_{\mathcal{S}}(c)$ and $t_{\mathcal{S}}(c')$ votes in their tallies, is equal to half the difference between the tallies of c and c' rounded up.

$$LRM = \left\lceil \frac{|t_{\mathcal{S}}(c) - t_{\mathcal{S}}(c')|}{2} \right\rceil \quad (2)$$

In the design of our adversary, we consider the concept of an elimination margin (EM) – the margin by which a candidate is eliminated in a given round. We find that a manipulator can be much more effective when controlling the elimination margin of the two runner-ups (the candidate eliminated just prior to the runner-up, and the runner-up themselves) rather than just the LRM of the runner-up and winner. In settings where there is a genuine three candidate race for winner, successful manipulations typically require the elimination of a specific candidate in third place whose preferences will mostly flow to the desired winner. Consider a manipulation designed to elect a candidate a with a certain margin, but where tallies of the candidates up for elimination in that crucial third place position are similar. Given uncertainty in what types of ballots a manipulator will see as it is changing votes on the fly, it is likely to be more successful it aims to eliminate the required candidate in third place with a larger elimination margin.

Definition 4. Elimination Margin (EM) The EM in a round of counting i , in which candidate c_i is eliminated and candidates $\mathcal{S} \setminus \{c_i\}$ are still standing, is equal to half the smallest difference in tallies between $t_{\mathcal{S}}(c_i)$ and $t_{\mathcal{S}}(c')$ for $c' \in \mathcal{S} \setminus \{c_i\}$ rounded up.

$$EM_i = \min_{c' \in \mathcal{S} \setminus \{c_i\}} \left\lceil \frac{t_{\mathcal{S}}(c') - t_{\mathcal{S}}(c_i)}{2} \right\rceil \quad (3)$$

Example 1. Consider the election with ballots shown in Table 1(a). The true election result is shown in Table 1(b), with c the winner with a last round margin of 13. The margin of victory is actually only 5. Changing five ballots from $[c, a]$ to $[b, c]$ results the totals shown in Table 1(e). In this election a wins as shown in Table 1(f). In the true election, b is eliminated with an EM of $\lceil 9/2 \rceil = 5$ votes, and a with an EM of 13. \square

³ Square brackets have been used to denote a multiset.

Table 1: IRV example, with (a) the number of votes cast with each listed ranking over candidates a, b, c , and (b) tallies after each round of vote counting (c) the ballot count 25% into the election and the estimated complete election (d) the tallies of each round of vote counting in the estimated election (e) the number of votes recorded after manipulation, and (f) the tallies after each round of vote counting in the manipulated election

Ranking	Count \mathcal{B}
[a]	55
[c, a]	30
[b, c]	36
[c]	15

(a)

Candidate	Round 1	Round 2
a	55	55
b	36	—
c	45	81

(b)

Ranking	Partial Count $\mathcal{B}_{\mathcal{T}}$	Estimated Completion $\hat{\mathcal{B}}$
[a]	11	44
[c, a]	9	36
[b, c]	10	40
[c]	4	16

(c)

Candidate	Round 1	Round 2
a	44	44
b	36	—
c	52	88

(d)

Ranking	Manipulated Count
[a]	55
[c, a]	25
[b, c]	41
[c]	15

(e)

Candidate	Round 1	Round 2
a	55	80
b	41	41
c	40	—

(f)

2.1 Modifying an Election Outcome

In this paper we make use of algorithms for determining minimal manipulations of IRV elections in order to change their outcome, developed by Blom *et al* [2]. They consider the case where all recorded ballots can be manipulated *after* they have all been scanned. This is clearly a necessary starting point for determining a more restricted manipulation.

These algorithms are based on the *margin-irv* algorithm for computing the true MOV of an IRV election. A description of *margin-irv* can be found in Blom *et al.* [3]. We summarise the algorithm in this section, and explain how it can be modified to compute the smallest number of vote changes required to: (i) bring about a change in the outcome of the election; (ii) produce a manipulated election with certain properties, modelled as side constraints; and (iii) produce a manipulated election that minimises discrepancies between a manual hand count of first preference tallies (based on non-manipulated paper ballots) and the first preference tallies of the manipulated election.

The *margin-irv* algorithm Consider an IRV election \mathcal{B} with candidates \mathcal{C} and winner $w \in \mathcal{C}$. The *margin-irv* algorithm starts by adding $|\mathcal{C}| - 1$ partial elimination sequences to a search tree, one for each of alternate or desired winner $c \in \mathcal{C} \setminus \{w\}$. These partial sequences form a frontier F , with each sequence containing a single candidate – an alternate winner. Note that a partial sequence [a, b, c] represents an election outcome

in which a and b are the last two candidates eliminated, and c the winner. All other candidates are assumed to have been eliminated in some prior round.

An adversary is likely to have a desired winner $a \neq w$. We can modify *margin-irv* by initializing the frontier F with the single sequence $[a]$. Then it only considers election sequences where a wins. The same idea is used by Blom *et al.* [1].

For each partial sequence $\pi \in F$, we compute a lower bound on the number of vote changes required to realise an elimination sequence that *ends* in π . These lower bounds are used to guide construction of the search tree, and are computed by both solving an Integer Linear Program (ILP), and applying several rules for lower bound computation. These rules are described in Blom *et al.* [3]. The ILP, denoted `DISTANCETO`, computes a lower bound on the smallest number of vote changes required to transform the election \mathcal{B} , with an elimination sequence π' , to one with an elimination sequence that ends in π . When applied to a complete order π , containing all candidates, `DISTANCETO` exactly computes the smallest number of votes changes required to realise the outcome π . The largest of the lower bounds computed by the rules of Blom *et al.* [3] and the `DISTANCETO` ILP is assigned to each partial sequence π as it is added to F . The `DISTANCETO` ILP is defined in Section 2.2. To enforce additional constraints on the nature of any manipulated election, we add these constraints to each ILP solved.

The partial sequence $\pi \in F$ with the smallest assigned lower bound is selected and *expanded*. For each $c \in \mathcal{C}$ that is not already present in π , we create a new sequence with c appended to the front. For example, given a set of candidates e, f , and g , with winning candidate g , the partial sequence $\pi = [f]$ will be expanded to create two new sequences $[e, f]$ and $[g, f]$. We evaluate each new sequence π' by assigning it a lower bound on the number of votes required to realise any elimination order ending in π' .

While exploring and building elimination sequences, *margin-irv* maintains a running *upper bound* on the value of the true margin. Without any side constraints designed to inject desirable properties into a manipulated election, this upper bound is initialised to the last round margin of the original election. To enforce additional constraints on the properties of any manipulated election, we need to manipulate at least as many, and often more, votes than required to simply change the original outcome. Consequently, we must set the upper bound maintained by *margin-irv* to a higher value. In this context, we set the initial upper bound to the total number of votes cast in the election. This is clearly always a correct upper bound on any manipulation.

When a sequence π containing all candidates is constructed, the `DISTANCETO` ILP computes the exact number of vote manipulations required to realise it, while satisfying all desired side constraints. If this number is lower than our current upper bound, the upper bound is revised, and all orders in F with a lower bound greater than or equal to it are pruned from consideration (removed from F). This process continues until F is empty (we have considered or pruned all possible alternate elimination sequences). The final value of the running upper bound is the true electoral MOV (with side constraints).

2.2 `DISTANCETO` with Side Constraints

We now present the `DISTANCETO` Integer Linear Program (ILP) used to compute lower bounds on the degree of manipulation required to realise an election outcome ending in a given candidate sequence, and the (exact) smallest number of vote changes required

to realise a given (complete) alternate elimination sequence. This ILP, without added side constraints, was originally presented by Magrino *et al.* [5].

We consider additional side constraints to inject desirable properties into any manipulated election, and describe how we can minimise both the total number of vote changes required to elect a desired winner, and the total changes made to first preference tallies between the true and manipulated election profiles.

Let \mathbf{R} denote the set of possible (partial and total) rankings R of candidates \mathcal{C} that could appear on a vote, N_R the number of votes cast with ranking $R \in \mathbf{R}$, and N the total number of votes cast. Let $\mathcal{R}_{j,i}$ denote the subset of rankings ($\mathcal{R}_{j,i} \subset \mathbf{R}$) in which c_j is the most preferred candidate still standing (i.e., that will count toward c_j 's tally) at the start of round i (in which c_i is eliminated). For each $R \in \mathbf{R}$, we define variables:

- q_R integer number of votes to be changed into R ;
- m_R integer number of votes with ranking R in the unmodified election to be changed into something other than R ; and
- y_R number of votes in the modified election with ranking R .

Given a partial or complete order π , the DISTANCETO ILP is:

$$\min \sum_{R \in \mathbf{R}} q_R \quad (4)$$

$$N_R + q_R - m_R = y_R \quad \forall R \in \mathbf{R} \quad (5)$$

$$\sum_{R \in \mathbf{R}} q_R = \sum_{R \in \mathbf{R}} m_R \quad (6)$$

$$\sum_{R \in \mathcal{R}_{i,i}} y_R \leq \sum_{R \in \mathcal{R}_{j,i}} y_R \quad \forall c_i, c_j \in \pi . i < j \quad (7)$$

$$n \geq y_R \geq 0, \quad N_R \geq m_R \geq 0, \quad q_R \geq 0 \quad \forall R \in \mathbf{R} \quad (8)$$

Constraint (5) states that the number of votes with ranking $R \in \mathbf{R}$ in the new election is equal to the sum of those with this ranking in the unmodified election and those whose ranking has *changed to* R , minus the number of votes whose ranking has been *changed from* R . Constraint (7) defines a set of *special elimination constraints* which force the candidates in π to be eliminated in the stated order. Constraint (6) ensures that the total number of votes cast in the election does not change as a result of the manipulation. The objective minimises the total number of ballot changes required to manipulate the election and enforce a desired elimination sequence.

The above ILP does not include any additional side constraints – properties that we want the manipulated election to satisfy besides resulting in a different winner to that of the original election. Manipulated elections found by *margin-irv* in this setting are almost always evidently close, with a last round margin of 0 or 1 vote. This makes sense as the algorithm is trying to manipulate as few votes as possible, breaking any ties in favour of an alternate outcome. An adversary with the ability to modify electronic records of cast votes, however, will want to create a manipulated election that is not evidently close. An election with a tie in the final round of counting, or a difference of several votes in the tallies of the final two remaining candidates, is likely to be closely

scrutinised. Australian IRV elections in which the final tallies of the last two candidates differ by less than 100 votes, for example, trigger an automatic recount.

Given the widespread use of the last round margin as the indicator of how close an IRV election is, rather than the true MOV of the election, our adversary can use this to their advantage. Consider a candidate elimination sequence π , containing at least n candidates from a set \mathcal{C} . Let the last n candidates in the sequence π be denoted by $c_k, c_{k+1}, \dots, c_{k+n}$, with c_{k+n} denoting the winning candidate according to π . Adding the following side constraint to DISTANCETO ensures that margin by which each candidate c_i for $k \leq i \leq k+n-1$ is eliminated (the EM, Definition 4) is at least Δ votes. This allows us to ensure that both the last round margin of the manipulated election, and the elimination margin of any number of prior rounds, is at least a certain size.

$$\sum_{R \in \mathcal{R}_{i,i}} y_R \leq \sum_{R \in \mathcal{R}_{j,i}} y_R + 2\Delta \quad \forall i \in \{k, \dots, k+n-1\} \quad (9)$$

An important side constraint we will make use in the design of our manipulators is *limiting the manipulation*. Since in our scenario the attacker will modify ballots in the middle of the election its important that we calculate manipulations that do not remove ballots already recorded. Suppose C_R is the number of ballots already recorded for ranking R . We must ensure that the modified election has at least C_R ballots with ranking R , since they cannot be changed. Adding the following constraint ensures this.

$$y_R \geq C_R \quad \forall R \in \mathbf{R} \quad (10)$$

2.3 Minimising Change to First Preference Counts

To both minimise discrepancies between any manual count of first preference tallies, and that computed by counting software applied to a manipulated election, we solve the DISTANCETO ILP of Section 2.2 twice for each complete elimination sequence π .

For every such π that *margin-irv* encounters (these are the leaves of the generated search tree), we first solve the ILP with an objective to minimise first preference tally discrepancies. Let $t_{\mathcal{C}}(c)$ denote the first preference tallies of each candidate $c \in \mathcal{C}$. The first preference tally for candidate $c \in \mathcal{C}$ in any solution of our ILP, $t'_{\mathcal{C}}(c)$, is given by:

$$t'_{\mathcal{C}}(c) = \sum_{R \in \mathbf{R}.c = \text{first}(R)} y_R \quad (11)$$

Let FPD denote the total number of first preference tally discrepancies between any manipulated profile found by DISTANCETO ILP, and the true election profile of the original election (Equation 12). For each complete candidate elimination sequence π , containing all candidates, we first solve the DISTANCETO ILP with the objective shown in Equation 13. Let fpd denote the value of FPD in the optimal solution found when solving DISTANCETO with this objective. We then constrain the ILP to ensure that any subsequently found solutions must satisfy the constraint in Equation 14.

$$FPD = \sum_{c \in \mathcal{C}} |t_c(c) - t'_c(c)| \quad (12)$$

$$\min FPD \quad (13)$$

$$FPD \leq fpd \quad (14)$$

We re-solve our DISTANCETO ILP with the objective shown in Equation 4 – to minimise total ballot changes given the constraint limiting permitted change to first preference counts. The result is a manipulated election profile that first aims to minimise the total number of discrepancies between the true and manipulated elections, as a first priority, and the total number of ballot changes as a second.

3 Manipulation Algorithms

We consider a setting in which all ballots are assumed to be scanned at a central location, and the resulting electronic records passed into counting software that computes the result of the election. We assume the scanner has been compromised, and the attacker is able to manipulate (change) the ranking on the electronic record of ballots at the moment they are scanned. Alternatively, we can also consider a setting where votes are scanned at disparate locations, but the attacker is able to intercept and modify the scanned record before it is passed to the counting software. In both cases the attacker manipulates the record of ballots before they are used for counting. In either case, first preference counts may have been completed manually at polling booths and/or at the central scanning location. We assume the attacker has some reasonable estimate of the total number of ballots E expected in the election. Note that in places like Australia with compulsory voting the expected number of ballots E is known with high confidence, in other jurisdictions there will be more variance. Across all experiments in which we simulate our manipulators, we use $E = |\mathcal{B}|$ where \mathcal{B} is the historical set of ballots.

