Report from Dagstuhl Seminar 25411

Trustworthy Evidence-Based Elections

Josh Benaloh*[1], Peter Rønne*[2], and Melanie Volkamer*[3]

1 **Microsoft Research – Redmond, US.** benaloh@microsoft.com
2 **University of Luxembourg, LU.** peter.roenne@uni.lu
3 **KIT – Karlsruher Institut für Technologie, DE.** melanie.volkamer@kit.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 25411 "Trustworthy Evidence-Based Elections". The seminar brought together 41 participants mainly from computer science, but also sociology and statistics, including members of organizations, industry, and national agencies; all experts on elections and secure voting. The five-day seminar had a full program primarily with talks presenting the latest research results, challenges, and observations including ample time for discussions. However, the presentation sessions were complemented with demo sessions and daily breakout groups continuing the most interesting discussions that emerged from the talks.

Seminar October 5–10, 2025 – https://www.dagstuhl.de/25411
2012 ACM Subject Classification Security and privacy → Cryptography; Security and privacy → Human and societal aspects of security and privacy; Applied computing → Voting / election technologies
Keywords and phrases elections, risk-limiting audits, Trust, usable security, verifiable voting, voting
Digital Object Identifier 10.4230/DagRep.15.10.1

1 Executive Summary

Josh Benaloh (Microsoft Research – Redmond, US)
Peter Rønne (University of Luxembourg, LU)
Melanie Volkamer (KIT – Karlsruher Institut für Technologie, DE)

License Creative Commons BY 4.0 International license
© Josh Benaloh, Peter Rønne, and Melanie Volkamer

This 5-day Dagstuhl Seminar 25411 was dedicated to "Trustworthy Evidence-Based Elections", a topic whose relevance and importance is witnessed by a number of democracies being under threat and elections becoming political battlegrounds in themselves rather than free and fair elections of political candidates. In this earnest backdrop, the seminar reviewed the state of art of secure elections, explored paths towards trustworthy elections, and the open challenges, in particular by sharing and elaborating on novel ideas, approaches, and use cases.

This report gives a description of the program, collects the abstracts of the talks in Sec. 3, and a selection of the breakout groups in Sec. 4.

Day 1

The Dagstuhl Seminar was opened with a round of short individual presentations, followed by a discussion about the organization of the seminar and the planning of topics and talks. All participants were allowed and encouraged to contribute with talks on recent research, election

* Editor / Organizer

Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

experiences, open problems, and ideas. In addition, it was agreed to end each day with breakout sessions and a final joint synthesis session with reports from each of the breakout groups. Due to the large number of individual talks, the breakout groups had to be limited to one-two sessions per day. Each talk was followed by active discussions.

The first day also managed to feature two sessions of talks, first a short session on Evidence-Based Elections and a longer one on Usability in voting. The breakout groups included a group on Commitment-Based Elections, Coercion-Resistance & Usability, and Education & Outreach.

Day 2

The second day had two sessions of talks on Cryptographic Voting Protocols and Risk-Limiting Audits. Furthermore, in the afternoon a demo session was held which presented an explanatory video[1] to enable informed decisions on whether or not to go for Internet voting, a demo of the Selene voting scheme, a demo on a protocol for voter enrollment in coercion-resistant elections, and a system to test paper ballots before elections in particular for logical errors. Finally, a demo on visualising the counting of individual ballots in STV was postponed until day 4.

The breakout groups included groups on Universally Verified elections and RLAs.

Day 3

The morning of day three had two sessions of talks on respectively Trust and Security Definitions. Further, it included a talk on e-collection of signatures for referendums, a topic sharing many security problems with e-voting.

In the afternoon, the participants split into two groups, one group explored the hiking trails around Schloss Dagstuhl while another group visited the Celtic Hillfort of Otzenhausen.

Day 4

Day four had three sessions: Cryptographic Voting Protocols (the second session on this topic), Verified Implementations & Formal Verification, and Trust & Elections. From the talks in these sessions, several breakout groups emerged, especially groups looking at boardroom voting, trust in voting, a group initiating the work towards the creation of an open source software developer kit (SDK) for e-voting, and trustworthy randomness-generation for election audits.

Day 5

The final day started with discussions on creating a white-paper on trustworthy evidence-based elections, the structure of this and responsible editors. In addition, two rounds of breakout groups were organized: both dedicated to new topics, following up on earlier groups and preparing the whitepaper topics. In particular, there were groups on Usability, Post-Quantum E-Voting, Formal Verification and Security Definitions.

[1] `https://secuso.org/2025-07-E-Voting-Explanation-Video/Video/paper_version/final_long_subtitles.mp4`

2 Table of Contents

3 Overview of Talks

3.1 My Journey to Commitment

Josh Benaloh (Microsoft Research – Redmond, US)

License Creative Commons BY 4.0 International license
© Josh Benaloh

Most verifiable election systems work by encrypting votes. However, replacing encryptions with commitments for in-person voting greatly simplifies the process by removing the need for keys and trustees. Using commitments also substantially reduces the size and complexity of election records – greatly simplifying the burdens on verifiers.

Sparse commitments achieve all of these benefits without substantially altering basic integrity and privacy properties or the underlying assumptions. Dense commitments can achieve ever greater efficiencies as well as everlasting privacy of votes at a cost of making the integrity of the election dependent upon a computation assumption – such as the difficulty of computing discrete logarithms at the time of the election.

