Voter Registration: A Security and Cryptography Perspective

by

Andrés Fábrega Gerbaud

B.S. Computer Science and Engineering
Massachusetts Institute of Technology, 2020

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2022

Author .....
Department of Electrical Engineering and Computer Science
May 14, 2022

Certified by .....
Ronald L. Rivest
Institute Professor
Thesis Supervisor

Certified by .....
Sunoo Park
Postdoctoral Fellow, Cornell Tech
Thesis Supervisor

Accepted by .....
Katrina LaCurts
Chair, Master of Engineering Thesis Committee

Voter Registration: A Security and Cryptography Perspective

by

Andrés Fábrega Gerbaud

Submitted to the Department of Electrical Engineering and Computer Science
on May 14, 2022, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

Security and transparency of voter registration systems are crucial properties that any electoral system must satisfy: without robust guarantees on the underlying voter data, trust in election results—and the system as a whole—is severely impacted.

In this thesis, we study two fundamental problems related to the security of voter registration. First, we formalize voter registration systems by providing a set of high-level definitions that characterize these systems in a general sense. To our knowledge, this is the first formal treatment of this sub-field of election security, which is (surprisingly) often neglected by the academic community. By abstracting away low-level implementation details, our work provides a clearer understanding of these complex systems; furthermore, it lays the formal groundwork and definitions which are useful to design secure technical protocols. Thus, we hope to pave the way for more research in this area.

Secondly, we give a brief overview of an ongoing work-in-progress consisting of a new design for voter registration systems with stronger transparency guarantees, where voters are able to independently verify that their data has not been tampered with, even in the presence of untrusted election officials. We hope that our eventual system increases voter confidence in the electoral system, and helps detect (and, thus, mitigate) attacks that target voter registration databases.

Thesis Supervisor: Ronald L. Rivest
Title: Institute Professor

Thesis Supervisor: Sunoo Park
Title: Postdoctoral Fellow, Cornell Tech

Acknowledgments

I first want to thank both of my supervisors, Ron Rivest and Sunoo Park, for their amazing mentorship and guidance. Both inside and outside the scope of this thesis, they have both provided me with invaluable advice and support that has helped me in uncountable ways. Thank you for being two incomparable role models who showed me what academic excellence looks like, and who taught me the importance of being passionate about your research. I am deeply inspired by your passion for using cryptography to benefit society, and working with you both has had a profound impact on my academic life and interests.

Thank you to Mike Specter and Jack Cable for being amazing collaborators. Working with you two was an absolute pleasure, and I will always cherish the (technical and non-technical) conversations we had.

Thank you also to my friends, both from back home and from MIT, for supporting me when times were tough and for always reminding me about the importance of balancing work and life.

Lastly, I am eternally grateful to my family for their unconditional love, and for always providing me with the support I need to pursue my dreams. ¡Gracias por todo!

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Voter registration is one of the most fundamental links of any voting system; without an accurate list of eligible voters, the security and correctness of the entire electoral system is affected. A good understanding of such systems is thus very important for election officials, the election security community, and the general public.

Voter registration systems are very complex, as they are composed of multiple subsystems involving various entities, are very dynamic and constantly evolving, and are subject to nuanced judicial requirements. Furthermore, the details of them vary wildly across jurisdictions (even from state to state, in the case of the US), making their study even more complicated. This has led to a surprising lack of attention from the election security community, particularly by researchers and academics. One big research gap, which further exacerbates this lack of work, is that there are no precise technical definitions of voter registration systems and their security; without a formal definition and high-level description of voter registration systems, it is very difficult to understand (and, thus, study) them in a general sense. Furthermore, there is no formal groundwork upon which to base technical work, such as the design of cryptographic protocols or data structures.

This thesis is mostly composed of a research paper that studies voter registration systems from a definitional point of view. In this joint work with Sunoo Park, Jack Cable, and Michael Specter, we tackle the aforementioned gap, and lay the formal foundations for studying voter registration systems by introducing a set of definitions

and security properties that abstractly model these systems in a generic way. To do this, one core aspect of our work is that we define a set of abstract objects that encapsulate jurisdiction-specific details, upon which our definitions depend, allowing one to study voter registration systems in depth while being able to ignore low-level implementation features. Furthermore, we explain how to, if desired, instantiate these objects using concrete parameters, in order to gain insight into specific jurisdictions. We provide an example of this by doing a case study of the state of Colorado. Hence, our work serves as a framework, which is flexible enough to study voter registration systems in a black-box way or specific implementations of these.