The algorithm applied by our manipulators is as follows. The parameter α , determining how often the attacker computes a set of manipulation rules, and k , the number of rounds on which to apply an elimination margin constraint, are given as input.

1. Let E be the expected number of ballots for the election. Let $n = \lceil \alpha E \rceil$ be a proportion α of this expected number of votes.
2. Collect the first n ballots $\mathcal{B}_{\mathcal{T}}$, passing them on to the counting software unmanipulated. Let $\mathcal{B}_{\mathcal{M}} = \mathcal{B}_{\mathcal{T}}$, the current profile of the manipulated election.
3. Compute an approximate complete election profile $\hat{\mathcal{B}}$ by extending the (manipulated) ballots $\mathcal{B}_{\mathcal{M}}$ processed so far with ballots uniformly drawn from the set of true (unmanipulated) ballots seen so far $\mathcal{B}_{\mathcal{T}}$.
4. Use the methods of [2] to determine a minimal manipulation \mathcal{M} of $\hat{\mathcal{B}}$ in order to achieve the desired winner with an EM of Δ applied to the last k rounds. Note that this manipulation may be null if the desired winner already win $\hat{\mathcal{B}}$ by Δ .
5. Examining the minimal manipulation made in \mathcal{M} and the assumed unseen ballots $\mathcal{U} = \hat{\mathcal{B}} - \mathcal{B}_{\mathcal{M}}$, determine a set of manipulation rules \mathcal{R} which will ensure that applying \mathcal{R} to \mathcal{U} will result in manipulation \mathcal{M} .

6. Intercept the next n ballots. If an incoming ballot b matches one of the manipulation rules, $r \in \mathcal{R}$, replace b by $r(b)$ before passing it on to the counting software.
7. Let $\mathcal{B}_{\mathcal{T}}$ be the true ballots seen so far. Let $\mathcal{B}_{\mathcal{M}}$ be the manipulated ballots processed so far – the current state of the manipulated election profile. If all ballots have been processed, the algorithm is complete, otherwise we return to step 3.

We now examine the individual steps in the approach in detail.

Completing an Election At any given point during the scanning of ballots, our adversary has seen a proportion δ of the total number of expected ballots, E . The manipulated election profile at this point contains δE ballots. Some of these ballots have been manipulated (altered from their true state), and others have been left unchanged. To compute a set of manipulation rules to apply to future ballots, the adversary needs to estimate what a complete election profile could look like (i.e., a profile containing the current set of manipulated ballots, and an estimate of future ballots). Let $\hat{\mathcal{B}}$ denote this estimated profile. To compute $\hat{\mathcal{B}}$, we start with the current set of ballots in $\mathcal{B}_{\mathcal{M}}$ and add $E - |\mathcal{B}_{\mathcal{M}}|$ further ballots. These additional ballots are drawn uniformly at random (with replacement) from $\mathcal{B}_{\mathcal{T}}$ – the set of ballots, unmodified, that have been seen so far. Each sampled ballot is added to $\hat{\mathcal{B}}$.

Example 2. Suppose the first quarter of ballots of the election in Table 1(a), $\mathcal{B}_{\mathcal{T}}$, is as shown in Table 1(c). Then an estimation of the complete election $\hat{\mathcal{B}}$ might be determined as shown in the same table. The result of the estimated election is shown in Table 1(d). Candidate c remains the winner with a LRM of 22. \square

Computing a Manipulation We consider two methods of computing a set of manipulation rules to apply to future scanned ballots, given an estimate of what the eventual election profile could look like, $\hat{\mathcal{B}}$, assuming the adversary makes no further changes. Each method first simulates the outcome of $\hat{\mathcal{B}}$ to determine whether any further manipulation is needed. If the desired candidate wins, no manipulation rules are generated. Otherwise, a set of rules indicating what kind of ballots to look for during the scanning process, and what to replace them with when they are seen, are formed. When a rule is followed by either of our manipulators, that rule is removed from their rule set.

Our first method for generating such rules uses *margin-irv*, as described in Section 2.1, to determine a minimal manipulation \mathcal{M} of the election $\hat{\mathcal{B}}$ so that the candidate selected by the attacker wins. Note that we add the side constraints of Eq (10) where $C_R = |\{r \mid r \in \mathcal{B}_{\mathcal{M}}, r = R\}|$ is the current count of ballots of the form R .

We then translate this minimal manipulation into a set of rules for the attacker to follow. Our second method does not compute a minimal manipulation of $\hat{\mathcal{B}}$, but simply computes the difference between final tallies of the eventual winner w , and the desired winner w' , $\Delta_{w,w'}$. The adversary will seek to remove $\lceil \Delta_{w,w'} / 2 \rceil$ votes in which w is preferred first, and replace them with a vote in which w' is preferred first.

Example 3. Imagine the attacker wants candidate a to win with a last round margin of 20. One such manipulation (certainly not the minimal one) is to change the ballots to sum to the counts in Table 1(e) which results in election shown in Table 1(f).

The manipulation needs to remove one $[c]$ vote and 11 $[c, a]$ votes and add one $[b, c]$ vote and 11 $[a]$ votes, to end with these tallies from the estimated completion $\hat{\mathcal{B}}$. \square

MOV-based Manipulator A minimal manipulation \mathcal{M} found by *margin-irv* specifies, for each type of ballot $R \in \mathbf{R}$, the number of ballots of that type that should be *added to* or *removed from* the election profile to achieve a desired elimination sequence π . The manipulation \mathcal{M} simply records for each ranking R of candidates: how many ballots are modified q_R to take on the new ranking R , and how many ballots with ranking R are modified m_R to show a different ranking. Its not possible that both q_R and m_R are non-zero for the same R , otherwise there is a smaller manipulation with the same effect.

In order for such a manipulation to be found in a reasonable time frame, the ILP of Section 2.2 operates over *equivalence classes* of ballot rankings, $\tilde{\mathbf{R}}$, rather than all possible rankings over a set of candidates \mathcal{C} , \mathbf{R} . Given an elimination sequence to achieve, π , each ranking in \mathbf{R} is reduced to a ballot class in $\tilde{\mathbf{R}}$. The original ranking is reduced by removing all candidates that would be eliminated by the time that ballot could possibly be placed in their tally. All ballots in the same class will move between the tally piles of the same set of candidates, at the same times. For example, consider an election with candidates a, b, c , and d , and a desired elimination sequence $\pi = [a, c, d, b]$. Ballots with rankings $[c, a, b, d]$, $[c, b, a]$, and $[c, a, b]$, are reduced to the equivalent class $[c, b]$.

The minimal manipulation \mathcal{M} found by *margin-irv* defines: a candidate elimination sequence to be achieved, π , in which the desired winner is victorious; a set of ballots D , in equivalence class form, to remove from $\hat{\mathcal{B}}$; and a set of ballots A , in equivalence class form, to add to $\hat{\mathcal{B}}$. For each ballot to add to the profile, there is a ballot to remove – leaving the total number of cast ballots unchanged (i.e., $|A| = |D|$).

Our MOV-based manipulator creates a manipulation rule for every ballot in D . For the i^{th} ballot in D , d_i , a rule of the form:

$$reduce_{\pi}(b) = d_i \rightarrow a_i$$

is formed, stating that if the manipulator sees a ballot b with a ranking that could be reduced to the equivalence class d_i (assuming the eventual elimination sequence will be π), this ballot should be replaced with the i^{th} ballot in A , a_i .

Example 4. The elimination order desired by the attacker is $\pi = [c, b, a]$. The equivalence classes of the seen ballot types are $[a]$, $[c, a]$, $[b]$ and $[c]$ respectively. The manipulation in terms of these equivalence classes is $+11[a]$, $-11[c, a]$, $+1[b]$ and $-1[c]$. We end up with 12 rules: 11 copies of $[c, a] \rightarrow [a]$ and 1 copy of $[c] \rightarrow [b]$.

If we perform the manipulation on the *actual* remaining ballots we will find enough ballots to change resulting in a final manipulated count $B_{\mathcal{M}}$ of $[a] : 66$, $[c, a] : 19$, $[b, c] : 37$ and $[c] : 14$. The election will first eliminate c and then b with a winning with an LRM of 24. Because the initial ballots were less favorable to the attackers candidate than in the full election the manipulation is larger than required. \square

First Preference Manipulator Recall that our first preference manipulator seeks to take $\Delta_{w, w'}$ ballots (divided by two and rounded up) in which a candidate w is preferred first, and replace them with a ballot in which candidate w' is preferred first. This manipulator creates $\lceil \Delta_{w, w'} / 2 \rceil$ rules of the form:

$$[w, \dots] \rightarrow [w']$$

The pattern on the left hand side of this rule matches all ballots in which candidate w is preferred first. Such ballots are replaced by the ballot $[w']$ containing a single preference for w' . If the manipulator is seeking to achieve a last round margin of a given size, Δ , it creates $\lceil \Delta_{w,w'}/2 \rceil + \Delta$ rules of the above form. Note that this naive manipulator is able to influence the last round margin of an election, but not the elimination margin of losing candidates in prior rounds. The manipulation is also heuristic – there is no guarantee that it will result in a desired winner, as it does not consider how preferences might flow between candidates. For example, robbing the original winner of some of their primary vote may result in their early elimination, distributing enough votes to cause an alternate candidate $c \neq w'$ to win. The MOV-based manipulator has more control over potential outcomes as it can alter the later preferences on each ballot.

Example 5. The difference in tallies between the actual winner c of the estimated election \mathcal{B} and the desired winner a is 44. The first preference manipulator then requires moving 22 + 20 votes from c to a in order to attain a LRM of 20 for a . The manipulation is then 42 copies of $[c, \dots] \rightarrow [a]$.

When we apply this manipulation to the actual ballots we find there are only 21 $[c, a]$ and 11 $[c]$ votes arriving in the remainder of the election which are all converted to $[a]$ votes. The final manipulated count is $[a] : 87$, $[c, a] : 9$, $[b, c] : 36$, $[c] : 4$. The winner is a with a LRM of 30 over b . These rather gross manipulations may bring the election result into question, since the final tallies are far from the actual tallies. \square

4 Results and Conclusions

We take the cast ballot data available for 5 seats of the Australian New South Wales (NSW) 2015 Legislative Assembly election, and simulate the use of our first preference and more intelligent MOV-based manipulators. The goal of each manipulator is to bring about the election of a specific candidate. For each type of manipulator, we simulate its application in each seat over 100 trials. In each trial, the order which the cast ballots arrive at the scanner is randomised. We compute, over the 100 trials: the average number of ballot changes made by the manipulator; the average total change in first preference tallies resulting from the manipulations; the number of simulations in which the manipulator achieved its desired winner, and achieved its desired winner with a last round margin sufficient to avoid an automatic recount; and the average number of manipulation rules generated by the manipulator after the processing of each αE batch of ballots. The latter statistic corresponds to the average number of *intended* manipulations the adversary still expects to need at each stage of processing.

All experiments have been conducted on a machine with an Intel Xeon Platinum 8176 chip (2.1GHz), and 1TB of RAM. We have used $\alpha = 25\%$ across each batch of 100 simulations. Each batch of 100 simulations have been initialised with the same random seed controlling the order in which ballots arrive at the scanner.

We first consider the relative performance of our first preference manipulator, and a MOV-based manipulator that does not attempt to minimise change across first preference tallies in the true and altered election. We then consider the effectiveness of a

Table 2: Performance of the first preference manipulator (aiming to enforce a LRM of at least Δ) and MOV-based manipulator (enforcing an elimination margin of Δ over the last 2 rounds of counting). The MOV-based manipulator is not using first preference discrepancy minimisation. We report the number of simulations (/100) in which the manipulators are successful, the average number of ballot manipulations performed, the average resulting change to first preference counts, and the number of manipulation rules (int. ballot changes) generated after each 25% proportion of ballots has been seen.

2Δ	First Preference Manipulator				MOV-based Manipulator			
	100	300	500	1000	100	300	500	1000
Ballina: MOV of 1,130; LRM of 1,267; 47,865 ballots cast								
Desired winner achieved	100	100	100	100	76	98	91	95
Avg ballot changes	3198	3198	3198	3198	1171	1265	1328	1407
Avg FP count changes	6397	6397	6397	6397	2059	2283	2408	2591
Avg int. ballot changes (1)	4699	4799	4899	5149	1218	1310	1396	1629
(2)	0	0	0	0	670	774	754	824
(3)	0	0	0	0	272	310	291	145
Balmain: MOV of 1,731; LRM of 1,731; 46,952 ballots cast								
Desired winner achieved	84	94	98	100	71	94	93	97
Avg ballot changes	1860	1935	2004	2203	1779	1865	1907	2023
Avg FP count changes	3720	3871	4008	4405	3215	3443	3556	3875
Avg int. ballot changes (1)	1747	1847	1947	2197	1776	1891	1968	2223
(2)	66	54	34	0	313	304	287	146
(3)	47	34	22	5	111	87	50	29
Campbelltown: MOV of 3,096; LRM of 3,096; 45,124 ballots cast								
Desired winner achieved	90	98	99	100	72	92	93	100
Avg ballot changes	3232	3313	3392	3590	3161	3237	3294	3512
Avg FP count changes	6464	6627	6784	7181	6166	6319	6452	6903
Avg int. ballot changes (1)	3130	3230	3330	3580	3151	3272	3357	3586
(2)	75	63	41	5	615	673	770	1023
(3)	27	20	21	5	130	91	61	12
Heffron: MOV of 5,824; LRM of 5,835; 46,367 ballots cast								
Desired winner achieved	85	96	97	100	71	95	99	100
Avg ballot changes	5928	6006	6087	6327	5878	5973	6068	6255
Avg FP count changes	11856	12013	12175	12656	11357	11450	11744	12346
Avg int. ballot changes (1)	5862	5962	6062	6312	5883	5979	6047	6320
(2)	744	844	944	1194	2680	2820	2746	3028
(3)	50	29	10	0	589	566	601	747
Lismore: MOV of 209; LRM of 1,173; 47,208 ballots cast								
Desired winner achieved	100	99	98	100	81	91	90	95
Avg ballot changes	4651	4696	4727	4742	339	396	434	510
Avg FP count changes	9304	9392	9455	9486	601	716	796	974
Avg int. ballot changes (1)	4371	4475	4572	4814	294	242	514	730
(2)	192	195	151	110	150	150	150	117
(3)	89	44	44	0	52	52	51	23

MOV-based manipulation that attempts to minimise such discrepancies. The computational requirements of each of our manipulators varies. The first preference manipulator is able to compute a set of manipulation rules in less than a second, the MOV-based manipulators (without first preference tally change minimisation) several seconds, while minimising first preference tally changes extends rule generation time by up to a minute.