3.2 Risk-limiting Audits for Complex Elections: Progress and What's Next

Michelle Blom (The University of Melbourne, AU), Alexander Ek (KU Leuven, BE), and Vanessa Teague (Australian National University – Acton, AU)

License Creative Commons BY 4.0 International license
© Michelle Blom, Alexander Ek, and Vanessa Teague

In this talk we highlight the progress that has been made in the development of risk-limiting audits (RLAs) for complex elections, and the extent to which they have been adopted by practitioners. We discuss current challenges that are limiting further progress on RLAs for preferential and proportional elections, and propose several future directions for this research.

3.3 Breaking and fixing CHVote, a protocol for the Swiss contex

Véronique Cortier (LORIA, CNRS – Nancy, FR)

License Creative Commons BY 4.0 International license
© Véronique Cortier
Joint work of Véronique Cortier, Alexandre Debant, Pierrick Gaudry
Main reference Véronique Cortier, Alexandre Debant, Pierrick Gaudry: "Breaking verifiability and vote privacy in CHVote", IACR Cryptol. ePrint Arch., p. 80, 2025.
URL https://eprint.iacr.org/2025/080

CHVote is one of the main electronic voting systems developed in the context of political elections in Switzerland, where the regulation requires a specific setting and specific trust assumptions. In this setting, voters must receive return codes to check that their device did not modify their vote. CHVote is based on oblivious transfer to achieve this functionality. We have shown that actually, CHVote fails to achieve vote secrecy and individual verifiability (here, recorded-as-intended), as soon as one of the online components is dishonest, contradicting the security claims of CHVote. The attacks have been presented at Esorics'25. An open question is to enhance the oblivious transfer scheme in order to retrieve privacy (and verifiability).

3.4 A Practical and Fully Distributed E-Voting Protocol for the Swiss Context

Alexandre Debant (INRIA – Villers-lès-Nancy, FR)

License Creative Commons BY 4.0 International license
© Alexandre Debant
Joint work of Véronique Cortier, Alexandre Debant, Olivier Esseiva, Pierrick Gaudry, Audhild Høgåsen, Chiara Spadafora
Main reference Véronique Cortier, Alexandre Debant, Olivier Esseiva, Pierrick Gaudry, Audhild Høgåsen, Chiara Spadafora: "A Practical and Fully Distributed E-Voting Protocol for the Swiss Context", IACR Cryptol. ePrint Arch., p. 1625, 2025.
URL https://eprint.iacr.org/2025/1625

Internet voting in Switzerland for political elections is strongly regulated by the Federal Chancellery (FCh). It puts a great emphasis on the individual verifiability: security against a corrupted voting device is ensured via return codes, sent by postal mail. For a long time, the FCh was accepting to trust an offline component to set up data and in particular the voting material. Today, the FCh aims at removing this strong trust assumption. We propose a protocol that abides by this new will. At the heart of our system lies a setup phase where several parties create the voting material in a distributed way, while allowing one of the parties to remain offline during the voting phase. A complication arises from the fact that the voting material has to be printed, sent by postal mail, and then used by the voter to perform several operations that are critical for security. Usability constraints are taken into account in our design, both in terms of computation complexity (linear setup and tally) and in terms of user experience (we ask the voter to type a high-entropy string only once). The security of our scheme is proved in a symbolic setting, using the ProVerif prover, for various corruption scenarios, demonstrating that it fulfills the Chancellery's requirements and sometimes goes slightly beyond them

3.5 A trust/distrust theoretical framework for internet voting

David Dueñas-Cid (Kozminski University, PL)

License Creative Commons BY 4.0 International license
© David Dueñas-Cid

Trust is a complex concept that relies on numerous factors. Technology adds complexity to its understanding by providing a different sociotechnical context; on the one hand, the technical properties of technology contribute to the perception of trustworthiness in systems, but on the other, societal determinants affect individual perceptions with elements that extend beyond the systems' technicalities. Electoral technologies do not simplify this equation. Other aspects, such as power distribution, result acceptance, political interests, or complex management, put even more pressure on trust-related issues. This presentation proposes a theoretical framework to structure the concepts of trust and distrust within the field, facilitating further research on the topic.

3.6 Beyond voting: tool building for better democracy

Bryan Ford (EPFL Lausanne, CH)

License Creative Commons BY 4.0 International license
© Bryan Ford

In this talk I propose the need to broaden our problem scope from "end-to-end voting" to "end-to-end democracy", incorporating not just the process of casting, collecting, and counting votes, but also the process of democratic agenda-setting and decision-making. In particular, there are promising scalable processes such as random-sampled citizens' assemblies or mini-publics, and liquid democracy, that show the potential to enable people to participate more directly at large scale, gather the information and diverse perspectives to make better decisions in everyone's broader interests, and leverage accountable expertise to perform deep analysis of complex issues. But to make these processes work at scale in borderless, permissionless deployment settings, we need the secure and privacy-preserving foundation of proof of personhood, and that foundation needs to be coercion resistant in order to ensure that participants are representing their own interests rather than someone else's. I see promise that this challenge is solvable in principle because at least in-person proof of personhood (pseudonym parties) are composable with coercion-resistance mechanisms based on fake credentials and in-person credentialing as developed in the TRIP/Votegral system.