This work will be of interest to many audiences. First, as mentioned already, to the election security community, as we hope this formal treatment will open the door for more work in this space. In particular, we are optimistic that our set of definitions will provide the foundation for technical proposals related to voter registration systems. Secondly, it will be of interest to election officials, which can instantiate the framework with the specific details of their jurisdiction to get an abstract representation of it. This can be useful when doing security audits, considering different proposals for system changes, evaluating their coverage of their threat model, and much more. In general, we hope our work will help them have a deeper and more comprehensive understanding of their own system. Lastly, it will be of interest to the general public, who will gain insight into the details of voter registration, and thus have more confidence on how their data is handled.

The last part of this thesis outlines a separate line of work, with the same set of collaborators as before, that explores the use of cryptographic techniques to provide stronger transparency guarantees to voters. After a member of the public registers to vote, there is, unfortunately, no concrete way for them to verify that their data was properly stored in the voter registration database. Indeed, there are various reasons why voters may be interested in verifying their data: fear of voter purges, external attacks, system errors, or simply lack of trust in the system. In this work, we attempt to design a new voter registration database that gives voters a way to verify that their data is present in the database and that it has not been tampered with.

That is, voters can request a cryptographic proof from election officials, which they can verify to ensure that their data is in its desired state, even if the election officials might provide incorrect information. This work is still in progress, so the treatment of it is more high-level and experimental.

We believe these transparency guarantees could be of tremendous value to the public, as we empower voters to be able to directly verify their own data instead of having to trust government authorities. Thus, this detection mechanism is an important step towards mitigating attacks and errors targeting the electoral rolls. We hope our results will help increase public confidence in the system, and legitimize the outcome of elections.

Chapter 2, which is the bulk of this thesis, consists of the aforementioned paper with Sunoo Park, Jack Cable, and Michael Specter, where we formalize voter registration systems. Chapter 3, which is much shorter, provides a brief overview of the ideas behind our new design for a voter registration database with stronger transparency guarantees. Lastly, Chapter 4 serves as a short conclusion, which summarizes the work contained in this thesis.

Chapter 2

Voter Registration Modelling

Voter registration is an essential part of almost any election process, and its security is a critical component of election security. Yet, despite a history of compromises of voter registration systems, relatively little academic work has been devoted to securing voter registration systems, compared to research on other aspects of election security.

In this chapter, we present a systematic treatment of voter registration system security. We propose the first rigorous definitional framework for voter registration systems, describing the entities and core functionalities inherent in most voter registration systems, the jurisdictional policies that constrain specific implementations, and key security properties. Our definitions are configurable based on jurisdiction-specific parameters and policies. We provide a template for the structured presentation of detailed jurisdictional policy information, via a series of tables, and illustrate its application with a detailed case study of Colorado's voter registration system. Throughout our research, with the aim of realism and practical applicability, we consulted current and former U.S. election officials, civil society, and non-profits in the elections space. We conclude with a list of critical questions regarding voter registration security.

2.1 Introduction

Voter registration systems maintain a list of eligible voters, and are a crucial component of almost any election process. Starting well before election day, jurisdictions are tasked with enrolling eligible voters' information — either automatically or on a voter's initiative — and must keep that information up to date and verifiable for use throughout the democratic process.

Public attention and academic research around election security often focus more intensely on the *casting and counting* processes that happen on and right after election day, than on voter registration and other non-voting processes. Yet voter registration security is critical to election security: a voter registration system failure can cause significant disruption to an election and the public's confidence. The results of failure could include disrupting voting processes (forcing voters to cast provisional ballots, if available), preventing voters from receiving absentee ballots, and the leakage and misuse of sensitive personal and political information.