Table 2 reports the performance of our first preference and MOV-based manipulators on the IRV elections held across our 5 case study seats: Ballina; Balmain; Campbelltown; Heffron; and Lismore. We report the true MOV, LRM, and number of ballots cast in each election alongside the effectiveness of each manipulator across 100 simulated trials. We have found that the MOV-based manipulator is more effective in achieving a change in winner when it seeks to enforce a reasonably sized elimination margin (EM) on the last two rounds of counting, rather than a margin on just the last round. In all our reported results, the MOV-based manipulators enforce an EM of Δ between the runner-up and winner, and the runner-up and their runner-up. The two settings – enforcing an LRM vs an EM on the last two rounds – are similarly effective in certain seats (Campbelltown, Heffron, and Lismore), with the latter more effective in Ballina. Note that the detailed results of this comparison have been omitted for brevity.

Ballina is a seat where the identity of the candidate coming in third significantly influenced which of the remaining two candidates won. The manipulator needed the Greens candidate to place third so that their preferences flowed to the manipulator’s desired candidate from the Country Labor Party. Simply enforcing a certain LRM in this instance resulted in manipulated elections in which the tallies of the two runner-ups were similar when determining third place. When forming a manipulation, there is no guarantee that all generated manipulation rules will be applied, and no guarantee that these rules will bring about the desired change in winner. The rules are computed based on hypothetical completions of partially known election profiles. Forming manipulation rules designed to ensure the Greens candidate was eliminated in this third-last position, with a reasonably sized elimination margin, more reliably achieved the desired result.

Table 2 shows that in general, our MOV-based manipulator requires less ballot changes on average to reliably (more than 90% of the time) bring about a desired change in winner. The first preference manipulator is often more successful, as more of its manipulation rules are likely to be applied in practice. A rule that is looking for a ballot with a certain candidate ranked first is more likely to be applied than one that looking for a ballot with a specific ranking of candidates. While a manipulation based on first preferences may, in some circumstances, underestimate the number of ballot changes required to alter an election outcome, it generally overestimates the degree of manipulation required. By focusing on preference flow in Lismore, the MOV-based manipulator reliably achieves a desired outcome by changing orders of magnitude less ballots, on average. In Balmain and Campbelltown, the first preference manipulator generates a slightly smaller ‘intended manipulation’, forming less manipulation rules, on average, after seeing the first 25% batch of ballots. The MOV-based manipulator applies fewer of its generated rules, in these two contests, leading to fewer ballot changes on average.

It may be the case that ballots are manually examined to compute first preference tallies for each candidate, while the full count is performed by a computer. In this setting, a large discrepancy between the first preference tallies reported by the software, and that of the manual count, is likely to arouse suspicion. Table 3 compares the average discrepancy in first preference counts (between the true, unmanipulated elections, and those altered by our manipulators) when using a MOV-based manipulator that *does not* focus on minimising these discrepancies, and one that *does*. Note that if a number of ballots N is shifted from the first preference tally of one candidate to another, this

Table 3: Number of trials (/100) in which two MOV-based manipulators achieve a desired winner change. The first does not minimise first preference count discrepancies, while the second does. We report the number of successful manipulations in which the resulting election avoided an automatic recount ($LRM \geq 100/2 = 50$ votes).

2Δ	MOV-based Manipulator No FP change minimisation				MOV-based Manipulator Minimise FP count changes			
	100	300	500	1000	100	300	500	1000
Ballina: MOV of 1,130; LRM of 1,267; 47,865 ballots cast								
Desired winner achieved	76	98	91	95	72	99	100	88
Avoid auto recount	76	98	91	95	72	99	100	88
Avg ballot changes	1171	1265	1328	1407	1185	1296	1463	1735
Avg FP count changes	2059	2283	2408	2591	230	324	512	714
Avg int. ballot changes (1)	1218	1310	1396	1629	1216	1368	1457	2042
(2)	670	774	754	824	789	870	998	1235
(3)	272	310	291	145	398	434	502	403
Balmain: MOV of 1,731; LRM of 1,731; 46,952 ballots cast								
Desired winner achieved	71	94	93	97	50	93	100	92
Avoid auto recount	48	82	83	94	26	75	100	87
Avg ballot changes	1779	1865	1907	2023	1784	1958	2147	2502
Avg FP count changes	3216	3443	3556	3875	849	970	1079	1136
Avg int. ballot changes (1)	1776	1891	1968	2223	1783	1954	2130	2585
(2)	313	304	287	146	1173	1241	1316	1429
(3)	112	87	50	29	577	612	637	545
Campbelltown: MOV of 3,096; LRM of 3,096; 45,124 ballots cast								
Desired winner achieved	72	92	93	100	59	93	100	95
Avoid auto recount	57	85	89	100	31	79	100	92
Avg ballot changes	3161	3237	3294	3512	4034	4115	4245	4366
Avg FP count changes	6166	6319	6452	6903	3051	3122	3329	3654
Avg int. ballot changes (1)	3151	3272	3357	3586	4213	4325	4380	4661
(2)	615	673	770	1023	2134	2108	2173	2116
(3)	130	91	61	12	875	901	906	758
Heffron: MOV of 5,824; LRM of 5,835; 46,367 ballots cast								
Desired winner achieved	71	95	99	100	45	93	100	100
Avoid auto recount	49	89	96	100	31	83	100	100
Avg ballot changes	5878	5973	6068	6255	6595	6719	6897	7651
Avg FP count changes	11357	11450	11744	12346	10607	10910	11414	12029
Avg int. ballot changes (1)	5883	5979	6047	6320	6635	6748	6904	7611
(2)	2680	2820	2746	3028	2976	3164	3291	3392
(3)	589	566	601	747	1402	1466	1568	1697
Lismore: MOV of 209; LRM of 1,173; 47,208 ballots cast								
Desired winner achieved	81	91	90	95	72	91	96	95
Avoid auto recount	77	83	86	87	67	91	96	95
Avg ballot changes	339	396	434	510	339	438	521	621
Avg FP count changes	601	716	796	974	215	283	368	542
Avg int. ballot changes (1)	294	242	514	730	384	535	719	1140
(2)	150	150	150	117	206	293	354	240
(3)	52	52	51	23	97	86	64	77

is viewed as a discrepancy of $2N$ votes. The discrepancy minimising manipulator was able to reduce change in first preference counts by 2.5 times, on average, between a factor of 1.1 and 9, while requiring only a small increase in total ballot changes. In Ballina,

we are able to reliably realise a desired winner change while producing a first preference count discrepancy that is significantly lower than the MOV or LRM of the election. Heffron and Campbelltown, with their large margins of victory, are more challenging to manipulate in a non-obvious manner.

Irrespective of whether first preference count changes are being minimised or not, our MOV-based manipulator can successfully alter the outcomes of elections, while avoiding an automatically triggered recount.

Minimising first preference count changes requires a more subtle manipulation, with the later preferences on ballots being altered more often. The result is that the manipulator must aim to achieve larger elimination margins in the last two eliminations to reliably achieve a desired winner change. The manipulator can be too ambitious however, and try to achieve elimination margins that are not realistically achievable. A limitation of the MOV-based manipulators is that if they cannot find a manipulation of a given hypothetical complete election profile that satisfies all desired side constraints, they give up and fail to generate a set of manipulation rules. A more effective strategy would relax these constraints – with smaller requirements on elimination margins – until the algorithm of Section 2.1 is able to find a manipulation.

Conclusions The experiments show that it is quite feasible for an attacker to manipulate an election to change the winner with high confidence in the scenario we examine. Using MOV-based manipulation and minimising first preference changes the attacker can avoid an automatic recount, and often significantly reduce the number of first preference changes. Hence we can conclude that rigorous risk limiting audits of elections is warranted, since simple counting based approaches to auditing can be defeated.

References

1. Michelle Blom, Peter J. Stuckey, and Vanessa Teague. Computing the margin of victory in preferential parliamentary elections. In *Proceedings of the E-Vote-ID 2018: Third International Joint Conference on Electronic Voting*, 2018.
2. Michelle Blom, Peter J. Stuckey, and Vanessa Teague. Election manipulation 100. In *Proceedings of the Fourth Workshop on Advances in Secure Electronic Voting (Voting'19)*, 2019.
3. Michelle Blom, Vanessa Teague, Peter J. Stuckey, and Ron Tidhar. Efficient computation of exact IRV margins. In *Proceedings of the 22nd European Conference on Artificial Intelligence*, 2016.
4. Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating manipulations in voting with partial information. In *AAAI Conference on Artificial Intelligence*, 2011.
5. T.R. Magrino, R.L. Rivest, E. Shen, and D.A. Wagner. Computing the margin of victory in IRV elections. In *USENIX Accurate Electronic Voting Technology Workshop: Workshop on Trustworthy Elections*, USENIX Association Berkeley, CA, USA, 2011.
6. R. Richie. Instant Runoff Voting: What Mexico (and Others) Could Learn. *Election Law Journal*, 3:501–512, 2004.
7. Lirong Xia. Computing the margin of victory for various voting rules. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC '12*, pages 982–999, New York, NY, USA, 2012. ACM.

On practical aspects of coercion-resistant remote voting systems

Kristjan Krips^{1,3} and Jan Willemson^{1,2}

¹ Cybernetica AS, Ülikooli 2, 51003 Tartu, Estonia

{krisjan.krips,jan.willemson}@cyber.ee

² STACC, Ülikooli 2, 51003 Tartu, Estonia

³ Institute of Computer Science, University of Tartu, J. Liivi 2, Tartu, Estonia

Abstract. Coercive behaviour is hard to control in the remote electronic voting setting. This is why a number of protocols have been proposed that aim at mitigating this threat. However, these proposals have remained largely academic. This paper takes the practical viewpoint and analyses the most common assumptions that are required by the various schemes, together with the exact level of coercion-resistance they provide.

1 Introduction

With introduction of Australian secret ballot into the voting process in mid-19th century, the threat of voter coercion was significantly reduced. Voting in a private booth surrounded by a controlled environment became the “gold standard” which has served democratic societies around the world well for over a 100 years.

However, several developments in recent decades have undermined the effect of Australian ballot as a coercion-resistance measure. First, technology of recording the private events within the voting booth (both with the voter cooperation and stealthily) has become readily available [3,17,18]. And second, human mobility has increased to an extent where expecting all the voters to come to a controlled environment on a particular day is less and less of an option [32].

These problems have motivated research and development in the field of coercion-resistant (remote) voting solutions. However, only a few of these solutions have actually been implemented in practice, leaving practical considerations such as usability or technical complexity of satisfying necessary assumptions often out of scope.

Another issue with the notion of coercion resistance is that it does not have a single clear interpretation. Thus it is not always immediately clear which levels of coercion resistance are achieved by various proposals.

The current paper aims at narrowing these gaps. We have selected seven different schemes from recent proposals and analyse them from two viewpoints. First, we identify common technical and organisational assumptions that these schemes rely on and assess their practical satisfiability. Second, we gather different interpretations of coercion resistance and analyse to what extent each one of the considered schemes achieves them.

We do not claim full coverage of all coercion-resistant schemes that have ever been proposed, but we have made an attempt to put together a representative selection of different approaches used for remote voting. Also, voting schemes often come in families. In this case we have selected members of such families for which coercion resistance and/or usability issues have been addressed the most.

2 Notions of voting freedom

One of the fundamental requirements of democratic elections is that the voter should be able to express her true preference freely, i.e. without being coerced. This broad statement has several possible interpretations, leading to more fine-grained requirements. E.g. following [4], we can identify the following properties.

- *Basic ballot privacy* guarantees that no one can learn how a voter voted (if she is not coerced and is willing to keep her vote secret). All the voting schemes studied in this paper satisfy this requirement.
- *Receipt-freeness* ensures that a malicious voter is unable to produce a proof for the value of her vote, making coercion essentially inefficient.
- *Coercion resistance* means intuitively that the voter should be able to cast a vote reflecting her true preference even if being monitored by the coercer for (most of) the voting period. To distinguish this property from the generic term, we will also call it *over-the-shoulder coercion resistance* in this paper.

Juels *et al.* [16] go even further and state three additional requirements that a fully coercion-free voting system should correspond to.

- The coercer should not be able to force the voter to abstain from elections.
- The coercer should not be able to force the voter to cast an invalid vote.
- The coercer should not be able to cast a valid vote if he gets access to the voter’s credentials.

3 Coercion-resistant schemes and their assumptions

The threat of coercion depends on many aspects: type of elections, properties of the voting protocol, assumptions on the voting system and environment, awareness and coercibility of the voters, capabilities of the attacker, etc.

Typically, voting protocols aiming at some form of coercion resistance must make trade-offs between different goals. In the following, we describe and classify existing coercion resistant voting protocol proposals according to their assumptions, usability and applicability for different types of elections.

3.1 Re-voting based schemes / Estonian scheme

Re-voting is a metatechnique that can be used on top of other voting systems to provide voter with an option of changing her vote in case she was coerced

during the first attempts(s). An example of a pure re-voting-based protocol is the Estonian scheme, where this is the only anti-coercion measure in use [22].

The biggest problem with such schemes is that the coercer might stay with the voter until the end of the voting period (either physically or virtually [3]) to make sure that she does not cast a re-vote. To mitigate this threat (and also some other risks of remote voting), Estonia has chosen to end the Internet vote submission two hours before the polling stations are closed on the last day of advance voting period. The rationale is that if the voter feels coerced, she still has some time to submit her vote on paper and the paper vote cancels the e-vote. However, if the voter resides far from any of the polling stations (and enabling this scenario is one motivation of Internet voting), she can not submit an uncoerced vote. The whole system operates under the assumption that the share of such events is insignificant.

In addition, the re-voting functionality can affect integrity of the cast vote as an active attacker may use it to overwrite the previous vote.

On the positive side, enabling re-voting does not need extra setup on the client side, and the process is easy to understand for an average voter.

Aside from that, the Estonian system relies on significant technical assumptions, most notably voter credential pre-distribution. This is implemented via the national digital identity mechanisms (ID-card and mobile-ID), with the corresponding public keys being available via national PKI. Thus, even though the Estonian scheme relies on special client-side hardware, these devices are already very widely in use.