3.7 Demo session: Coercion-resistant E-voting with in-person credentialing

Bryan Ford (EPFL Lausanne, CH)

License Creative Commons BY 4.0 International license
© Bryan Ford

Joint work of Louis-Henri Merino, Simone Colombo, Rene Reyes, Alaleh Azhir, Shailesh Mishra, Pasindu Tennage, Mohammad Amin Raeisi, Haoqian Zhang, Jeff R. Allen, Bernhard Tellenbach, Vero Estrada-Galiñanes, Bryan Ford

Main reference Louis-Henri Merino, Simone Colombo, Rene Reyes, Alaleh Azhir, Shailesh Mishra, Pasindu Tennage, Mohammad Amin Raeisi, Haoqian Zhang, Jeff R. Allen, Bernhard Tellenbach, Vero Estrada-Galiñanes, Bryan Ford: "TRIP: Coercion-resistant Registration for E-Voting with Verifiability and Usability in Votegral", in Proc. of the ACM SIGOPS 31st Symposium on Operating Systems Principles, SOSP '25, p. 834–874, Association for Computing Machinery, 2025.

URL https://doi.org//10.1145/3731569.3764837

This talk briefly summarized the design and usability study results for TRIP, the coercion-resistant in-person registration process for end-to-end verifiable electronic voting in the Votegral system. The slides summarize the design and key usability study results, and the linked video is one of the three instructional videos we used to explain the system to participants in the usability study.

3.8 Which roadmap for a Software Development Kit suitable for Internet Voting?

Pierrick Gaudry (CNRS – Nancy, FR)

License Creative Commons BY 4.0 International license
© Pierrick Gaudry

In this talk, we propose to explore the possibility to build a Software Development Kit (SDK) that could be the basis for constructing various Internet Voting systems. Several questions are to be answered: which features do we want to support? Which programming language? Which licence? How to organize the governance and the maintenance? And last but not least, can we build this SDK on top of existing SDK like ElectionGuard (for in-person voting) or Verificatum (for mixnets), or by extracting parts of existing full system implementations like CHVote, SwissPost, Helios, Belenios, Polyas, etc?

3.9 Some thoughts about security definitions

Kristian Gjøsteen (NTNU – Trondheim, NO)

License Creative Commons BY 4.0 International license
© Kristian Gjøsteen

To make security results comparable, we would like to have one security definition that covers (almost) all systems, even paper-based systems and systems with more than one casting method. We also want to model all aspects of the system, including human behaviour/mistakes.

To define security, we first need to model the class of voting systems. The thing every system has in common is that there are voters that want to cast ballots, there should be a result, and there will be an adversary.

As usual, an experiment simulates the honest parties and interacts with an adversary. The adversary directs the system, including starting honest parties, scheduling message delivery, and directing honest voter operations (casting, verifying, reporting results, etc.) and signalling the next election phase (assumed to be synchronised). The adversary may also corrupt parties.

The experiment keeps track of the honest voters' beliefs about what ballots they have successfully cast. For integrity, our goal is two-fold: any accepted result is consistent with the honest voters' beliefs, and any result is unique. For accountability, parties that evaluate the result output evidence that anyone can evaluate to point to a party, and our goal is that any party pointed at is corrupt. For coercion resistance, the adversary defines two groups of voters and attempts to coerce all of them. One group resists coercion, the other does not. The adversary's goal is to distinguish the groups.

3.10 The Unstated Problem

Thomas Haines (Australian National University – Acton, AU)

License Creative Commons BY 4.0 International license
© Thomas Haines

What do the vulnerabilities we see in deployed voting systems tell us about our definitions? It suggests that we constrain attacks on state too much in definitions and not enough in our protocols.

3.11 Fuzzing E-Voting Protocol Implementations

Lucca Hirschi (LORIA & INRIA – Villers-lès-Nancy, FR)

License Creative Commons BY 4.0 International license
© Lucca Hirschi
Joint work of Max Ammann, Lucca Hirschi, Steve Kremer
Main reference Max Ammann, Lucca Hirschi, Steve Kremer: "DY Fuzzing: Formal Dolev-Yao Models Meet Cryptographic Protocol Fuzz Testing", in Proc. of the IEEE Symposium on Security and Privacy, SP 2024, San Francisco, CA, USA, May 19-23, 2024, pp. 1481–1499, IEEE, 2024.
URL https://doi.org//10.1109/SP54263.2024.00096

Today's information society crucially relies on cryptographic protocols. These protocols leverage cryptographic primitives to ensure confidentiality, integrity, or other security goals. A widespread class of vulnerabilities plaguing such protocols is logical attacks, which exploit flawed protocol logic. We examine two formal model-based methods for preventing logical attacks in protocol designs and implementations.

First, we present automated formal verification methods based on Dolev-Yao (DY) models, which formally define and excel at finding such flaws on protocol design specifications. Nevertheless, these methods alone cannot secure protocol implementations, as bugs may introduce implementation-level logical attacks. We then present a recent research avenue that aims to integrate formal DY models with fuzz testing techniques to capture logical attacks in cryptographic protocol implementations.

Finally, we explore the idea of fuzzing e-voting protocols, possibly using DY fuzzing. DY fuzzing seems ideally placed to address the unique challenges posed by e-voting systems, which constitute particularly interesting targets for several reasons: they are based on many different sub-protocols run by different actors, they aim to achieve extremely strong security properties under strong threat models, and they are critical pieces of software which have not received much attention from the fuzzing community so far.

3.12 Common trust assumptions in internet voting – how comfortable are we with them?

Audhild Høgåsen (Schweizerische Post – Bern, CH)

License Creative Commons BY 4.0 International license
© Audhild Høgåsen

We discuss how comfortable we are with some common trust assumptions for internet voting. Further, we discuss the following question: "Which do you consider preferable for security for a scheme being implemented today; a protocol which security is based on lattice assumptions

on a protocol which security is based on classical assumptions?" We also discuss questions like "Would you, in order to get coercion resistance, accept the risk that some voters are prevented from voting?" Finally, we ask ourselves: "If you must choose only 2 of the following, which do you choose?": "Vote secrecy against malicious voter device", "Individual verifiability (against malicious voter device)", "Coercion resistance", "Post-quantum security (on top of classical security, so assume some hybrid scheme)".