Recognizing the importance of voter registration system security, the U.S. Department of Homeland Security's designation of election infrastructure as national critical infrastructure explicitly includes voter registration systems [9]. At least three U.S. states and numerous other countries and regions have suffered publicized compromises of their voter registration systems [7, 8, 11, 48, 68, 71], underscoring the value of registration systems as targets for attack and as potential sources of damage to electoral integrity and confidence. Some of these security incidents arose from software errors (e.g., [8]); others were perpetrated by foreign adversaries (e.g., [11]). Yet other compromises may have gone undetected or unreported.

At first glance, the voter registration problem might appear to be addressed by known solutions in the distributed and accountable systems literature. For example, maintaining a canonical, audited database has been studied in a variety of settings including distributed consensus systems [60], the HTTPS ecosystem [40], and, most recently, decentralized currencies [52]. However, voter registration systems are complex and specialized systems with functionality requirements and security challenges

not encapsulated by generalized database management and security. For example, the availability requirements on a voter registration database on election day are unusually demanding and time-constrained. Voter registration systems also have unusual accessibility requirements, as they must accommodate *any eligible voter* in the relevant electorate: a highly diverse set of people from whom no technical expertise must be required (since that should not be a requirement to vote). Relatedly, voter registration is commonly facilitated by third-party intermediaries — neither the election office nor voters — that relay communication between the election office and voters, such as department of motor vehicles[1] or nonprofit organizations. Election administrators are also often under-resourced, so it bears note that even basic security practices may be difficult to implement [53].

Currently, the security research community lacks a precise and systematic shared understanding of the scope and security challenges of voter registration. The infrequent security and cryptography publications that focus on voter registration have scoped out specific sub-problems and offered valuable technical approaches to them, but hardly any prior work has addressed voter registration security with a more holistic perspective alongside technical depth aimed for a research audience (see Section 2.2 for more discussion on prior work).

One barrier to such a systematic approach may have been the large variation between voter registration systems' implementations and requirements across jurisdictions. Even within the United States, every state manages its own voter registration system subject to its own state election law (in addition to federal law, which is fairly limited in scope), resulting in significant differences in implementation. The types of information collected and treated as public or confidential, registration methods offered, voter authentication methods, and conditions for updating or removing voter information are all subject to these jurisdiction-dependent regulations. Across countries, of course, an even wider range of laws apply.

This paper provides a systematic treatment of voter registration security. Our

[1] Under the U.S. National Voter Registration Act, states must offer voter registration opportunities at certain offices, including public assistance and disability offices. [5]

aim is to serve as a reference for the security research community in (1) identifying research questions in voter registration security, (2) framing voter registration functionalities and security definitions in shared terminology, (3) assessing the applicability of security approaches across different jurisdictions, and (4) effectively obtaining and organizing detailed information about a particular jurisdiction's voter registration requirements, to facilitate jurisdictionally tailored designs and security analyses.

To this end, we provide definitions of the *categories of entities*, *core functionalities*, and *security requirements* inherent to voter registration. These definitions, while rigorous, are formulated at a high enough level of abstraction to capture the features common to all fifty U.S. states and many other countries. We also provide a systematic exposition of the *jurisdiction-specific parameters* and *policies* that, when combined with the more abstract definitions just described, yield detailed lists of entities, functionality descriptions, and security requirements tailored to a particular jurisdiction.

The jurisdiction-specific parameters and policies effectively *instantiate* the abstract definitional framework to represent particular real-world implementations and security needs. The separation between the abstract definitions and the jurisdiction-specific parameters and policies highlights which aspects can be treated as common to most registration systems, and which aspects will need to be configured per jurisdiction.

To further illustrate how our framework might yield a jurisdiction-specific instantiation of our definitions, we explore a specific *case study* of a voter registration system deployed in one U.S. state: Colorado. Based on information from a range of public sources (such as Colorado's election code), we compile a detailed description of Colorado's parameters and security policies as relevant to voter registration, which we then validated with a Colorado state election official. We hope that this case study provides a useful template for researchers to organize jurisdiction-specific voter registration parameters and policies of interest, as well as for transparency-minded election officials to provide information useful to security researchers and the public

in a structured, extensive form amenable to comparison between jurisdictions.

To ensure that our work is grounded in the reality of how voter registration systems work in practice, we gathered feedback from a range of election experts, checked our definitions' compatibility with a range of U.S. states' and other countries' voter registration systems using public compilations of comparative data, and confirmed our definitions' applicability by conducting the detailed case study of Colorado mentioned above. See Section 2.3.2 for more details.