3.2 JCJ/Civitas family

Formal study of coercion resistance in voting systems was initiated in 2002 by Juels, Catalano and Jakobsson [15]. They gave a definition of coercion resistance and proposed the first scheme satisfying it, later becoming known as the JCJ scheme [16]. This research introduced fake credentials which the voter can use under coercion, but the coercer is unable to distinguish from the genuine ones.

In 2008, the JCJ scheme was extended by Clarkson, Chong and Myers by introducing distributed trust assumptions and improving the performance. The resulting protocol was called Civitas [8].

Neither of the JCJ and Civitas proposals specified how exactly the voter should select the appropriate credentials. Neumann and Volkamer noted in 2012 that this action is non-trivial, and may lead to both usability and security issues when implemented carelessly. Improving the specification of Civitas, they proposed an implementation based on smart cards and readers with PIN-pads and trusted displays [28]. Selection between a fake and a real credential would be accomplished by entering either a real or a fake PIN into the reader.

Essentially, Neumann-Volkamer proposal encapsulates all the critical voter-side operations into special hardware, which has to be trusted. While in principle such an approach can make credential handling more secure, it does not really move us much closer to a practical implementation. Smart card readers with

trusted preview are not commonplace on the market, and the smart cards would require a lot of non-standard functionality.

In a later research Neumann *et al.* have shown that, in principle, modern smart cards have sufficient performance required to implement such functions [27]. However, performance is not the only bottleneck in the practical deployment. The software implementing the protocol functionality needs to somehow get onto the cards.

Roughly speaking, election organisers have two approaches to tackle this problem. First, they can approach a large smart card vendor and convince it to implement the required functionality as part of the card firmware. Our interview with a representative of Gemalto (previous supplier of Estonian ID-cards) revealed that smart card vendors are quite reluctant to include limited-use applications on their products, and prefer implementing only general-purpose cryptographic primitives like standardised asymmetric signatures. One reason for this is that in many applications (likely including voting as well) the customers require certification, testing and validation of the security features of smart cards (for example, according to Common Criteria standard or FIPS-140-2). Such processes are expensive and time-consuming, and the vendor cannot earn this investment back selling limited-use cards.

Another option would be using programmable cards and implementing the functionality oneself in the spirit of [27]. The drawback of this approach is the need to support the whole software development life cycle locally. While it may give better control over the implementation, the risks are also higher. In case a bug is discovered, updating applications on the cards that have been distributed to numerous remote voters is a nightmare. Also, the whole expense of certification (in case it is desired) needs to be carried locally.

We conclude that while special-purpose smart cards provide an appealing option for a “poor man’s HSM”, their deployment has problems that are not necessarily easier to solve than the original challenge they were designed to meet.

As a part of the registration procedure, the Neumann-Volkamer protocol also depends on availability of anonymous channels (e.g. Tor is suggested by the authors). We refer to Section 3.4 for a more elaborate discussion on difficulties of achieving anonymous channels (using Tor) in practice.

Another branch of JCJ was developed by Araújo *et al.* in 2010 [2]. They introduced shorter credentials and provided a formal proof of coercion-resistance, although their proof relied on a non-standard number-theoretic assumption. In 2018, Neto *et al.* conducted usability studies for the CIVIS system [26], which is an implementation of the protocol proposed by Araújo *et al.* [2]. The study revealed that more than 90% of the test participants did not understand the functionality of casting fake votes. Also, they did not feel comfortable with the result, being unable to distinguish whether their submitted vote was real or fake. This brings the whole concept of using fake credentials under question.

3.3 Helios family

The original proposal of Helios by Adida [1] was explicitly targeted towards low-coercion environments. During later research, several extensions have been developed to enhance its coercion resistance.

KTV-Helios Kulyk, Teague and Volkamer have extended the Helios voting system to provide private eligibility verifiability, i.e. the property that anyone can verify that only votes from eligible voters are included in the tally, without revealing who actually submitted them [20,19]. As a by-product, they achieve receipt-freeness in the sense that the voter can not prove how she voted as she can undetectably re-vote. However, the authors stated that the protocol is susceptible to forced abstention and randomisation attacks. Following the authors' initials, the scheme is known as KTV-Helios.

The core idea of Kulyk *et al.* is to hide the true votes among dummy ones. Receipt-freeness is achieved allowing the voter to cast differential vote updates, so that the final vote would be a combination (e.g. product) of the votes cast. A similar approach was independently developed by Locher and Haenni [21].

Even though the dummy votes can be cast by any voter, most of them would probably not bother to do so. Hence a specific party called posting proxy or posting trustee is introduced by Kulyk *et al.*, and its task is to submit the dummy votes. In order to prevent timing side channels (see Section 3.4), posting trustee must operate in a randomised fashion.

Regarding the practical implementation aspects, the authors of KTV-Helios admit themselves that the understandability and usability issues remain largely unsolved [20]. Seeing many votes submitted onto the bulletin board on her behalf probably makes an average voter quite anxious. We add here a potential legal problem of voter impersonation, even if there are cryptographic proofs certifying that the extra votes do not change the final tally.

BeleniosRF In the original version of Helios, the voter can present encryption randomness as a receipt for the coercer. BeleniosRF uses re-randomisable ciphertexts and signatures, with part of the randomness being out of the voter's control, making it impossible for a voter to produce such a receipt [4].

The ballot is signed by the voter and re-randomisation of the ballot by the server does not invalidate the corresponding signature. Thus, the voter can verify the signature to make sure that the vote has not been changed. However, this applies only when re-voting is not enabled. The authors of BeleniosRF state that in case of re-voting the voters would not be able to check which of their ballots were re-randomised by the server. Therefore, BeleniosRF does not allow re-voting and thereby does not provide protection against over-the-shoulder coercion. However, vulnerability to in-person coercion is one of the major objections against remote electronic voting in the first place [10,11,25,24,13,14].

The authors of BeleniosRF argue that changing one's vote is a legally grey area anyway, and most of the countries would need to go through a complicated legal process before they can support it.

While we agree that legislative changes are necessary to support re-voting, we feel that the authors of BeleniosRF over-estimate the complexity of this process.

For example, extensive social and legal debate concerning constitutionality of re-voting took place in Estonia when Internet voting was introduced there. A few months before the first Internet-enabled elections, the President of Estonia brought Internet voting provisions to the Supreme Court for constitutional review, arguing that the possibility to change Internet votes gives advantages to Internet voters in comparison with paper voters. The decision of the Supreme Court did not support this point of view, reaching the conclusion that merely a technical option of casting multiple votes does not put Internet voters into any kind of advantage [23].

While the outcome of a similar legal discussion may be different in other jurisdictions, we feel that re-voting as an easy-to-implement and relatively efficient anti-coercion measure is important enough to review some of the legislative principles. Changing legislation in order to catch up with technological advancements is an unavoidable process anyway.

3.4 Selene

The primary design goal of the Selene scheme proposed by Ryan *et al.* [31] is achieving a user-friendly end-to-end vote verification protocol. As too strong of a verification mechanism brings along a threat of coercion, the authors of Selene have also paid a lot of attention to mitigating this threat. They propose using cryptographic tracking numbers which are first committed to a bulletin board using trapdoor commitments. After the end of the voting period, clear-text votes with clear-text tracking numbers are displayed on the bulletin board as well. The (voter-controlled) trapdoor can later be used to open the commitment to any tracking number of coercer's liking.

The voter, of course, still needs to somehow identify the real tracking number of her own vote. This is facilitated by sending her the correct decommitment value α . In order to fool the coercer, the voter can produce an alternative decommitment value α' that is cryptographically indistinguishable from α and points to any vote requested by the coercer.

However, cryptographic indistinguishability is not sufficient, as the attacker potentially has a number of side channels available to separate the true α from voter-generated α' -s. The authors of Selene acknowledge this problem and state that α -terms should be transferred over an unauthenticated and private channel.

Unfortunately, implementing such a channel is non-trivial. Note first that in order to mitigate the threat of coercion, it is not sufficient just to drop strong authentication mechanisms like signatures. For example, if α (or its shares coming from the trustees) is sent via regular, otherwise unauthenticated email, it has to carry sender's email address. There are both legal and usability issues that suggest using a fixed official address rather some randomly generated ones. Email is just an example here, similar problems would occur if other taggable delivery channels like instant messaging or web bulletin board would be used.

In principle, the process of preparing a false α' can also include sending it from the official address. In this case there is still the timing side channel that the coercer can use to distinguish the genuine α . In order to counter this, the

genuine α -s would need to be sent at randomised moments, and the voter must prepare α' during this period. This is doable and is also proposed by the authors, but it complicates the voter's view of the protocol substantially.

We can also imagine genuine α -s being sent out via regular mail, printed on standard office paper. The voter can print α' out on her home printer, but this assumes using exactly the same kind of paper, printing resolution, etc. In addition, majority of modern colour laser printers mark the printed papers with tracking dots which can be used to identify the printer [30]. We can see that it could be possible to deliver α -s with the help of the postal service, but generating the fake values is not as easy as the authors of Selene probably foresaw. Note also that a vote buyer is typically after a number of votes and he can live with some of the voters being able to fool him as long as their share is not too high.

One can also utilise stronger anonymisation techniques, e.g. mixing or onion routing. These would only help if the full set of messages is larger than just the official α -terms, as otherwise we would have no sender anonymity. One may consider using an existing anonymisation network, say, Tor (as also recommended by Neumann and Volkamer [28]). However, due to significant illegal activity happening over it, utilising Tor for legally binding elections would be questionable.

We argue that this dilemma is at least partially inherent and not specific to Tor. On one hand, too small of an anonymisation set does not fulfil the goal, but fighting doubtful traffic in a large network is practically impossible.

Furthermore, by relying on Tor (let's still use it as a prime example) new problems are introduced. Referring to the objectionable content and general uncontrollable nature, several countries have attempted to block/filter Tor traffic⁴. This makes it hard for expatriates living in those countries to participate in the elections remotely, but supporting expatriate participation is one of the main reasons for introducing remote electronic voting in the first place.

Setting up private channels from the election organiser to all the voters is not a trivial task either. As Selene already relies on a PKI for vote signing, assuming additional access to an authentic public-private key pair for encryption and decryption is probably not a big extra. However, even the α -term encrypted with the voter's public key has to be delivered to her somehow. We conclude that channel privacy does not really help against the soft sender identification problem described above.

In 2019, Distler *et al.* performed an e-voting usability study based on a Selene protocol implementation [9]. Unfortunately, they left the steps related to coercion resistance (including preparing the fake α' and selecting it in the presence of the coercer) out of scope. We also note that their implementation relies only on a mobile device for both vote casting and verification. This means that verification is inefficient against the malicious device and does not thus fulfil the purpose of verification. We feel that in order to get a more realistic understand-

⁴ It is hard to get reliable statistics on the extent of Tor filtering, but there exists indirect evidence in the form of the share of users relying on Tor bridges (<https://metrics.torproject.org/userstats-bridge-table.html>) and observed irregularities (<https://metrics.torproject.org/userstats-censorship-events.html>).

ing of usability of Selene protocol the authors of [9] should have implemented a complete version, e.g. by using a second channel for verification. Adding extra channels and steps would have likely changed the user perception and feedback.

3.5 Eos

Patachi and Schürmann have proposed the Eos voting scheme based on a specific flavour of ring signatures, namely conditional linkable ring signatures [29]. As each voter can have multiple pseudo-identities in the scheme, conditional linkability allows the signer to choose if the signatures can be linked to the same identity by the verifier.

There are two main anti-coercion measures in Eos. First, the voter can use subliminal hinting (called selecting between “red” and “green” envelopes or alternative pseudo-identities in [29]) while preparing the encrypted vote. In practice, such hinting would be implemented by presenting either a real or pseudo-PIN to a special-hardware voting device or to the coercer who controls the device.

Second, if the actively coerced voter had to cast a vote using a valid PIN, she may later re-vote to update the vote. However, in that case the public bulletin board will contain multiple encrypted votes given by the same pseudo-identity, which may be known to the coercer. In that case, the voter may have to lie to the coercer that the coercer was the last one to cast the vote.

The protocol makes several non-trivial assumptions. First, to get rid of side-channels during submitting the ring-signed votes, one would need to use anonymous channels, but achieving these is quite tricky in practice (see Section 3.4).

Second, special hardware tokens would be needed to implement the client-side operations (key management, PIN validation, identity selection, and signature computation). The paper [29] suggests that hardware wallets designed for storing the keys for cryptocurrencies could be used in this role. It might be possible to reprogram such hardware, but distributing the hardware or the private keys to the voters is a non-trivial task.

As the selection between identities would happen by entering a real or pseudo-PIN, we also have all the regular problems of pseudo-PIN management – if the user enters a wrong PIN, the device can not give any feedback (as the coercer might be watching), and would quietly submit a vote that the voter did not intend to (e.g. in the scenario where the voter wanted to use a pseudo-PIN, but accidentally used a real one).

3.6 Selections

A special form of fake credentials called *panic passwords* has been proposed by Clark and Hengartner in 2008 [5]. The essence of panic passwords is what the name says – the user can select a true password together with a set of alternative ones that can be used to covertly alert the system that the user is under abnormal circumstances, e.g. coercion.

The latter is an important threat scenario in case of remote voting, so the same authors have built a coercion-resistant voting scheme called Selections around their core idea [6].

Unfortunately, making human-memorable passwords to work as fake credentials is even more problematic than in case of cryptographic credentials.

First, a complex registration process is needed. Of course, it has to take place in a controlled, coercion-free environment, but this is a standard assumption. The registration procedure can even be implemented bare-handed (i.e. not requiring the voter to perform computations by heart). An Internet-enabled computer is still required inside the controlled registration booth to print out a voter preparation sheet. This is meant as a countermeasure “. . . in the event that an adversary ensured she entered the registration process without her sheet” [6].

The only way the coercer can achieve this is to search through the voter’s belongings and walk together with her until the door of the registration booth. But if the coercer is prepared to do this much, he can also request the voter to record all her actions with a camera or even send a live stream [3]. As a result, the effect of controlled registration environment will be significantly reduced.

During the registration process, the previously selected and encrypted panic passwords are re-randomised. The voter selects one of the re-randomised encryptions which is posted to a public roster. It is assumed in the protocol that the voter deletes the randomness used for re-randomisation and does not record it. Building security properties on the assumption that some value is deleted is always questionable. There may exist side channels that the coercer forces the voter to use to record or stream the value. If the coercer took part in creating the voter preparation sheet and has access to it, then the re-encrypted panic password on the public roster can be matched with the encrypted panic password on the preparation sheet. Thus, the randomness gives a way to prove the validity of the password given to the coercer.