3.13 Defining vote privacy ... again

Steve Kremer (INRIA – Villers-lès-Nancy, FR)

License Creative Commons BY 4.0 International license
© Steve Kremer

Defining vote privacy is tricky. In this talk we review some classical game-based definitions, recall some known short-comings of these definitions, and show a few new problems. In particular, we demonstrate that one classical definition, generally believed too strong for some counting functions, may actually be too weak. We also observe that the BPriv definition, regarded as one of the most mature definitions, is too strong for protocols that chain votes (which may be needed to avoid a recent attack). By supposing an "extract function" (a hypothesis that already exists in BPriv) we propose a simple variant of a classical definition by Benaloh that seems to overcome the mentioned problems.

3.14 Usable coercion-resistant voting

Oksana Kulyk (IT University of Copenhagen, DK)

License Creative Commons BY 4.0 International license
© Oksana Kulyk
Joint work of Christina Nissen, Tobias Hilt, Jurlind Budurushi, Melanie Volkamer, Oksana Kulyk
Main reference Christina Nissen, Tobias Hilt, Jurlind Budurushi, Melanie Volkamer, Oksana Kulyk: "Voting Under Pressure: Perceptions of Counter-Strategies in Internet Voting", in Proc. of the Electronic Voting – 10th International Joint Conference, E-Vote-ID 2025, Nancy, France, October 1-3, 2025, Proceedings, Lecture Notes in Computer Science, Vol. 16028, pp. 158–174, Springer, 2025.
URL https://doi.org//10.1007/978-3-032-05036-6_10

The risks of voter coercion in online voting has lead to a body research proposing coercion-resistant solutions, that would allow a coerced voter to adopt a counter-strategy that manages to deceive the coercer while at the same time allowing the voter to vote according to their actual wishes. The usability of these solutions, and voters' ability to successfully apply them in case of coercion, however, remains an open question. In this talk I present our ongoing research on the usability of different coercion-resistant strategies, concluding that while certain strategies can potentially be made usable, more research is needed to address their limitations.

3.15 E-Collecting in Switzerland

Reto König (Bern University of Applied Sciences, CH)

License Creative Commons BY 4.0 International license
© Reto König

After a real attack on the collection process for votes on initiatives and referenda ("Unterschriftenbschiss"), Switzerland faces is urged to implement a more secure and thus trustworthy process on how to collect such votes – electronically. Even though the process seems to result in a "lightweight" e-voting protocol, the security requirements are quite hefty. In order to protect Swiss Citizens from losing all their privacy, the system is required to provide "unconditional" ever lasting participation privacy. Furthermore, the system needs to be sound regarding eligibility throughout the entire procedure which can last up to 18 month. In seeking help from our fellow e-voting experts, we propose a system based on set-membership ZKPs with a single-use rate limiter. But will it be enough?

3.16 Verifiable Tally-Hiding E-Voting and zkSNARKs for Ballot Validity

Ralf Küsters (Universität Stuttgart, DE)

License Creative Commons BY 4.0 International license
© Ralf Küsters

Modern electronic voting systems (e-voting systems) are designed to provide not only vote privacy but also (end-to-end) verifiability. Several verifiable e-voting systems have been proposed in the literature. Most existing systems reveal not just the voting result but also the full tally, consisting of the exact number of votes per candidate or even all single votes. There are several situations where this is undesirable. So-called tally-hiding systems solve this problem. In this talk, several forms of tally-hiding are discussed and two systems, namely Ordinos and Kryvos, are presented to achieve verifiable yet tally-hidng e-voting, including a post-quantum secure system. Moreover, a new approach for ensuring ballot validity is presented based on zkSNARKs.

3.17 Trusting implementations

Vincent Laporte (LORIA & INRIA – Villers-lès-Nancy, FR)

License Creative Commons BY 4.0 International license
© Vincent Laporte

The overall security of a voting infrastructure relies, to some extend, on the quality of software implementations. This talk reminds some desirable features of an implementation and explores whether voting systems could benefit from high-assurance implementations. It also presents the Jasmin system for high-assurance cryptography.

3.18 Election Audits: There's More We Can Do

Jennifer Morrell (University of Minnesota – Minneapolis, US)

License Creative Commons BY 4.0 International license
© Jennifer Morrell

Election audits are vital to accountability and public confidence, yet their methods have become increasingly politicized. This presentation advances nine standards for post-election tabulation audits – covering independence, transparency, reconciliation, and voter privacy – and connects them to trusted frameworks in other policy domains. Drawing on experimental research, we show that process-focused visual audit summaries significantly increase voter confidence, especially among independents. We translate these findings into practical templates to help election officials communicate results more effectively. The presentation concludes by exploring automation and AI-assisted tools that can enhance audit quality, transparency, and accountability to voters.