Finally, with a view to facilitating effective communication between security experts, election officials, and the public about security-relevant issues in voter registration, we provide a collection of critical questions as a starting point for anyone looking to gather information about strengths, weaknesses, and potential for improvement in the security of proposed or deployed voter registration systems.

In summary, our contributions are as follows:

1. We provide the first definitional framework for voter registration system security, comprising core technical functionalities, entities, jurisdictional parameters, and security policies (Sections 2.4–2.5).

2. We define security properties required by voter registration systems. Our definitions are *configurable* to accommodate jurisdictional policy variations (Section 2.5).

3. To map these general definitions to real-world implementations, we provide a template for the structured presentation of detailed jurisdiction-specific policy information, via a series of tables (Section 2.6).

4. We conduct a case study of Colorado's voter registration system, showing the instantiation of our definitions with concrete jurisdictional parameters (Section 2.6).

5. We offer a collection of critical questions regarding security in voter registration systems (Section 2.7).

2.2 Relation to Prior Work

To our knowledge, there has been no systematic treatment of voter registration system security that provides precise problem definitions and system (security) requirements. Furthermore, there has been no treatment that captures the realistic constraints and operation of voter registration systems on the ground today. This paper aims to fill that gap: that is, to provide a detailed and systematic exposition of the challenges of voter registration security in practice, laying the groundwork for the security community to better contribute its expertise to pressing issues in voter registration.

2.2.1 Systematizing voter registration system security

The most extensive prior overviews of security considerations in voter registration systems are a 2006 report commissioned by the ACM U.S. Public Policy Committee on "accuracy, privacy, usability, security, and reliability issues" related to "statewide databases of registered voters" [18], and a 2019 report by the MITRE Corporation on "recommended security controls for voter registration" [25]. These two reports have very different emphases, as summarized next; they provide important perspectives complementary to our work.

The ACM report was produced at a time when U.S. states were adopting statewide voter registration databases to comply with then-new federal legislation [18]. The report's focus is much more policy-oriented, compared to our focus on definitions and systematization: for example, it lacks technical definitions of core functionalities or security properties. Within its broad policy-oriented scope, the ACM report is remarkably comprehensive, detailed, and thoughtful about security issues.

The MITRE report has a more technical focus. The bulk of the report overviews security measures and best practices[2] broadly applicable beyond the scope of voter registration. The MITRE report also presents a generalized voter registration system architecture and parties involved therein, in less detail than (but consistent with) our model; however, unlike this paper, it neither formalizes functionality and security

[2]E.g., firewalls, TLS, VPNs, and multifactor authentication.

requirements nor engages with variations in jurisdictional policy.

Additionally, the Electoral Knowledge Network's website on voter registration [33] is a rich source of information about how voter lists are operated across the world. Its focus is broader than security or technology: instead, it offers detailed information on operational and administrative issues, as well as a range of case studies and practitioners' perspectives on voter registration in specific regions.

Other (mostly policy-focused) reports that discuss security in voter registration systems are generally less comprehensive, and tend to have less technical detail than the ACM and MITRE reports. These include: an excellent series by the Brennan Center for Justice, including but not limited to [43, 47, 62, 65, 72]; the U.S. Election Assistance Commission's resources on voter registration systems [31,32]; a 2008 report by the National Research Council of the National Academies of Sciences, Engineering, and Medicine [58]; and a 2020 report by the Center for Election Innovation and Research [26]. These are valuable resources to understand specific aspects of modern voter registration systems, potential security issues, and the concerns of those managing the systems on the ground. Further information of this type may be found in policy-oriented resources discussing election infrastructure security more broadly, such as [19, 27, 53, 70].

2.2.2 Technical work

Another area of related work comprises technical proposals, such as secure protocols (e.g., [50]) or statistical techniques (e.g., [24]), that may improve voter registration system security. Beyond academia, a number of non-governmental organizations offer innovative technological solutions to improve the integrity of voter registration data. Examples include the Electronic Registration Information Center (ERIC), a non-profit that helps identify voters who have moved, died, or have duplicate registrations across U.S. states [2], and VoteShield, a non-profit that provides tools to monitor changes to voter data for anomalies [3].