We also noticed that the registration protocol differs significantly when comparing the full paper (e-print) [7] to the conference paper [6]. In the e-print version, the registration protocol allows the voter to rewind the process back to the re-randomisation phase. In the conference paper, the registration protocol allows the voter to rewind the process back to the beginning, i.e., to selecting new panic passwords. However, the difference is important as some of the coercion protections depend on the rewinding functionality.

Additionally, Selections suffers from the typical problems of password-based systems. Even though [6] proposes measures to increase password memorability, the scenario of voting stretches these boundaries. The idea of [6] was to go through the complex registration process once and then use the credentials over several events. However, elections typically only happen once in a few years, and many voters are likely to forget their passwords over this time, no matter how good of a mnemonic is used. To counter this problem, humans tend to write the passwords down, increasing their coercibility as a result.

4 Other coercion properties

In this section, we discuss the extra coercion properties (i.e. forced abstention, casting an invalid vote, and forced surrender of credentials) of the schemes.

A voter can be forced to not take part in the elections if a coercer has a way to check if the voter abstained from voting. As potential attackers, we also consider corrupt election officials and democratically elected politicians who decide to deviate from fair election practices. Such an attacker would be able to indirectly manipulate a large portion of the electorate.

Forcing the voter to cast an invalid vote can benefit the coercer in (at least) two ways. First, in case the voter is supporting a party opposing the coercer's views, the invalid vote would have no effect and the voter would effectively abstain from the elections. Second, if the invalid vote would be posted to a bulletin board, the attacker could remotely check if the voter behaved according to the instructions. Even if the invalid vote would not be published, it may still be possible that election officials are able to see the value of the vote and thus be able to play the role of a coercer.

In case another person would be able to use the voter's credentials, it would be possible to cast the vote on behalf of the voter. Juels *et al.* [16] refer to this type of an attack as a *simulation attack*.

The rest of this Section is devoted to the discussion of these coercion properties. Table 1 summarises the main assumptions used by different coercion-resistant protocols proposals together with their level of coercion-resistance in respect to the requirements listed in Section 2. The only exception is the basic ballot privacy that all the considered schemes trivially satisfy.

4.1 Re-voting based schemes / Estonian scheme

The Estonian voting system provides protection against standard versions of these coercion attacks. More specifically, an outside third party is not able to detect if a voter cast a vote online or abstained as there is no public proof of the vote casting. There is a private bulletin board in the Estonian voting system, which is only accessible to the election officials and auditors. The official voting client software does not support casting an invalid vote. Finally, the signing key of the voter is stored inside of a smart card, hence the coercer would need to have physical access to use the credentials.

However, the situation gets more complicated in case of an attacker who has insider information. The voting system has to verify the ballot signatures to make sure that only the votes of eligible voters are accepted. Thus, insiders could check if a certain voter abstained.

There is also an insider threat when an invalid vote is cast. To cast an invalid vote, either the voter or the coercer would have to create a non-standard voting client. In case the invalid vote would have a correct format and would correspond to a non-existing candidate number of a suitable district, the vote would be decrypted during the tallying process. Writing a voting client that would allow casting such votes is possible as the voting protocol and the communication API

Table 1. Cross-table of assumptions and achieved coercion resistance properties

	<i>Estonia</i>	<i>NV-Civitas</i>	<i>KTV-Helios</i>	<i>BeleniosRF</i>	<i>Selene</i>	<i>Eos</i>	<i>Selections</i>
Special client hardware	● ¹	●	●	○	○	●	○
Anonymous channels	○	●	●	○	●	●	●
PKI / key distribution	●	● ²	●	●	● ²	● ²	○
Subliminal password/PIN hinting	○	●	○	○	○	●	●
Casting a re-vote	●	●	●	○	○ ³	●	●
Non-trivial registration	○	○ ⁴	○	○	○	○	●
Receipt-freeness	○	●	●	●	○ ⁵	●	○ ⁶
Over-the-shoulder coercion resistance	●	●	○ ⁷	○	○ ⁸	●	●
Resistance to forced abstention	○ ⁹	●	○ ¹⁰	○	○ ¹¹	●	○ ¹²
Resistance to casting an invalid vote	○ ⁹	●	○ ¹³	○ ¹⁴	○ ¹⁵	○ ¹⁶	○ ¹⁷
Resistance to simulation attack	○ ¹⁸	●	○ ¹⁹	○	○ ²⁰	○ ²¹	○ ²²

● = is assumed / holds ○ = is not assumed / does not hold ○ = may hold
 ○ = depends on the implementation

¹ Smart card based ID-cards are mandatory in Estonia and widely in use.

² PKI is not explicitly mentioned, but its functionality is implicitly described.

³ Whether re-voting is allowed in Selene depends on the used policy [31].

⁴ Information about the registration process of NV-Civitas can be found in [28].

⁵ Selene’s receipt-freeness depends on the anonymous channel, see Section 3.4.

⁶ Whether Selections is receipt free depends on how the re-randomisation randomness is handled during registration. For more information, see Section 3.6.

⁷ The property depends on how the coercer prevents re-voting, see Section 4.3.

⁸ The property depends on the re-voting policy in the implementation of Selene [31].

⁹ The attack can be implemented by an insider, see Section 4.1.

¹⁰ KTV-Helios is susceptible to forced abstention only in the case of an active attacker.

¹¹ For information about the implementation of Selene, see Section 4.4.

¹² It is not clear whether Selections is resistant to forced abstention, see Section 4.6.

¹³ In KTV-Helios invalid votes can be cast, but they will be removed by plaintext equality tests before votes are published in the bulletin board.

¹⁴ See Section 4.3 for information about the coercion properties of BeleniosRF.

¹⁵ Vote casting procedure is not specified in Selene, see Section 4.4 for more details.

¹⁶ Whether it is possible to cast an invalid vote depends on the version of Eos. More information can be found from Section 4.5.

¹⁷ It is not specified how vote is encoded and how votes are tallied in Selections [6].

¹⁸ The coercer might be able to get physical access to the smart card. However, it is possible to re-vote as described in Section 4.1.

¹⁹ The coercer might be able to get physical access to the smart card. However, the voter may be able to re-vote to cancel the coerced vote as described in Section 4.3.

²⁰ It is not specified how keys are managed in Selene [31]. In case Selene is used as an add-on, then key management may be specified by the underlying voting protocol.

²¹ The possibility of casting a valid vote with the voter’s HSM depends on the configuration of the HSM. For more information, see Section 4.5.

²² If the registration process and thus the credentials are remotely monitored then the voter has the option to revoke the registration and vote in person. For more information, see Section 4.6.

is public. Now, if a voter would be able to cast such an invalid vote then either the members of the election committee or the auditor who audits the election result might be able to read the invalid value. Thus, the coercer would have to cooperate with the election officials or the auditor to see the vote value.

In order to get a hold of the signing keys, the coercer would have to take the possession of all of the digital ID-s of the voter along with the corresponding PIN codes. Still, the voter could use a non-digital ID to cast a paper vote in the polling station that overwrites the e-vote. Thus, all non-digital ID-s would also have to be collected by a coercer in case the coercer would like the voter to abstain from participating in the elections. Such an attack could be applied on selected individuals, but this approach does not scale.

4.2 NV-Civitas

NV-Civitas was the only one of the protocols that we analysed not susceptible to the three aforementioned coercion attacks. Forced abstention is impossible as the ballots are not signed and are delivered over an anonymous channel. Invalid ballots are either rejected by the smart card or by the voting system after checking the proof of vote well-formedness [28]. It is also impossible to force the voter to surrender the credentials as the voter can give the coercer the smart card with a fake PIN, which would create a ballot with invalid credentials.

4.3 Helios family

KTV-Helios While it is possible to cast invalid ballots in KTV-Helios, they do not end up on the bulletin board. Invalid ballots are removed before tallying with the help of plaintext equality tests. Thus, invalid votes are not decrypted.

The authors state that casting an invalid vote can cause the voter to abstain from the elections. The attack would always work in an active scenario where the attacker waits until the end of the voting period to force the voter to cast an invalid vote. In this case the invalid vote would be discarded before the tally and the voter would not have enough time to re-vote. However, if the value of the invalid vote would be known to the voter and there would be time to re-vote, the voter may be able to cancel the previous vote.

As the signing keys are stored on smart cards, it is in principle possible to force the voters to give up the cards, but such an attack would not scale well.

BeleniosRF BeleniosRF uses a fixed message space for encoding the vote and tallying is done homomorphically. Thus, the possibility of casting an invalid vote depends on the implementation. In case the message space is not used up to encode the candidates, it might be possible to cast an invalid vote that would be published.

The other two coercion attacks could be applied in the case of BeleniosRF. It is possible to force the voter to abstain from voting as there is public proof of participation in the voting event. The signature of the randomised public ballot can be verified by the voter. In case the voter's public key is accessible to the coercer, the latter is able to verify all the ballots on the bulletin board. Also, the

voter ID is verified before a ballot is accepted and re-randomised by the bulletin board. Thus, the election officials could coerce voters to abstain.

A coercer might also be able to force the voter to surrender her secret key as no special hardware is used for storing the secret key. However, the voter is only able to give one vote, so the coercer would have to get access to the signing key before the voter casts her vote.

4.4 Selene

Whether Selene is safe from forced abstention attack depends on the implementation of the protocol. The basic scheme is vulnerable as the ballots signed by the voters are published on the bulletin board. However, the optional enhancement of using pseudonymous credentials enables giving signatures without revealing the identity of the voter. Thus, the extended scheme is resistant to forced abstention attack if the coercer can not access the voter's pseudonymous credentials.

Similarly, the ability to cast an invalid vote depends on the implementation of the vote casting procedure and is not fixed on the protocol level. Selene can be used as an add-on on top of another voting system, which may remove invalid votes. E.g., Selene combined with JCJ is resistant to casting an invalid vote [12].

Still, Selene is susceptible to forced surrender of credentials as no hardware token is proposed for storing the secret key. Also, re-voting policy is not fully specified, thus it is not clear if voter's initial choice could be overwritten.

4.5 Eos

Eos is resistant to the forced abstention attack. It uses ring signatures to hide voter identities from the election officials. Also, an anonymous channel is used to cast the vote. Thus, it won't be possible to detect if a specific voter has voted.

The authors of Eos acknowledge that in the basic version of the protocol a coercer could force a voter to cast an invalid vote [29]. As a solution, they propose using a disjunctive zero-knowledge proof protocol, such that the voter could prove that her vote is in the set of valid votes. In that case, invalid votes could be removed before they are tallied and published.

It would be difficult to force a voter to surrender the credentials as that would require getting physical access to the voter HSM. However, the possibility can not be excluded as it is not clear if the correct PIN code could be extracted from the voter or HSM. It might be possible to try out all PIN code combinations in order to give a valid vote. It is also not specified in [29] if the HSM would allow to change the valid PIN codes. A successful change of the PIN would probably reveal the real PIN code. If changing PIN codes is not possible, then the usability aspect of the HSM would come under question. Even if the coercer could use the HSM, the attack would not scale well.

4.6 Selections

It is not clear whether Selections is resistant to the forced abstention attack. While the votes are cast over an anonymous channel and the passwords are re-randomised, there are some questions that can not be answered based on the

protocol description. First, the protocol allows to revoke voter registration before pre-tallying, but it is not specified how it could be implemented. The authors of Selections also state that the revocation process might not be covered by coercion resistance. Second, during the registration, the randomised encryption of the password is posted to the roster along with the VoterID. However, it is not stated what the VoterID is or how it is assigned to the voters. Thus, the coercer might be able to use the VoterID to check if the coerced voter registered to use Selections. Third, it is assumed that during the registration process, the voter does not copy or remember the randomisation of the selected password. However, modern technology makes it quite easy to copy and broadcast information. Rewinding some of the registration steps would not help in case the coercer forces the voter to live broadcast the process.

The protocol does not specify the way how the vote is represented or how the votes are tallied. Thus, the possibility of casting an invalid vote depends on the specific implementation of the protocol.

If the coercer would like to get access to the valid credentials, the voter would have to record or broadcast the registration process. However, in that case the voter could revoke the registration before pre-tallying and thus invalidate the credentials given to the coercer together with the vote. After revoking, the voter could go to the polling station to vote in person.

5 Conclusions and further work

Developing a voting protocol to meet the requirements of a given jurisdiction is a complex task. On one hand, we would like the protocol to be secure against all critical attacks, but this security comes with a price of increased implementation complexity and technical assumptions that need to be satisfied.

This paper focused on coercion-resistance properties of various voting protocols proposed in academic literature from the practical system developer viewpoint. As academic proposals are not required to include real-life deployments, it is very easy to leave some of the implementation details out of consideration. Unfortunately, there are many devils hidden in these details.

During our research we identified six main (groups of) popular technical assumptions. Some of them (like existence of PKI or ability to cast a re-vote) indeed have readily accessible practical instantiations. At the same time, the requirements to set up anonymous channels or distribute special-purpose client hardware are easy to write down on paper, but quite tricky to implement.

Subliminal hinting using fake credentials is one of the oldest methods to achieve provable coercion-resistance properties, but a recent usability study by Neto *et al.* [26] found that more than 90% of the test participants did not understand this functionality. This questions the whole idea of using fake credentials.

In general, there is a lack of usability studies that focus on the coercion-resistance aspects of voting protocols. We see this as an important open question that requires further research.

Another general shortcoming of the current proposals is under-specification. On several occasions, it was impossible to determine susceptibility to certain

attacks as this would have depended on specific implementation aspects. Sure, a 16-page academic paper can not fit all the details, but we encourage future scholars to accompany their proposals with deployed implementations. This would help identifying potential problems in an earlier stage of academic discussion.

Acknowledgments The research leading to these results has received funding from the Estonian Research Council under Institutional Research Grant IUT27-1 and the European Regional Development Fund through the Estonian Centre of Excellence in ICT Research (EXCITE) and the grant number EU48684.