3.19 Post-Quantum Proof of Shuffle

Rafieh Mosaheb (University of Luxembourg, LU)

License Creative Commons BY 4.0 International license
© Rafieh Mosaheb
Joint work of Thomas Haines, Rafieh Mosaheb, Johannes Müller
Main reference Thomas Haines, Rafieh Mosaheb, Johannes Müller, Reetika: "Zero-Knowledge Proofs from Learning Parity with Noise: Optimization, Verification, and Application", in Proc. of the 38th IEEE Computer Security Foundations Symposium, CSF 2025, Santa Cruz, CA, USA, June 16-20, 2025, pp. 441–456, IEEE, 2025.
URL https://doi.org//10.1109/CSF64896.2025.00010

In this talk, we present the design of a post-quantum e-voting system based on the hardness of the Learning Parity with Noise (LPN) problem. We begin by motivating our choice of this approach and highlighting the importance of diversifying the underlying post-quantum assumptions in the context of e-voting. Next, we describe the phases of our proposed system and the cryptographic primitives employed. We then detail the construction of our code-based proof of shuffle. The talk concludes with a discussion of the current limitation of this proof of shuffle, specifically, its inability to distribute trust among multiple talliers, requiring reliance on a single trusted tallier to preserve ballot privacy. We are exploring directions to overcome this limitation, either by developing a code-based proof of shuffle that supports distributed trust or by rethinking the structure of the e-voting phases to achieve this goal.

3.20 Composing (provable) fine-grained definitions into an (understandable) E2E-verifiability definition

Florian Moser (famoser GmbH – Allschwil, CH)

License Creative Commons BY 4.0 International license
© Florian Moser

For an intuitive understanding of the overall security guarantees of an internet voting system, a single and short high-level definition such as E2E-Verifiability is useful (e.g., roughly, all ballots are correctly tallied exactly as cast). In contrast, when proving the security of such

a system, or when specifying precise behaviour of some part of the system, fine-grained definitions are more practical. However, proving certain fine-grained properties does not necessarily imply that the targeted high-level property is indeed fulfilled, too.

In this work, we aim to bring these two approaches together. We propose precise fine-granular definitions, inspired by usual notions in literature. Then, we compose these fine-granular definitions and observe the exact E2E-Verifiability notion reached. Finally, we propose avenues for future research, which includes most crucially a proof of our claims.

3.21 GI Elections: Individual Verification for POLYAS

Florian Moser (famoser GmbH – Allschwil, CH)

License Creative Commons BY 4.0 International license
© Florian Moser
URL https://github.com/famoser/polyas-verification

When voting over the internet, the voter needs to be able to verify that their vote was cast correctly. This avoids trusting their voting device fully. This project implements voter-side verification (also called individual verification) for POLYAS 3.0 (Version 1.3.2, 31 July 2023), following the second device spec (Version 1.2-SNAPSHOT, 22. September 2024).

This project has been developed for the Gesellschaft für Informatik (German) and has been supported by the Université de Lorraine, CNRS, Inria, and LORIA (Nancy, France). In collaboration with other researchers, a scientific publication documents the experience of implementing this verifier (and others), and formulates recommendation for similar projects: `https://inria.hal.science/hal-04663997`.

3.22 Verifying cryptographic software in Lean

Michael Naehrig (Microsoft – Redmond, US)

License Creative Commons BY 4.0 International license
© Michael Naehrig

In a collaborative effort, we are modernizing Microsoft's cryptography library SymCrypt by transitioning implementations of cryptographic primitives to Rust. These Rust implementations are then formally verified using the Aeneas tool chain to produce an equivalent Lean representation that is shown to be functionally correct according to the primitive's specification. This talk describes the methodology and discusses the reasons for this approach.

3.23 From individual verifiability to universally verified

Peter Y. A. Ryan (University of Luxembourg – Esch-sur-Alzette, LU)

License Creative Commons BY 4.0 International license
© Peter Y. A. Ryan

For End-to-end verifiable voting systems a key element of the assurance argument is that the vote is *cast as intended*: voters should be confident that their intended vote was encrypted. This often takes the form of an audit of the ballot: the encryption of the vote is opened to reveal the plaintext. It is often thought that an audited ballot cannot then be cast as this

would violate receipt-freeness: auditing typically requires revealing the randomisation used for the encryption. This leads to the use of some form of cut-and-choose protocol between the voter and the voting device, for example, Benaloh challenges.

In this talk we argue that a cut-and-choose procedure may not in fact be ncessary: if an audited ballot is re-encrypted before being posted then receipt-freeness can be preserved. The voter still needs a way to identify their ballot on the BB, but this can be achieved, for example, using suitable malleable signatures that remain valid under re-encryption of the signed text. The assurance provided by being able to cast an audited ballot is far higher than that possible via cut-and-choose techniques.

A further observation is that if this approach is used in conjunction with a scheme that involves pre-commitment to encryptions of the votes these can be independently pre-audited. This removes the need for voters to perform their own ballot audits. Given that such audits are generally regarded as burdensome this will contribute to improved usability, and means that the assurance in the outcome is less reliant on a good proportion of voters performing the checks diligently.

3.24 Sign-Then-Encrypt

Peter Rønne (University of Luxembourg, LU)

License Creative Commons BY 4.0 International license
© Peter Rønne

In many voting schemes the plaintext vote choice is being encrypted with zero-knowledge proofs of ballot correctness to form a ballot which is then digitally signed. This digital signature facilitates voter eligibility checks and prevents ballot stuffing. In this talk we suggest improving cast-as-intended verifiability by instead signing the plaintext vote directly using one device after which a second device encrypts the choice and provides a zero-knowledge proof that it has seen a valid signature of an allowed vote choice. To maliciously change a vote undetected, an attacker would need to corrupt both devices. This is especially interesting if the voting system can rely on an already existing electronic ID infrastructure and corresponding secure signing devices. This generalises ideas from schemes using blind, re-randomisable and linearly-homomorphic signatures.