There is also a body of technical work proposing approaches to improve the security of election infrastructure other than voter registration systems, such as ap-

proaches and systems for secure casting and tallying (e.g., [17, 20, 22, 64]) or post-election auditing (e.g., [21, 44, 45]). A related literature warns of serious security risks entailed by certain technical approaches — such as Internet voting — if used in high-stakes political elections, given the limitations of the current state of the art in computer security (e.g., [53, 61, 67]).

2.2.3 Beyond security

Many aspects of voter registration are beyond the scope of this paper, because our focus is on system security. Important security-adjacent considerations include usability, privacy practices, software engineering practices, and personnel training. For an overview of these broader topics, we recommend [18] (an ACM report on registration systems) and [53] (a National Academies report on election systems generally).

2.3 Background & Methodology

2.3.1 Background

Voter registration is the act of maintaining an accurate list of voters who are eligible to vote in an election. While most countries have some of voter registration, practices vary widely. Countries may institute compulsory voter registration, in which voters are either automatically registered (such as in Argentina, Chile, Hungary, Israel, and the Netherlands) or required by law to register (such as in New Zealand and Tonga) [12]. In other cases, including the United States and India, qualified residents are not required to register to vote by law, though generally must be registered to vote in order to vote.

While a straightforward premise, maintaining voter registration databases (VRDBs) is complicated by a number of practical and legal concerns. Election administrators must allow voters to register or update their registration in a variety of means, including in-person, mail, fax, email, and via web portals. This list must then be accessible to election officials when the voter requests their ballot, both for access control, and

to allow the official to customize the ballot for the various contests available to a particular voter in that election. Election administrators must also perform complicated maintenance on the database when voters become inactive or ineligible, which may be the result of any number of life occurrences (e.g. when a voter dies or leaves the jurisdiction). Finally, the voter registration database may have a number of transparency requirements. Members of the public, including various entities and candidates, may be allowed to review the database's contents to ensure that the list is accurate.

Voter registration databases must therefore allow access and maintenance by a variety of entities of varying trustworthiness and technical ability. This includes state election officials, local election officials, and poll workers (many of whom are only temporarily employed). For instance, poll workers must have access to the voter registration database (or a local copy of it) in an electronic pollbook in order to check in voters on election day – this brings its own security challenges, as every electronic pollbook may have an entire copy of the VRDB on it.[3] In some cases, states may provide third parties access to their voter registration databases or data therein, such as third parties like ERIC to review databases.

Broad legal landscapes govern voter registration databases. In the United States, for instance, voter registration systems are run separately by each state,[4] and registration occurs in most states solely on the initiative of the voter.[5,6] Like all election administration, voter registration in the U.S. is heavily decentralized, with implementations dependent on state and local election laws and policies [32]. The National Voter Registration Act of 1993 (NVRA, also known as the Motor Voter Act) required states to use a unified voter registration form for federal elections, allow voters to register to vote while applying for driver's licenses, and allow voters to register to vote by mail [5]. The Help America Vote Act of 2002 (HAVA) mandated that states base their voter registration systems on a computerized voter registration database [6]. States

[3] An electronic pollbook is a device used by poll workers to review voter registration lists. If using electronic pollbooks, election officials often maintain paper copies in case of failures.

[4] All states, with the exception of North Dakota, require voter registration to vote.

[5] "In many democracies, citizens are automatically registered to vote. The requirement in the United States that citizens take the initiative by registering is not only atypical, but also costly to administer." [38]

[6] As of 2022, 22 U.S. states had implemented automatic voter registration. [57]

have taken three primary approaches: *top-down* databases maintain a central, authoritative database statewide; *bottom-up* databases have local jurisdictions maintain authoritative registration databases, which are then compiled in a statewide database; and *hybrid* systems give local offices discretion to either maintain an authoritative list locally, or rely on a statewide database.

Voter records are made available to various third parties, often including the public, in all 50 states [54]. Laws and policies governing access vary widely – in some states, the voter registration list (excluding certain fields) is made publicly available for download (e.g. North Carolina [59]), while in others data is restricted to political parties and other organizations, such as in Maine [46]. The data may be available either for free or for purchase, and commercial vendors sell compiled "voter files" that contain records of most American voters for political outreach and advertising purposes [10].