References

1. Adida, B.: Helios: Web-based Open-Audit Voting. In: Proceedings of the 17th USENIX Security Symposium. pp. 335–348. USENIX Association (2008)
2. Araújo, R., Rajeb, N.B., Robbana, R., Traoré, J., Youf, S.: Towards practical and secure coercion-resistant electronic elections. In: CANS 2010, Proceedings. LNCS, vol. 6467, pp. 278–297. Springer (2010)
3. Benaloh, J.: Rethinking voter coercion: The realities imposed by technology. *USENIX Journal of Election Technology and Systems (JETTS)* 1, 82–87 (2013)
4. Chaidos, P., Cortier, V., Fuchsbauer, G., Galindo, D.: BeleniosRF: A Non-interactive Receipt-Free Electronic Voting Scheme. In: Proceedings of 2016 ACM CCS. pp. 1614–1625. ACM, New York, NY, USA (2016)
5. Clark, J., Hengartner, U.: Panic Passwords: Authenticating under Duress. In: HotSec’08, Proceedings. USENIX Association (2008), http://www.usenix.org/events/hotsec08/tech/full_papers/clark/clark.pdf
6. Clark, J., Hengartner, U.: Selections: Internet Voting with Over-the-Shoulder Coercion-Resistance. In: Danezis, G. (ed.) FC 2011, Revised Selected Papers. LNCS, vol. 7035, pp. 47–61. Springer (2011)
7. Clark, J., Hengartner, U.: Selections: Internet voting with over-the-shoulder coercion-resistance. *Cryptology ePrint Archive*, Report 2011/166 (2011), <https://eprint.iacr.org/2011/166>
8. Clarkson, M.R., Chong, S., Myers, A.C.: Civitas: Toward a Secure Voting System. In: 2008 IEEE Symposium on Security and Privacy (S&P 2008). pp. 354–368. IEEE Computer Society (2008)
9. Distler, V., Zollinger, M.L., Lallemand, C., Roenne, P.B., Ryan, P.Y.A., Koenig, V.: Security – Visible, Yet Unseen? In: CHI 2019. pp. 605:1–605:13. ACM, New York, NY, USA (2019)
10. Gerck, E., Neff, C.A., Rivest, R.L., Rubin, A.D., Yung, M.: The Business of Electronic Voting. In: FC 2001, Proceedings. LNCS, vol. 2339, pp. 234–259. Springer (2001)
11. Hoffman, L.J., Cranor, L.F.: Internet voting for public officials: introduction. *Commun. ACM* 44(1), 69–71 (2001)
12. Iovino, V., Rial, A., Rønne, P.B., Ryan, P.Y.A.: Using Selene to Verify Your Vote in JCJ. In: Financial Cryptography and Data Security. LNCS, vol. 10323, pp. 385–403. Springer (2017)
13. Jefferson, D.R., Rubin, A.D., Simons, B., Wagner, D.A.: Analyzing internet voting security. *Commun. ACM* 47(10), 59–64 (2004)
14. Joaquim, R., Ribeiro, C., Ferreira, P.: Improving Remote Voting Security with CodeVoting. In: Towards Trustworthy Elections, New Directions in Electronic Voting. LNCS, vol. 6000, pp. 310–329. Springer (2010)

15. Juels, A., Catalano, D., Jakobsson, M.: Coercion-Resistant Electronic Elections. Cryptology ePrint Archive, Report 2002/165 (2002), <https://eprint.iacr.org/2002/165>
16. Juels, A., Catalano, D., Jakobsson, M.: Coercion-resistant electronic elections. In: Proceedings of WPES 2005. pp. 61–70. ACM (2005)
17. Krips, K., Willemsen, J., Väriv, S.: Implementing an Audio Side Channel for Paper Voting. In: E-Vote-ID 2018, Proceedings. LNCS, vol. 11143, pp. 132–145. Springer (2018)
18. Krips, K., Willemsen, J., Väriv, S.: Is your vote overheard? A new scalable side-channel attack against paper voting. In: Proceedings of Euro S&P 2019. pp. 621–634. IEEE (2019)
19. Kulyk, O.: Extending the Helios Internet Voting Scheme Towards New Election Settings. Ph.D. thesis, Technische Universität Darmstadt (2017)
20. Kulyk, O., Teague, V., Volkamer, M.: Extending Helios Towards Private Eligibility Verifiability. In: VoteID 2015, Proceedings. LNCS, vol. 9269, pp. 57–73. Springer (2015)
21. Locher, P., Haenni, R.: Receipt-free remote electronic elections with everlasting privacy. *Annals of Telecommunications* 71(7) (Aug 2016)
22. Madise, Ü., Martens, T.: E-voting in Estonia 2005. The first Practice of Country-wide binding Internet Voting in the World. In: Krimmer, R. (ed.) *Electronic Voting 2006*. LNI, vol. 86, pp. 15–26. GI (2006)
23. Madise, Ü., Vinkel, P.: Internet voting in Estonia: from constitutional debate to evaluation of experience over six elections. In: *Regulating eTechnologies in the European Union. Normative Realities and Trends*, pp. 53–72. Springer (2014)
24. Mitrou, L., Gritzalis, D., Katsikas, S.K.: Revisiting Legal and Regulatory Requirements for Secure E-Voting. In: SEC2002. IFIP Conference Proceedings, vol. 214, pp. 469–480. Kluwer (2002)
25. Mohen, J., Glidden, J.: The case for internet voting. *Commun. ACM* 44(1), 72–85 (2001)
26. Neto, A.S., Leite, M., Araújo, R., Mota, M.P., Neto, N.C.S., Traoré, J.: Usability Considerations For Coercion-Resistant Election Systems. In: Proceedings of the 17th Brazilian Symposium on Human Factors in Computing Systems. pp. 40:1–40:10. IHC 2018 (2018)
27. Neumann, S., Feier, C., Volkamer, M., Koenig, R.: Towards A Practical JCJ / Civitas Implementation. Cryptology ePrint Archive, Report 2013/464 (2013), <https://eprint.iacr.org/2013/464>
28. Neumann, S., Volkamer, M.: Civitas and the Real World: Problems and Solutions from a Practical Point of View. In: ARES 2012. pp. 180–185. IEEE (2012)
29. Patachi, S., Schürmann, C.: Eos a universal verifiable and coercion resistant voting protocol. In: E-Vote-ID 2017, Proceedings. LNCS, vol. 10615, pp. 210–227. Springer (2017)
30. Richter, T., Escher, S., Schönfeld, D., Strufe, T.: Forensic analysis and anonymisation of printed documents. In: Proceedings of IH&MMSec '18. pp. 127–138. ACM, New York, NY, USA (2018)
31. Ryan, P.Y.A., Rønne, P.B., Iovino, V.: Selene: Voting with Transparent Verifiability and Coercion-Mitigation. In: FC 2016 International Workshops, Revised Selected Papers. LNCS, vol. 9604, pp. 176–192. Springer (2016)
32. Willemsen, J.: Bits or paper: Which should get to carry your vote? *Journal of Information Security and Applications* 38, 124–131 (2018)

PhD Colloquium

How can Internet Voting be implemented in Portuguese elections? A comparison with Estonia

Marlon Freire^{1 2}[0000-0003-4600-0746]

¹ Faculty of Engineering of University of Porto, Porto, Portugal

² Tallinn University of Technology, Tallinn, Estonia

marlonfreirephd@gmail.com

1 Introduction

In 2005, a test pilot of I-voting was introduced in the general elections for the Portuguese Parliament. The I-voting experiment was aimed at all citizens who were allowed to vote abroad using postal vote. From a total of 148,159 electors outside Portugal who were registered, only 4,367 voted through the Internet (12% of mailed votes) [4]. The experiment was voluntary and not valid for official elections results. Although the voters who used I-voting had a good experience (around 98%) [4], the Portuguese government decided to stop the I-voting experiments. However, in the same year, Estonia became the first country in the world to successfully introduce I-voting as an additional voting channel during Local Elections. In 2007, a new experience was made in Estonian Parliament Elections resulting in almost 30,000 voters used the I-voting method to cast their vote [5]. The Estonian I-voting system rests on two key elements: the transformation of the electoral law and the use in elections of the potential of its e-Governance ecosystem components and infrastructure – e-ID, Mobile-ID, Population Register and X-Road. The share of Internet voters grew from a mere 2% in 2005 to approximately 44% in 2019 [1], indicating the consolidation of the Estonian I-voting system.

Given this situation, the Estonian I-voting system seems to be a role model where other electoral systems can mirror, and it is relevant to determine which factors and aspects stemming from the Estonian case could be useful. In accordance with that, this research aims to build guidelines for the introduction of I-voting in Portuguese elections based on the experience of Estonia. For doing so, two dimensions will be considered: the legal and the digital administrative transformations implemented by the Estonian government. And, to shape it as research questions:

1. What adjustments would be necessary in the Portuguese electoral law to introduce a system as the Estonian I-voting model?
2. Are the Portuguese e-Governance ecosystem components and infrastructure enough to implement an I-voting system?

2 Research Methodology and Implementation

For the first question we will apply Business Process Model and Notation (BPMN¹) to model the Portuguese and Estonian electoral law's internal processes related to I-voting. It will allow to understand both electoral systems, as well as analyzing potential connections between them [3]. BPMN will allow to identify points where changes could be introduced in the Portuguese electoral system and understand its impacts in adding an I-voting channel in Portuguese elections. Additionally, we will use Unified Modelling Language (UML²) to model (through use cases, activity and state diagrams) the processes of the Estonian I-voting system in order to analyze whether what is written in the electoral law is what actually happens in practice. Both methods can provide detailed description and better understanding how the Estonian Internet voting activity proceeds in several aspects - technologically and administratively.

The second question is focused on the analysis of whether the Portuguese ecosystem is able to implement an I-voting system as the Estonian. It will be addressed by developing a case study on Estonia and Portugal, comparing the current situation and development of both e-Governance models paying attention to the technologic and administrative ecosystems components and infrastructures. Secondly, the UML and BPMN models will allow to connect the I-voting system with the e-Governance frames and to compare the Estonian and Portuguese cases. When the models will be done, we will conduct meetings and interviews with e-Governance experts connected to Estonian and Portuguese government and Academia in order to discuss the findings. The final outcome will be an analysis and model of the existing similarities and differences (strengths and weaknesses) between the Estonian and Portuguese e-Governance ecosystems, allowing to identify the necessary steps for the implementation of an I-voting system in Portugal.

References

1. Eesti Valimised – Estonian Elections Results, <https://ep2019.valimised.ee/en/voting-result/index.html>, last accessed 2019/06/11.
2. Henrik, C., Aino, C., Klaus, M. H.: An Approach to Software Architecture Description Using UML (2004).
3. Krimmer, R., Dueñas-Cid, D., Krivonosova, I., Vinkel, P., Koitmae, A.: How Much Does an e-Vote Cost? Cost Comparison per Vote in Multichannel Elections in Estonia (2018).
4. Piteira, S. R.: *Projecto Voto Electrónico, Voto Electrónico e Defesa da Privacidade* Workshop (Electronic Voting and Privacy Protection Workshop) (2006).
5. Solvak, M., Vassil, K.: E-voting in Estonia: Technological Diffusion and Other Developments Over Ten Years (2005 - 2015). Johan Skytte Institute of Political Studies, p. 244 (2016).

¹ BPMN is a model-driven approach for organizational engineering in which user interfaces are derived from business processes [3].

² UML is an approach to documenting and modelling systems based on diagrammatic representations and describing software architectures, components and processes [2].

How elections with Internet voting are administered? The case of the 2019 Parliamentary elections in Estonia

Iuliia Krivonosova¹[0000-0001-7246-1373]

¹Tallinn University of Technology, Tallinn, Estonia
iuliia.krivonosova@ttu.ee

1 Introduction

With 2019 Parliamentary elections, Estonia celebrates 15 years of continuous usage of Internet voting in all binding elections and reaches a new record, both in relative and absolute numbers. However, while Estonia frequently serves as a benchmark in Internet voting for other countries, there is lack of research on how elections with Internet voting are administered. Hence, this research focuses on administration of complex elections where multiple voting channels are available, and one of them is Internet voting. The issue of election administration has both practical and theoretical implications. For the theoretical implications, election administration is one of the most under-researched fields of public administration [6], and on the example of Internet voting administration, a wide array of issues such as coordination problem and principal-agent problem can be studied. This research also adds to the literature on convenience voting and usage of new voting technologies. As for the practical implications, this research might help election administrations to organize elections more securely, efficiently and transparently.

2 Methodology

This study builds on the methodology developed in [4]. This methodology reflects the multidisciplinary approach as it embraces theories of election administration, business process management and business process re-engineering. As for the case-study, the 2019 Parliamentary elections in Estonia are selected, for the reason defined in [4], as well as being the first case of elections where Internet voting has become the most used channel among the 10 available. The methods used for this case-study include electoral legislation analysis, supported by direct observation of electoral processes, interviews with stakeholders, electoral data analysis, procurement contracts analysis, and mapping of processes and actors (mostly with BPMN software).

3 Preliminary findings

- A trend for **outsourcing** electoral tasks to the private sector: besides well-documented involvement of Internet voting system vendors, we also mapped such actors

as individual contractors, ICT companies, construction companies, supermarkets, and private post offices. While outsourcing electoral tasks to the private sector is a global trend, the implementation of this approach in Estonia results in decrease of transparency and accountability, as most of the contracts with private actors are classified, making the content of contracts, including price paid for a service, not accessible. The contracts regarding Internet voting are sometimes classified even when concluded between public bodies. These findings go in contrast with the policy of “aggressive openness” [5] which Estonia presumably applies to delivery of Internet voting by publishing source code and allowing Internet voting public observation;

- **a multi-stakeholder approach** is frequently applied to electoral activities: one electoral activity is delivered by a wide range of actors. However, this approach results in some resource-consuming and prone to mistake tasks being delegated to the lowest level of election administration, without allocation of required resources to deliver the task properly. Overall, it increases the complexity of election administration, creating coordination problems;
- **significant regional disparities** in salaries, equipment and working tasks in local election administrations, revealed through stakeholders’ interviews and observation of electoral processes across the country. Such disparities might provoke security threats.

Overall, the study of the 2019 Parliamentary elections in Estonia demonstrates: (1) the growing number of actors involved into delivery of elections, at least partially caused by a trend for outsourcing electoral tasks to the private sector; (2) a multi-stakeholder approach combined with the delegation of electoral tasks to the lowest level of election administration; (3) regional disparities in financing elections.