3.25 Boardroom voting with their smartphones: Experiencing Verifiability

Kazue Sako (Waseda University – Tokyo, JP)

License Creative Commons BY 4.0 International license
© Kazue Sako

In my talk, I will call for proposals in a specific, real-life voting scenario. We had been discussing about in-person voting where casting vote takes in a shared computer at a polling place, and remote voting where people cast their vote through their individual computers/smartphones. The scenario I propose today is a mixture – it will be in-person, in a boardroom where people in the room are presumably authenticated, but 300 voters will be using their own smartphones to do the voting. The primary purpose is to have a shorter voting/tallying period, but this aims to provide voters the value of verifiability that is hard to achieve with plain paper ballots.

3.26 Some Usability Issues in Coercion-resistant and E2E Solutions

Roberto Samarone Araujo (Federal University of Pará – Belém, BR)

License Creative Commons BY 4.0 International license
© Roberto Samarone Araujo

Coercion resistance and end-to-end verifiability represent two of the most promising approaches for enhancing the security and transparency of modern elections. Despite their potential, these technologies still face significant challenges that hinder their use in real-world electoral settings. In this talk, we examine key usability issues associated with these approaches and discuss two potential strategies that could mitigate such limitations.

3.27 Selene Voter Experience

Steve Schneider (University of Surrey – Guildford, GB)

License Creative Commons BY 4.0 International license
© Steve Schneider
Joint work of Mohammed Alsadi, Steve Schneider
Main reference Robert Krimmer, Melanie Volkamer, Bernhard Beckert, Ardita Driza Maurer, David Dueñas-Cid, Stéphane Glondu, Iuliia Krivonosova, Oksana Kulyk, Ralf Küsters, Beata Martin-Rozumilowicz, Peter Rønne, Mihkel Solvak, Oliver Spycher: "Fifth International Joint Conference on Electronic Voting – E-Vote-ID 2020 – Taltech Press Proceedings", 2020.
URL https://www.researchgate.net/publication/344471755_Fifth_International_Joint_Conference_on_Electronic_Voting_-_E-Vote-ID_2020_-_Taltech_Press_Proceedings#page=302

Selene is a verifiable voting protocol proposed by Ryan, Roenne and Iovino. It uses trackers to provide a way of enabling voters to check their vote in plaintext at the end of the election. It can be added to existing voting schemes to add verifiability functionality.

This talk described a trial "in the wild" with a commercial internet voting provider, whereby their existing (non-verifiable) system was augmented with additional functionality based on Selene verifiability (though with stronger trust assumptions) to give voters the opportunity to check their votes. This was used in two real online elections in order to explore voters' experience of verifiability and their opinions of it, by means of a questionnaire. The majority of the respondents found the system easy to use and considered that it maintained their vote privacy. They appreciated the opportunity to verify their vote and felt it gave them confidence in the system, but they were split on seeing the point of verifiability. This leaves open the question of how to engender trust in electronic voting systems and to what extent this needs to relate to understanding.

3.28 Voting Channel Security

Carsten Schürmann (IT University of Copenhagen, DK)

License Creative Commons BY 4.0 International license
© Carsten Schürmann

Elections are designed to achieve a peaceful resolution to a civic battle about power and governance. Implied is that elections should ensure that all voters – supporting winners and losers alike – accept the election outcome and develop collective trust. In my talk, I argue that if an election offers multiple voting channels, their interactions might be subtle, and they can have negative effects on the collective trust. Politicization driven by policitcal parties, activism, and foreign influence of existing voting channels, malicious or accidental, can trigger security violations that puts this collective trust at risk.

3.29 Quantum-Safe Electronic Voting: Status and Challenges

Tjerand Silde (NTNU – Trondheim, NO)

License Creative Commons BY 4.0 International license
© Tjerand Silde

Quantum computers are coming and will break the (long-term) privacy and (online) integrity of electronic voting schemes; hence, we need to build new schemes from quantum-safe assumptions. This talk will overview the quantum threat, lattice-based cryptography, and our current state-of-the-art schemes.

3.30 Electronic Voting in Brazil: How E2E-V might work there (a work in progress)

Marcos Simplicio (University São Paulo, BR)

License Creative Commons BY 4.0 International license
© Marcos Simplicio

The talk briefly shows how electronic voting in Brazil currently works: it is based on Direct Recording Electronic (DRE) machines empowered by Trust Platform Modules (TPM) for protecting its underlying software – so they are called "T-DRE" machines. To improve the system's transparency, it is possible to integrate end-to-end verifiability (E2E-V) features into the process, thus giving voters the ability to check if their individual votes where cast-as-intended, recorded-as-cast, and tallied-as-recorded. The presentation shows the proposed procedures, which include: tracking codes generated by voting machines via cryptographic commitments (e.g., using Pedersen commitments for ever-lasting privacy); Benaloh challenges on those commitments, done under the supervision of auditors, with simulated votes, to preserve vote privacy; signature of cast votes, aiming to facilitate dispute resolution if a tracking code does not appear in the tally; and publication of all relevant records. It also discusses the confidence level of the election results, depending on the number of challenges performed and the difference in the number of votes received by each candidate.

3.31 Colorado IRV RLAs

Vanessa Teague (Australian National University – Acton, AU)

License Creative Commons BY 4.0 International license
© Vanessa Teague
Joint work of Vanessa Teague, Michelle Blom, Andrew Conway
URL https://github.com/democracydevelopers

We'll talk through the educational materials, and possibly also demo the code, at `https://github.com/democracydevelopers`.