States also offer protections to voters whose safety would be threatened by the public release of their voter registration information, such as victims of domestic violence [54]. Most commonly known as an Address Confidentiality Program, voters may request to have a substitute address listed in their record. This allows participants to vote without fear for their safety, and hence protecting their private information is critical.

Threats to voter registration databases.

Following the 2016 election, attention has grown towards the security of voter registration databases. U.S. intelligence officials have confirmed that hackers from the the GRU, Russia's foreign military intelligence agency, targeted all 50 states' voter registration systems in the run-up to the 2016 election, succeeding in two states, including Illinois [11]. In Illinois, the hackers exfiltrated hundreds of thousands of records – including social security numbers – before being caught. While there is no evidence that the hackers modified voter records in these cases, the threat remains that voter records could be surreptitiously modified. Threats to the availability of voter registration databases may also pose a threat, for instance by preventing election officials from looking up voters on election day.

Additionally, voter purges occur when jurisdictions remove voters from registration lists for illegitimate regions, often discriminating against certain groups of voters [13]. Between 2014 and 2016, according to a Brennan Center for Justice report, states purged nearly 16 million voters from registration lists [23]. As many of these states do not implement same-day voter registration, voters who are unaware that their records have changed may be unable to vote.

Finally, voter registration stuffing may occur when an attacker surreptitiously adds voters that are fake, ineligible, or dead. This can be mitigated by a number of controls, including interstate programs such as ERIC and public transparency of voter registration lists. In practice, numerous studies have found voter registration stuffing to be extremely rare [35, 51].

2.3.2 Methodology

To construct our model, we began by performing a survey of publicly available documentation of voter registration systems used in the United States, including comprehensive overviews of systems in all 50 states and the District of Columbia via the National Conference of State Legislatures (NCSL) [23, 31, 43, 47, 54–57, 62, 72].[7] We also reviewed compilations of information on several international systems [15, 33, 37, 41, 42, 65, 66].

We then conducted a series of informal discussions with a variety of current and former U.S. elections officials, civil society organizations, and non-profits in this space. Discussions focused on understanding the voter registration process, perceived risks, and functional requirements of voter registration, filling in gaps from the available documentation. We then iteratively developed our framework through repeated feedback from these stakeholders to ensure that our abstraction maps correctly to the real-world application of these systems.

Finally, we conducted a case study focused on applying our model to the state of Colorado to both further validate our models, and provide a worked example of

[7] Throughout the paper, though we mainly focus on the 50 U.S. states, we recognize that the District of Columbia and U.S. territories also maintain voter registration lists.

their practical application. We completed tables containing jurisdictional parameters and security policies based on a review of publicly accessible laws, policies, and documentation detailing Colorado's voter registration system. We then validated the tables with an official from Colorado's Department of State. For more information, see Section 2.6.

2.4 Core Policies, Entities, and Functionalities

This section presents our definitional framework. First, we present definitions of the *types of entities* (Section 2.4.1) and *core functionalities* (Section 2.4.2) inherent to most voter registration systems. Then, turning to jurisdiction-specific aspects of voter registration, we non-exhaustively define core *parameters* (Section 2.4.3) and *security policies* (Section 2.4.4) that are determined according to jurisdictional policy. The jurisdictional parameters and policies serve to *instantiate* the core functionality definitions to match with concrete implementation and security needs in a particular jurisdiction.

Later, in Section 2.5, we present security definitions that build upon the entities, core functionalities, and jurisdictional policies defined in this section.

2.4.1 Entities

We identify six types of entities that are involved in most voter registration systems. The specific lists of entities that belong in each category will vary between jurisdictions.

We use the term "entity" to encompass individuals, organizations, and hardware/software systems (such as devices or databases). This is a convenient shorthand that is common in the security literature;[8] however, we emphasize that devices, systems, and organizations *do not act of their own accord*, and responsibility for their management and conduct must be ascribed to individuals via well-defined chains of

[8]The security literature usually uses the term "parties" rather than "entities," but we prefer the term "entities" here in order to avoid confusion with political parties in the elections context.

responsibility according to jurisdictional policy (as further discussed in Section 2.4.4).