References

1. Alvarez, R. M., Hall, T. E.: *Point, Click, and Vote: The Future of Internet Voting*. Brookings Institution Press (2003).
2. Vassil, K., Solvak, M., Vinkel, P., Trechsel, A. H., Alvarez, R. M.: The diffusion of internet voting. Usage patterns of internet voting in Estonia between 2005 and 2015. *Government Information Quarterly*, 33(3), 453-459 (2016).
3. Solvak, M., Vassil, K.: Could Internet Voting Halt Declining Electoral Turnout? New Evidence That E-Voting Is Habit Forming. *Policy & Internet*, 10(1), 4-21 (2018).
4. Krimmer, R., Duenas-Cid, D., Krivososova, I., Vinkel, P., & Koitmae, A.: How much does an e-Vote cost? Cost comparison per vote in multichannel elections in Estonia. In *International Joint Conference on Electronic Voting* (pp. 117-131). Springer, Cham (2018).
5. Past, L.: All Elections are Hackable: Scalable Lessons from Secure I-Voting and Global Election Hacks. *European Cybersecurity Journal*, 34-47 (2017).
6. Clark, A.: Public Administration and the Integrity of the Electoral Process in British Elections. *Public Administration* 93(1): 86–102 (2015).

Modelling Strategic Capabilities in Tamarin: Pros and Cons

Damian Kurpiewski

Institute of Computer Science Polish Academy of Sciences
d.kurpiewski@ipipan.waw.pl

1 Modelling Voting Protocols

Verification of voting protocols is deep and interesting topic. A lot of work has been done in this field and a number of properties have been formalized, including ballot privacy and receipt freeness [2], coercion-resistance [4], and voter-verifiability [6].

Verifying strategic properties. When verifying a voting protocol, a standard approach is to assume that all the participants (voters, election authority, tellers) follows the protocol rules. The question answered in that process is "Is the protocol secure?". The question we want to ask, is "Does the voter has a strategy to...". For example, "Does the voter has a strategy to verify his vote". More than that, maybe coercer working together with the voter will have a strategy to change the results of the votes. To this end we want to use ATL [1] for modelling and specifying the desired properties.

Modelling Voting Protocols. Tamarin is a theorem prover commonly used for verification of security protocols. It is as tool one can use to model cryptography in the protocol. Voting protocols often involve a lot of cryptography, so does that mean we can easily use Tamarin for the verification of the voting protocols? While similar in some respects, security and voting protocols have some differences between them. The main difference lies in the agents actions. In the security protocol agents just follow the protocol rules. There is no place for choices. In the voting protocol however, voter can make many choices: which candidate to vote for, when to vote, to verify his vote or not.

2 Pros and cons of using Tamarin

The first question to ask would be is it a good decision to use Tamarin for the verification of strategic properties in the voting protocol. While there exists some model-checkers created exactly for that purpose, like MCMAS [5], it is problematic to model cryptography in the standard model-checking techniques. Of course, the cryptography in the model can be simplified or even omitted, but by doing this we may lose some important information.

Modelling Voting Protocols. Depending on the protocol, it may involve many different cryptographic operations and have many stages, like ballot mixing, distribution, keys preparation and more. At the same time, interaction of the agents with the protocol is often simple: get your ballot, fill it with your vote, post it and then verify it on

public bulletin board. Compared to other ATL models, voting protocols have different structure, more like that of a security protocol. One of the pros of using Tamarin for modelling of the voting protocol is an easy to understand specification language and its compactness. Moreover, it is relatively simple to model cryptography in Tamarin, as well as coercer in the form of a built-in Dolev-Yao attacker. The issue here is to remember, that there are no explicit agents in Tamarin, although one could create the specification in such a way to simulate agents in the protocol.

Observational Equivalence. When verifying properties like vote-privacy in tamarin, the standard approach would be to use observational equivalence, like in [3]. The problem here is that Tamarin uses much stronger notion of observational equivalence, than one commonly used in model-checking. Because of that, when using Tamarin observational equivalence, one should very carefully create the model of the verified protocol. Even potentially significant changes may have a great impact on the result. If possible, it would be better to use trace lemmas.

Information flow and knowledge. In the Tamarin information passed in the protocol evolution plays crucial role. Any information send on the insecure channel can be intercepted by the intruder and added to his knowledgebase. Using this information intruder can compute new data and use it for encryption, decryption, or to send some messages. In Tamarin, when we say that intruder knows something, it means that he can learn or compute it. While working on ATL models, when we say that agents know some property p , it means that p is true in every state indistinguishable for that agent. Furthermore, in the voting protocol, we are often interested in asking if it is possible for the coercer to learn how the voter has voted. But how to specify some-thing like that in Tamarin? We can ask if coercer knows a pair consisting of the voter id and some candidate id, but unless voters ids and candidate ids are private, coercer can compute exactly that pair. That of course doesn't mean that he knows what that pair of data means, but he can compute it. One should be very careful when asking about knowledge in Tamarin.

3 References

1. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time Temporal Logic. In: Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS). pp. 100{109. IEEE Computer Society Press (1997)
2. Benaloh, J., Tuinstra, D.: Receipt-free secret-ballot elections. In: Proceedings of the twenty-sixth annual ACM symposium on Theory of Computing. pp. 544{553. ACM (1994)
3. Bruni, A., Drewsen, E., Schürmann, C.: Towards a mechanized proof of selene receipt-freeness and vote-privacy. In: Proceedings of E-Vote-ID. pp. 110{126 (2017). https://doi.org/10.1007/978-3-319-68687-5_7
4. Juels, A., Catalano, D., Jakobsson, M.: Coercion-resistant electronic elections. In: Proceedings of the 2005 ACM workshop on Privacy in the electronic society. pp. 61{70. ACM (2005)
5. Lomuscio, A., Qu, H., Raimondi, F.: MCMAS: An open-source model checker for the verification of multi-agent systems. International Journal on Software Tools for Technology Transfer 19(1), 9{30 (2017). <https://doi.org/10.1007/s10009-015-0378-x>
6. Ryan, P.Y.A., Schneider, S.A., Teague, V.: End-to-end verifiability in voting systems, from theory to practice. IEEE Security & Privacy 13(3), 59{62 (2015). <https://doi.org/10.1109/MSP.2015.54>

Voters' Understanding of the Coercion Mitigation Mechanism in Selene

Marie-Laure Zollinger, University of Luxembourg

1 Introduction

One specificity of the Selene voting protocol is to propose a coercion mitigation mechanism to allow any voter to fake the tracking number used to verify their vote. This tracking number is kept secured until the end of the election, and is leading to the plaintext vote of the voter. While the impact of the verification phase has already been explored in a previous user study [1], we want to explore the ability of a voter to fake his tracking number in case of coercion. Furthermore, from the previous study, many participants have provided feedback regarding the verification process and their lack of understanding. In particular, seeing the entire bulletin board while they could just see their own vote was troubling. By providing a complete experience to the voter, we want to evaluate again the understanding and the usability of Selene.

2 Methodology

To evaluate the understanding and the usability of Selene, we design a user game inspired by game theoretic experiment in a similar way than Llewellyn et al. in [2]. The following steps are a first sketch of the possible methodology for a user test study. We will run the study online using a crowd-sourcing website to recruit participants. They will receive 3 tasks to perform described in the 3 phases below.

Preparation In order to have the participants understanding the next phases, we need to start the study with an explanation of the Selene protocol. The tracking number retrieval results from the combination of a public commitment, a secret alpha-term, and a trapdoor key owned by the voter. The secret is sent at the end of the election to the voter, who is able to combine and retrieve the tracking number. To fake a tracking number, a new secret can be computed by the participant, from the public commitment, the trapdoor key and a tracking number of the participant's choice. To achieve this, the provided application will have a feature that will perform the computation automatically.

Phase I: Vote and Verify The participants vote for a favorite candidate using a browser. The scenario explains that they are able to verify the vote after the election is over¹. Then, the participants connect to the Bulletin Board where the pairs (tracking number, vote) have been published. They have the possibility to

¹ All those steps are performed at once for convenience. Indeed, even if we can ask the same participants again, we cannot ensure that they will continue the test.

ask for a fake tracking number in case they are coerced. The provided application has a feature helping the voter to fake the tracking number, by letting him/her choose a new tracking number to be notified with. In the scenario, they are invited to explore the features. Finally, the participants are invited to verify their vote.

Phase II: User Experience Questionnaire Once the test of the application is done, participants are redirected on a new page where the user experience questionnaire (UEQ) [3] is displayed. It must take 3 to 5 minutes to fill out.

Phase III: Vote-buying Game Finally, after the UEQ, we redirect the participant on a platform to play a vote-buying game. The participants will be given randomly the role of buyer or voter. As a buyer, he can try to buy the vote of the voter, knowing that the voter can fake his tracking number. As a voter, when the buyer asks for the tracking number, the participant can choose between giving a fake or her real tracker. A last question asking why the buyer/voter chose or not to buy/fake the vote will be asked in order to evaluate their understanding.

3 Data analysis

Understanding of the Bulletin Board Unlike the previous study, the participants will receive explanation on why they can see other tracking numbers. They will also have the ability to choose a tracker to be notified with, instead of their real tracking number, in case of coercion.

Understanding of/Trust in the security behind the tracking number Participants will have to play a game where they can have either the role of the vote buyer, or the role of the voter who fake his tracker. Then we ask them to explain their choice in the game.

User experience The UEQ will help us to evaluate the users' perceptions of the voting system and verify the assumption that the voting system fits their needs.

4 Work in progress

This paper is a first user study design that needs to be detailed. The feature for faking the tracker has to be designed and developed and the utility functions for the game must be described.

References

1. Distler, V., Zollinger, M., Lallemand, C., Rønne, P.B., Ryan, P.Y.A., Koenig, V.: Security - visible, yet unseen? In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (2019)
2. Llewellyn, M., Schneider, S., Xia, Z., Culnane, C., Heather, J., Ryan, P.Y.A., Srinivasan, S.: Testing voters' understanding of a security mechanism used in verifiable voting. In: 2013 Electronic Voting Technology Workshop / Workshop on Trustworthy Elections (2013)
3. Schrepp, M.: <https://www.ueq-online.org/> (2018)

DEMO Session

Features of Polys Blockchain Voting: A Survey

Roman Alyoskin

Polys.me

roman.alynoskin@polys.me

Abstract. Voting is the cornerstone of any democracy. Today the vast majority of voting is still done offline. Such a voting process is inefficient both for organizers and voters who have to contend with barriers to registration, long lines at the voting booth, difficulties in establishing proof of identity and flawed counting practices. And this is the case for elections at all levels, from universities and political parties to municipal and national governments. Blockchain technology has the potential to solve those problems and change the way people vote. Polys, a blockchain-based voting system, has emerged as a tool claiming to increase the efficiency, transparency, and trustworthiness of the voting process. Here we will consider the technology at the core of Polys, its capabilities as well as real-life examples of its use

Keywords: Blockchain, e-Voting, Online Election, Transparency, Polys.

1 Introduction

Online voting imposes stringent security requirements on the voting process because, in serious elections, the temptation of large-scale manipulation is too great and the risk simply too high. That's why such elections cannot be carried out using black-box voting – the process must be clear and transparent for all participants. It should always be possible to monitor the process and easily verify its results. Blockchain technology offers features that can meet these requirements.

What is Blockchain?

A blockchain can be thought of as a distributed database. Unlike regular servers where data can be hacked or manipulated, data in a blockchain is stored in information blocks on the computers of all the network participants. This means it's impossible to hack because to do so you would have to hack all the machines in a network. It is also impossible to steal the data because it's encrypted. Transactions in a blockchain are immutable, so it's impossible to change them once they are placed in a block. That's what makes blockchains so transparent and trustworthy.

What is Polys?

Polys is the blockchain-based voting system. To maintain the anonymity of a vote and ensure it's not possible to find out how people voted, Polys uses crypto algorithms to

encrypt preliminary results and voter identity. Imagine a black box where all the votes are collected.

2 Definition

Polys comprises an application for election deployment and a client application for casting votes. It's flexible and scalable in terms of use cases and capable of conducting elections at all levels – from small community votes to countrywide elections. Voters can vote online via desktop computers or smartphones. Another key feature of the Polys system is the availability of voting machines. Obviously, the switch from paper to online voting cannot happen instantly, but to make the transition as smooth as possible, Polys has designed voting machines that are also based on blockchain technology.

3 Implementation of Polys Voting

In May 2019, 82,500 residents of the Volgograd region in southern Russia chose which projects would receive funding from the municipal budget via a blockchain-based online voting platform. This was one of the largest votes ever conducted using blockchain technology (as of August 2019). This showed that people are much more willing to vote on important public issues if it doesn't require going to a physical location. It's far more convenient for them to contribute with the help of modern technologies.

4 Conclusion

In this paper, we have presented some of the main features of the Polys blockchain voting system. Of course, the introduction of blockchain platforms to public voting systems will require time and a well-thought-out approach. It's necessary to study all the weaknesses and prepare a comprehensive legislative framework that determines the status of the new technology. Nevertheless, electronic voting on a blockchain has enormous potential and is capable of increasing the reliability, security, and transparency of elections, while simplifying the voting process itself.

5 References

1. Polys — Online voting system, Whitepaper. [online] Available: https://polys.me/assets/docs/Polys_whitepaper.pdf
2. Preethi Kasireddy, "How does ethereum work anyway", [online] Available: <https://medium.com/@preethikasireddy/how-does-ethereum-work-anyway22d1df506369>
3. Roman Alyoshkin "Blockchain 2.0. The Purpose of Blockchain", [online] Available: <https://medium.com/polys-blog/blockchain-2-0-the-purpose-of-blockchain-e84e5a95cdd9>

Unsupervised electronic voting machines and methods

Promitheas Christophides

myevoter.com
promitheas@christophides.com

These electronic voting machines and methods to be presented at the Fourth International Joint Conference on Electronic Voting 1 – 4 October 2019 come in two versions so as to address all possible needs of voters with or without disabilities. Specifically, version (A) also addresses the needs of voters visually impaired, while version (B) also addresses the needs of voters with hand impairments. Both versions address also the needs for voters with mobility issues and the needs of voters with reading disabilities such as dyslexia or voters with functional illiteracy. Both versions are intended for use remotely via internet or in kiosks in uncontrolled environments, not excluding controlled environments like polling stations during a transitional period when moving from traditional ballot systems to uncontrolled electronic environments.

In order to accommodate voters with disabilities mentioned above, each version uses different verifiability methods as described below. Both versions, however, employ also human body verifiability to safeguard voter privacy and voting secrecy and satisfy in every aspect the Secret Ballot (Australian Ballot) requirements. Both versions are presently in the fully working prototype stage and have been privately tested in uncontrolled environments involving close to 300 volunteer voters and beta testers.