This extends Colorado's Risk Limiting Audit code to IRV, based on the RAIRE system (Blom, Stuckey, Teague, https://arxiv.org/abs/1903.08804).

The code will run in Colorado's RLA this November.

3.32 Work in Progress: SoK on usable verifiable electronic voting

Melanie Volkamer (KIT – Karlsruher Institut für Technologie, DE)

License Creative Commons BY 4.0 International license
© Melanie Volkamer

Based on our literature review on usability studies for verifiable electronic voting systems, we identified more than 20 factors in which the study setting differs, such as the type of mechanism studied, whether it was a remote or a lab study, what instructions they got, and what type of manipulations were studied. Given the number of papers we found, we conclude that it is difficult to compare the results. We learned that the most challenging part is to measure the effectiveness of the verifiability mechanism: One can observe e.g. by eye tracking whether voters go for the verification step at all but does this mean they would notice the manipulation and if yes would they report it properly.

3.33 Audit (Process) Efficiency for Plurality Elections

Poorvi Vora (George Washington University – Washington, DC, US)

License Creative Commons BY 4.0 International license
© Poorvi Vora

We provide an example where a mathematical model which better represents the physical audit process greatly reduces the expected number of ballots drawn in a ballot polling RLA, and thus reduces the number of person-hours required for the audit. We hence argue that better models of the true audit process on the ground will generally lead to (different) statistical tests representing "better" real audits.

4 Working groups

This section presents abstracts of selected breakout groups from the Dagstuhl Seminar.

4.1 Evidence-Based Elections (commitments)

Josh Benaloh (Microsoft Research – Redmond, US)

License Creative Commons BY 4.0 International license
© Josh Benaloh

This discussion focused primarily on using cryptographic commitments rather than encryption for verifiable voting systems. This technique is most useful for in-person voting where it enables the complete elimination of cryptographic keys and the trustees who would generate and manage the keys.

The discussion was generally quite favorable, although an important point was raised that backup can create privacy risks that are not present with the prior approach of using encryption to protect the privacy of votes. With the encryption approach, ballots are independently encrypted and no intermediate tally data is retained. However, the commitment approach effectively maintains a running tally after each new ballot is cast,

and a full sequence of intermediate tallies reveals the contents of each individual vote. The concern is, however, mitigated by the fact the the privacy risk is no greater than that posed by traditional ballot tallying systems which also maintain running tallies.

4.2 Coercion Resistance

Constantin Catalin Dragan (University of Surrey – Guildford, GB)

License Creative Commons BY 4.0 International license
© Constantin Catalin Dragan

Coercion has typically been viewed as "persuading someone to do something by using force or threats" (Oxford dictionary). The voting community has proposed some solutions for when the voter can leave the coercive environment (e.g., fake credentials in JCJ/Civitas or re-voting in Belenios) or under extreme coercion – can never leave the environment and can only annul their vote.

One of biggest challenges is to define what makes acceptable manipulation versus actual coercion. The voting community aims to technically and ethically support governments in tackling this challenge. One potential context is we've seen a recent increase in electoral influence via social media that can have drastic consequences on elections, for example the cancelled 2024 Romanian presidential election or the Cambridge Analytica scandal on influencing the 2016 US presidential elections. Another big challenge is on understanding the trade-offs between problems and solutions, with a focus on designing relevant baselines for comparison (e.g., failure rates, voter expectations, and real-world testing). It is exceptionally hard to test any developed solutions under high-stress conditions to simulate the real-world voting environment.

4.3 Methods for Trustworthy Randomness Generation

Jennifer Morrell (University of Minnesota – Minneapolis, US) and Peter Rønne (University of Luxembourg, LU)

License Creative Commons BY 4.0 International license
© Jennifer Morrell and Peter Rønne

This working group discussed methods for trustworthy randomness generation used in audits of elections. The current most-used method consists in rolling dice for seed to a pseudo-random function, where the ceremony is public or recorded and broadcasted. Different attacks were discussed such as deliberate biasing (or predicting) of the sample, but also the accusation that the sample is biased or predictable (when it isn't). New constructions and protocols were discussed, especially with active contribution from people who wish to participate in the ceremony.

4.4 Post-Quantum Cryptography and E-Voting

Rafieh Mosaheb (University of Luxembourg, LU) and Peter Rønne (University of Luxembourg, LU)

License Creative Commons BY 4.0 International license
© Rafieh Mosaheb and Peter Rønne

Currently, there is a large ongoing transformation to post-quantum cryptography due to potential threats from upcoming quantum computers against classical asymmetric cryptography.

This breakout group discussed the current status of quantum-safe e-voting, threats and solutions including partial solutions such as voting with perfect privacy that protects against store-now-decrypt-later attacks. In particular, possible constructions for hybrid verifiable encryption were discussed as a way to protect election privacy by combining classical and post-quantum cryptography, ensuring security even if a PQ scheme is later found to be flawed, while still providing verifiability against quantum attackers.

4.5 Education and Outreach

Florian Moser (famoser GmbH – Allschwil, CH)

License Creative Commons BY 4.0 International license
© Florian Moser

The participants identified different target public that need to be reached with education and outreach projects. The voters need to know how to use the system, but additionally need to be able to react to unforeseen states (e.g., a failed verification check), and to battle misinformation (e.g., claims about frauds of the election). The providers need to be able to continuously improve their system, given their respective resources and expertise. Last but not least, also non-experts that interact deeply with these systems need to be empowered to take informed decisions, which includes the election organizers, activists and observers.