- ***Voters:*** People who are legally allowed to cast a vote in the corresponding jurisdiction (possibly limited to particular kinds of elections).[9]

- ***Election infrastructure:*** All entities affiliated with — and controlled by or answerable to — the election office. We highlight three common types of sub-entities:

  - *election officials*, who are responsible for conducting elections, including maintaining the lists of voters and of those who are eligible to vote;
  - *poll workers*, who work for election officials to aid in conducting a specific election, and typically have much more limited expertise, system access, and responsibilities: e.g., confirming voter eligibility in pollbooks and issuing provisional ballots in case a voter's eligibility cannot be determined; and
  - the *voter registration database (VRDB)*, where voter records are stored.

  These sub-categories are non-exhaustive. Election infrastructure systems may be run by the government, external contractors, or a combination of both.

- ***Maintenance entities:*** Entities external to the election office, who work with election officials to maintain voter registration lists. Each jurisdiction has its own list-maintenance strategies, but common maintainer entities in the U.S. are the United States Postal Service (USPS) via the National Change of Address system (NCOA), a state's Department of Motor Vehicles and Bureau of Vital Statistics, and other states' VRDBs via the Electronic Registration Information Center (ERIC).

- ***Oversight entities:*** Entities external to the election office, who examine voter data or other components of a voter registration system, in order to verify that

[9]While some legal systems may define electors as those eligible to vote and voters as those who actually vote, we use the colloquial definition of voter as one who is eligible to vote throughout this paper.

the voter registration system is operating as intended. There may be three types of oversight entities:

  - ***general*** *oversight entities* who, on their own initiative, examine publicly available information; and
  - ***designated*** *oversight entities* who, on their own initiative, examine non-public information that is available to them because they meet certain generally applicable criteria; and
  - ***official*** *oversight entities* who, on request from or under contract with an election office, examine non-public information made available to them for the purpose of a system review or audit.

  Designated and official oversight entities are relatively rare in practice, at least in the United States. Watchdog organizations interested in monitoring voter registration are more common, and can be considered general oversight entities. Definitionally, any member of the public can be a general oversight entity; however, we consider the term useful to refer to those entities that actually *do* (not only *could*) engage in oversight activities.

  Oversight mechanisms *within* the election office are also important: e.g., internal logging, auditing, and accountability procedures. We refer to entities involved in such internal oversight as part of the election infrastructure rather than as separate "oversight entities."

- ***Intermediaries:*** All other entities that handle voter registration data at any point during registration, updating registration, proving registration, or maintenance and oversight of a voter registration system. E.g., an organization like vote.org that helps register voters by mail.

- ***The public***: All entities, whether listed above or not. (The scope of "the public" is not jurisdiction-specific and includes foreign entities.)

A given entity may fall within multiple of the above categories, depending on the context. For example, USPS serves as a maintenance entity when aiding states in

the process of finding voters who moved out of state, and it can also serve as an intermediary when a voter mails paper registrations to their election official.

2.4.2 Core functionality modules

Next, we define five *modules* that together make up the core functionality of a voter registration system. These modules represent the basic components that our research has found common to most voter registration systems. Real-world voter registration systems can be thought to *implement* these modules while taking into account jurisdiction-specific policy decisions and constraints. Real systems may also contain additional functionalities not described here; our model is intended to be inclusive rather than comprehensive.

The line between the voter registration system and other parts of an election system (e.g., casting and tallying systems) is not clear-cut, as many parts of the broader election system interact with the registration system. For this work, we focus on aspects of election infrastructure that more directly concern registration, as described by the following modules.

- Registration: The processes involved in checking an individual's eligibility to vote when their information is not already in the VRDB, and if they are determined to be eligible, entering their information into the VRDB.

- UpdateRegistration: The processes involved in applying voter-initiated edits to a voter record that is currently present in the VRDB. Note that this includes a voter removing themselves from the VRDB.

- ProveRegistration: The processes involved in determining whether an individual is registered to vote, based on information that the individual presents for this purpose (e.g., when "checking in" at a polling place).