More information on both versions can be obtained at www.myevoter.com, via email request at info@myevoter.com.

1 (A) Fingerprint authentication (touch ID) version

Version (A) relates to a device and a method associated with the device. With respect to the device, it is an electronic mobile device which allows voters to vote in an election or poll or survey privately and secretly. This device can be used by voters to enter electronically their votes and submit them online to a central database. This can be done even in the presence of other persons without risk of breach of secrecy and voters will have no way of providing proof to any other person of how and what they voted.

The core components of version (A) are an electronic device which can be used as desktop and/or mobile device comprising of a single board computer, a numbers keyboard, a fingerprint sensor, human body recognition sensors to safeguard from attempts to fraud the system, a monitor and a flash drive (USB stick) which stores the operating system required by the computer to function and required scripts for the device to operate as a voting machine and which, generally speaking, are configured as follows: Voters have their own flash drive which stores securely the operating system and scripts, an ID number and their fingerprint images. By inserting their personal flash drive in the computer's respective port and powering on the computer, voters can use

fingerprint authentication (touch ID) to go through a voting process on an online form or forms hosted on a central server.

With respect to the device it should be further noted that the numbers keyboard has essentially only 6 buttons which are numbers 1, 2 and 3, a delete button, a reset button and an enter/submit button. Essentially the 3 numbers button and the submit button are hidden from view inside a hood attached to the computer permanently. With respect to the associated method, in order to carry out the method the following core steps are followed: When voters successfully get access to the voting page after authentication they can vote following on-screen or voice instructions by means of the 3 numbers keyboard entering numbers which consist of any configuration that can be achieved using the said three numbers. Each of these number configurations will correspond to an election candidate, political party, coalition or individual or to specific answers in polls and surveys and are represented on screen by symbols such as dots or asterisks as in password fields.

To safeguard voting secrecy, the embedded human body sensors will interrupt the voting process and ask for re-authentication if voters decide to withdraw their hand before pressing the submit key.

Version (A) uses an environment very familiar to visually impaired voters who can vote unsupervised in an uncontrolled environment following voice instructions.

2 (B) Iris authentication headset version

The core components of t Version (B) are an electronic device housed in a headset which can be used as a mobile device comprising of a single board computer, a wireless computer trackball mouse, an iris recognition sensor, a human body temperature sensor to safeguard from attempts to fraud the system, a TFT screen, a magnifying lens, hearing aid such as earphone and a USB flash drive (memory stick) which stores the operating system required by the computer to operate and required scripts for the device to operate as a voting machine and which, generally speaking, are configured as follows: Voters have their own USB flash drive which stores securely the operating system and scripts, their ID number and their iris images.

It should be further noted that one of the two headset compartments, is reserved for the iris recognition camera sensor and the human body temperature sensor, while the other compartment points to a TFT screen observed by the voter via a 3X magnification lens. In order to carry out the method the following core steps are followed: When voters successfully get access to the online voting page after iris authentication, they can vote by means of navigating through the page or pages via the supplied wireless trackball mouse and make their selections onscreen from a list of names and/or images of candidates. Subsequently they are asked to click on a submit button to register their vote.

To safeguard voting secrecy, the embedded temperature sensor will shut down the computer if the voter decides to take off the headset before pressing the submit button.

With the help of an assistant who will be operating only the trackball mouse, voters with impaired hands or other mobility issues can vote in secrecy simply by directing the assistant to navigate the mouse cursor over their preferences.

Verifiability and security of Scytl's online voting system

Scytl Secure Electronic Voting
08008 Barcelona, Spain
www.scytl.com

Scytl's online voting system helped several governments on the introduction of verifiability in online voting for political elections. Starting from 2004 in Switzerland (Neuchâtel) and continuing in France (2009) and Canada (2014), Scytl's voting system included voter verifiability based on voting receipts¹, allowing voters to check that their votes were present in the final tally. In Norway, in 2011 municipal elections, Scytl's online voting system made major step by introducing individual verifiability for the first time in a national election. The Norwegian online voting system individual verifiability was based on using return codes². It also implemented counted-as-recorded verifiability based on using universal verifiable Mix-nets^{3,4}. In 2015, Scytl's voting system introduced a second verification mechanism designed for the State of New South Wales (Australia), based on a cast and decrypt approach (decryption of the vote in a trusted environment accessible by phone)⁵. Furthermore, in 2015, Scytl's individual verifiability based on return codes was adopted in Switzerland (Neuchâtel) and achieved in 2017 the Swiss certification for individual verifiable systems⁶. Scytl online voting system is currently in the process of achieving the Swiss complete verifiability certification level. At the beginning of 2019, Scytl's Swiss voting system went through the last phase of the process and went through a public scrutiny process⁷. After solving major found by researchers during this process, it is expected that the voting system will be ready for achieving the complete verifiability certification.

¹ Puiggalí, J., Morales-Rocha, V.: Independent voter verifiability for remote electronic voting. In: Proceedings of International Conference on Security and Cryptography (SECRYPT '07), pp. 333–336, Barcelona (2007).

² Puiggalí, J., Guasch, S.: Universally verifiable efficient re-encryption mixnet. In: Electronic Voting 2010 (EVOTE 2010), 4th International Conference, LNI, vol. 167, pp. 241–254, Austria (2010).

³ Puiggalí, J., Guasch, S.: Cast-as-intended verification in Norway. In: 5th International Conference on Electronic Voting 2012, (EVOTE 2012), LNI, vol. 205, pp. 49–63, Austria (2012).

⁴ Wikström, D.: A sender verifiable mix-net and a new proof of a shuffle. In: Proceedings of the 11th International Conference on Theory and Application of Cryptology and Information Security. pp. 273–292. ASIACRYPT'05, Springer-Verlag, Berlin, Heidelberg (2005).

⁵ Brightwell, I., Cucurull, J., Galindo, D., Guasch, S. An overview of the iVote 2015 voting system. Tech. rep. New South Wales Electoral Commission (2015).

⁶ Swiss Post: Audit certificates and reports. Available at <https://www.post.ch/en/business-solutions/e-voting/publications-and-source-code#auditcertificatesandreports>, last accessed 2019/09/15

⁷ Swiss Post: Swiss Post publishes the source code for its e-voting system. Press Release February 7th, 2019. Available at <https://www.evoting-blog.ch/en/pages/2019/swiss-post-publishes-the-source-code-for-its-e-voting-system?>

Finally, Australian ScytI's online voting system has been used again in 2019 by the State of New South Wales. This time, the individual verifiability cast-and-decrypt scheme was improved allowing voters to verify their voter through a mobile phone application (instead just a phone call) and included recorded as cast verifiability using a universal verifiable Mix-net⁸.

In the demo session, ScytI will explain in the different verification mechanisms implemented by its online voting systems and which additional functionalities are supported in addition to verifiability, like homomorphic tally or blockchain integration⁹. There will be also a demo system that will show the improved verification mechanism implemented this year in the Australian election held in New South Wales.

⁸ New South Wales Electoral Commission: iVote® refresh project for the 2019 NSW State election. February 2019.

⁹ Cucurull J., Puiggalí J.: Distributed Immutabilization of Secure Logs. In: Security and Trust Management. (STM 2016). LNCS, vol 9871. Springer, Cham. Greece (2016).

Call For Papers
E-Vote-ID 2020

Fourth International Joint Conference on Electronic Voting
6–9 October 2020 · Bregenz, Austria

E-Vote-ID 2020

General Chairs: **Krimmer, Robert** (Tallinn University of Technology, Ragnar Nurkse School, Estonia), **Volkamer, Melanie** (Karlsruhe Institut für Technologie, Germany)

Outreach Chairs: **Rønne, Peter** (University of Luxembourg, Luxembourg), **Krivososova, Iuliia** (Tallinn University of Technology, Estonia)

Submission Deadline:
15 May 2020

This is the fifth edition of one of the leading international events for e-voting experts from all over the world, taking place in Bregenz (Austria) in October 2020.

One of its major objectives is to provide a forum for interdisciplinary and open discussion of all issues relating to electronic voting. In the first 4 editions, up to 121 presentations had been discussed, gathering more than 400 participants.

The aim of the conference is to bring together e-voting specialists working in academia, politics, government and industry in order to discuss various aspects of all forms of electronic voting (including, but not limited to, polling stations, kiosks, ballot scanners and remote voting by electronic means) in the four following tracks below and a PhD colloquium:

Track on Security, Usability and Technical Issues

Chairs: **Beckert, Bernhard** (Karlsruhe Institut für Technologie, Germany) **Küsters, Ralf** (University of Stuttgart, Germany) and **Oksana Kulyk** (IT University of Copenhagen)

Design, analysis, formal modeling or research implementation of:

- ❖ Electronic voting protocols and systems;
- ❖ Voter identification and authentication;
- ❖ Ballot secrecy, receipt-freeness and coercion resistance;
- ❖ Election verification including end-to-end verifiability and risk limiting audits;
- ❖ Requirements;
- ❖ Evaluation and certification, including international security standards, e.g. Common Criteria or ITSEC;
- ❖ Human aspects of security mechanisms in electronic voting and in particular of verifiability mechanisms;
- ❖ Or any other security and HCI issues relevant to electronic voting.

Track on Administrative, Legal, Political and Social Issues

Chairs: **Duenas-Cid, David** (Tallinn University of Technology, Estonia / Kozminski University, Poland), **open position**

- ❖ Discuss legal, political and social issues of electronic voting implementations, ideally employing case study methodology;
- ❖ Analyze the interrelationship with, and the effects of electronic voting on democratic institutions and processes;
- ❖ Assess the cultural impact of electronic voting on institutions, behaviours and attitudes of the Digital Era;
- ❖ Discuss the administrative, legal, political and social risks of electronic voting;
- ❖ How to draft electronic voting legislations;
- ❖ Public administrations and the implementation of electronic voting;
- ❖ Understandability, transparency, and trust issues in electronic voting;
- ❖ Data protection issues;
- ❖ Public interests vs. PPP (public private partnerships).

Track on Election and Practical Experiences

Chairs: **Oliver Spycher** (Swiss Federal Chancellery, Switzerland) and **Beata Martin-Rozumilowicz** (International Foundation for Electoral Systems, USA)

- ❖ Review developments in the area of applied electronic voting;
- ❖ Report on experiences with electronic voting or the preparation thereof (including reports on development and implementation, case law, court decisions, legislative steps, public and political debates, election outcomes, etc.);

Contributions in this track will be published in TUT press proceedings only. These experience and practical reports need not contain original research, but must be an accurate, complete and, where applicable, evidence-based account of the technology or system used. Submissions will be judged on quality of review and level of analysis, and the applicability of the results to other democracies.



Track on Posters and E-Voting System Demo

Chairs: Rønne, Peter (University of Luxembourg, Luxembourg)

We invite demonstrations of electronic voting systems, to be presented in an open session on Tuesday 1 October before the welcome reception. Participation is open to all conference participants, but we request a Short Paper (two pages) by 15 September submitted via EasyChair describing the system's requirements and properties, such as:

- whether the system is intended for use in controlled (i.e. in polling stations) or uncontrolled environments (i.e. remotely via the Internet or in kiosks);
- which types of elections it accommodates;
- whether it addresses the needs of voters with disabilities;
- what sort of verifiability it provides;
- the extent to which it guarantees vote privacy;
- whether it has been deployed in a real election;
- where to go for more information.

PhD Colloquium

Chairs: Driza Maurer, Ardita (Zentrum für Demokratie Aarau/Zurich University, Switzerland) and Iuliia Krivonosova (Tallinn University of Technology, Estonia)

The goal of the colloquium is to foster understanding and collaboration between PhD students from various disciplines working on e-voting. To this end, the program allows plenty of space for discussion and initiating collaboration based on presentations by attendees.

Each interested participant should submit his/her research proposal (or alternatively ideas for papers, open problems, or other issues where feedback from colleagues would be helpful etc.) on the shaper of a short paper (two pages length) using the conference platform.

Format of the Conference

The format of the conference is a three-day meeting. The PhD Colloquium and Demo Session take places on the day before the formal conference begins. No parallel sessions will be held, and sufficient space will be given for informal communication.

Paper Submission

Paper submissions can be in two formats—either as a full paper or an abstract.

- Full paper submissions (max 16 pages in LNCS format all-in);
- Short Paper submissions (max 2 pages in LNCS format all-in).

All submissions will be subject to double-blind reviews.

Submissions must be anonymous (with no reference to the authors). Submissions are to be made using the EasyChair conference system at <https://www.easychair.org/conferences/?conf=evoteid2020>, which serves as the online system for the review process. During submission, please select the appropriate track or the PhD colloquium. The track chairs reserve the right to re-assign papers to other tracks in case of better fit based on reviewer feedback and in

coordination with other track chairs. LNCS style has to be used (see the Springer guidelines at <http://www.springer.com/gp/computer-science/lncs/conference-proceedings-guidelines>, including templates for LaTeX and Microsoft Word).

If you think that one or more of the programme committee members could have a conflict of interest with your submission, please let the general chairs know at conference-chairs@e-vote-id.org. In turn, according settings in the EasyChair system will be set, so that the respective member/s is/are not involved in the review process.

Key Dates for Submissions

Deadline for submission of papers for the Track on Security, Usability and Technical Issues and the Track on Administrative, Legal, Political and Social Issues: (Hawaiian time, hard deadline, no extension)

15 May 2020 – 23:59

Notification of Acceptance:

24 June 2020

Deadline for submission of papers for the Track on Election and Practical Experiences and the PhD Colloquium:

10 July 2020

Deadline for Camera-ready Paper Submissions: 24 July 2020

Deadline for Poster Submission and Short Papers for E-Voting System Demo Session:

15 September 2020

Publication

The conference proceedings will be available at the time of the conference. Full papers accepted for the tracks on security, usability, and technical issues, respective administrative, legal, political, and social issues will be published in Springer LNCS.

All other accepted publications, including full papers in the election experience track, accepted abstracts in any of the tracks, and from the submissions in the PhD colloquium will be published in proceedings with TUT press.

In case your academic host institution requires you to publish your research as open-access only, please contact the conference chairs for further information in which way it is intended to make accepted publications accessible.

Venue

The conference will be held in the Renaissance castle of Hofen at Lochau/Bregenz on the shores of Lake Constance in Austria.

On the evening of 6 October a welcome reception for all conference participants will be organized in castle Hofen, where also the conference dinner on 8 October will take place and feature the traditional “cheese road”.