However, many aspects of verifiable elections are hard to communicate. First, many results are non-intuitive, but nonetheless correct. For example, risk-limiting audits in elections with large margins need to audit very few ballots to achieve their guarantees (e.g. in the Kenyan elections 2017, 200 ballots for the whole country would have been enough for 99.9% of confidence). Also, the cryptography used to anonymize ballots before decryption are not well known outside of the internet voting domain, and may be unintuitive even for otherwise experts in IT security. The complexity and the "rough edges" of the current state-of-the-art systems may lead to distrust, or even the choice of an insecure but perceived simpler system.

The participants also discussed how to concretely move forward. Some of the measures discussed are to audit & update Wikipedia articles concerning verifiable elections, and to approach the development of an SDK for system providers (as also discussed in another session in this seminar).

4.6 Software Development Kit for E-Voting

Florian Moser (famoser GmbH – Allschwil, CH)

License © Creative Commons BY 4.0 International license
© Florian Moser

The working groups had a session discussing the development of a Software Development Kit for voting. The targets for the SDK should be verifiable elections, both for remote voting and in-person voting. While previous work exists, these SDKs usually cover only part of the cryptography used in verifiable elections, or are adapted to the needs of a specific system or vendor.

Some of the challenges discussed were the intended usability both for a target audience without a deep cryptographic background, while allowing customizability for cryptographers for special use cases. A basis for the the initial SDK would be public key encryption, committments, signatures and zero-knowledge proofs, as well as procedures for distributed key generation, distributed decryption, verifiable shuffles and homomorphic tallies. Further, it should provide building blocks for public bulletin boards and tamper-proof logging. Documentation should also include instructions for auditors of elections.

The governance aspects were also discussed, especially with a view towards long-term stability of the SDK by contacting and inviting stake-holders, from industry, administration and vendors, early in the process. The library does not need to start afresh, but can profit from existing libraries and their experiences, especially ElectionGuard.

4.7 Usability

Peter Rønne (University of Luxembourg, LU)

License © Creative Commons BY 4.0 International license
© Peter Rønne

The discussion int this working group was usability, user experience and trust. Usability encompasses several key aspects: the vote casting process, voter verifiability – including verifying that votes are both cast-as-intended and stored-as0cast but sometimes also more advanced mechanisms, e.g., ensuring resistance towards coercion. A problem with the current usability and user experience research is that it exhibits a high variability in study design, which demonstrates a need for clearer guidelines that address different manipulation types and the reporting of detected issues in the voting process especially for verifiability. Building trust in elections is a deeply interdisciplinary challenge that demands collaborative research. Future studies should explore how different types of information affect trust, particularly messages that encourage voters to verify or report issues. Additionally, while existing research focuses mainly on voters, more attention should be given to poll workers. Finally, two specific research gaps persist: effectively communicating new verification capabilities to voters accustomed to "black box" systems, and understanding the long-term effects of voter verifiability.

Participants

- Josh Benaloh
Microsoft Research – Redmond, US
- Michelle Blom
The University of Melbourne, AU
- Véronique Cortier
LORIA, CNRS – Nancy, FR
- Alexandre Debant
INRIA – Villers-lès-Nancy, FR
- Thi Van Thao Doan
University of Louvain, BE
- Constantin Catalin Dragan
University of Surrey – Guildford, GB
- David Dueñas-Cid
Kozminski University, PL
- Alexander Ek
KU Leuven, BE
- Bryan Ford
EPFL Lausanne, CH
- Pierrick Gaudry
CNRS – Nancy, FR
- Kristian Gjøsteen
NTNU – Trondheim, NO
- Thomas Haines
Australian National University – Acton, AU
- J. Alex Halderman
University of Michigan – Ann Arbor, US
- Lucca Hirschi
LORIA & INRIA – Villers-lès-Nancy, FR
- Audhild Høgåsen
Schweizerische Post – Bern, CH
- Reto König
Bern University of Applied Sciences, CH
- Steve Kremer
INRIA – Villers-lès-Nancy, FR
- Ralf Küsters
Universität Stuttgart, DE
- Oksana Kulyk
IT University of Copenhagen, DK
- Vincent Laporte
LORIA & INRIA – Villers-lès-Nancy, FR
- Karola Marky
Ruhr-Universität Bochum, DE
- Jennifer Morrell
University of Minnesota – Minneapolis, US
- Rafieh Mosaheb
University of Luxembourg, LU
- Florian Moser
famoser GmbH – Allschwil, CH
- Michael Naehrig
Microsoft – Redmond, US
- Olivier Pereira
University of Louvain, BE
- Thomas Peters
University of Louvain, BE
- Bart Preneel
KU Leuven, BE
- Sylvain Ruhault
ANSSI – Paris, FR
- Mark D. Ryan
University of Birmingham, GB
- Peter Y. A. Ryan
University of Luxembourg – Esch-sur-Alzette, LU
- Peter Rønne
University of Luxembourg, LU
- Kazue Sako
Waseda University – Tokyo, JP
- Roberto Samarone Araujo
Federal University of Pará – Belém, BR
- Steve Schneider
University of Surrey – Guildford, GB
- Carsten Schürmann
IT University of Copenhagen, DK
- Tjerand Silde
NTNU – Trondheim, NO
- Marcos Simplicio
University São Paulo, BR
- Vanessa Teague
Australian National University – Acton, AU
- Melanie Volkamer
KIT – Karlsruher Institut für Technologie, DE
- Poorvi Vora
George Washington University – Washington, DC, US