- Maintenance: The processes involved in election officials (with the aid of maintenance entities) editing, marking inactive, or removing voter records in the VRDB, without initiation by the concerned voter(s).

- Oversight: The processes involved in oversight entities assessing voter records and identifying discrepancies (such as voters who were incorrectly marked inactive), alerting either the public or election officials.

Section 2.5 describes each module in much more detail, framed as an interactive protocol parametrized by jurisdictional policies, and defines security properties for each module.

2.4.3 Jurisdictional parameters

In this section, we outline the core *jurisdictional parameters* of voter registration systems, which describe the variables of voter registration systems that vary across jurisdictions. Many of these parameters result from law or policy decisions that vary by jurisdiction. Jurisdictional parameters could include, but are not limited to, the following:

- p_{elig}: the voter eligibility criteria
- $p_{\text{reg-acts}}$: required actions from the voter in order to register
- $p_{\text{reg-methods}}$: the list of registration methods, such as the DMV, election office, registration website, etc. In particular, those that support AVR (typically only the DMV) get marked as such.
- $p_{\text{voter-info}}$: the types of voter information that are collected and stored
- $p_{\text{freeze-reg}}$: the period before election during which new registrations may not be processed
- $p_{\text{freeze-db}}$: the period before election during which systematic registration removals or maintenance are not allowed
- $p_{\text{keep-logs}}$: the period after an election for which a snapshot and activity logs of the VRDB for that election are kept[10]

[10]U.S. Federal law requires voter registration records to be kept for at least 22 months after a federal election [1].

- p_{auth}: the voter authentication criteria: How voters are authenticated for various stages of the VRDB process (registering to vote, updating voter registration record, checking in at a pollbook)

We refer to the ACM report [18] for a thoughtful policy perspective on how to set these parameters. To keep the scope manageable and to separate the technical from the policy aspects, in this paper, we do not suggest specific jurisdictional parameters. Instead, we focus on how to securely implement a voter registration system conditioned on given jurisdictional parameters. Whatever the jurisdictional parameters, secure implementation is an important goal.

2.4.4 Security policies

In addition to jurisdictional parameters, the rest of the jurisdiction specific details come in the form of *security policies*. A security policy governs the operations of a VRDB that affect its security. In the descriptions below, we outline a few items that would make sense to include in each security policy. Note that we can not aim to provide an exhaustive list of items contained of each security policy: since these are different across jurisdictions, there is no such thing as a complete description of each policy. Hence, we limit ourselves to a few important elements that serve as examples.

The different types of security policies relevant to voter registration are the following:

1. $\mathbb{P}_{\text{access}}$ denotes the ***access control policy***, which specifies which voter data specific entities may access.[11]

    - Types of voter information that are public
    - Description of which pieces of voter data are available to which entities
    - Whether there's an option to hide certain fields of a voter's information upon application for privacy/safety reasons

[11] Unlike the other security policies in this section, access control policies have been extensively studied, see, e.g. [39].

2. $\mathbb{P}_{\text{sys-chg}}$ denotes the ***system change control policy***, which specifies how election officials may modify the system, such as changing the system configuration, security policies, and database design.

    - How often is the system evaluated for upgrades?
    - Who needs to grant authorization before a system change?
    - What is the specific sequence of steps for implementing a system change?
    - What are the backup plans in case parts of the system go down during a system change?

3. $\mathbb{P}_{\text{data-chg}}$ denotes the ***data change control policy***, which specifies steps election officials must take when changing voter data, including authorization, execution, and logging.

    - Who needs to authorize a change of voter data?
    - What type of data can be changed?
    - Who triggers a change of voter data?
    - Who actually modifies the voter data?
    - How are such changes logged?

4. $\mathbb{P}_{\text{data-use}}$ denotes the ***voter data use policy***, which specifies guidelines related to uses to which public and non-public voter info can be put.

    - Which pieces of voter data are available for which uses?
    - What use cases are prohibited?

5. $\mathbb{P}_{\text{notif}}$ denotes the ***voter notification policy***, which specifies how jurisdictions notify voters when their data or registration status changes.

    - List of events for which a voter must be notified
    - Protocol by which voters are notified, including the amount of time a voter has to respond to a notification, if necessary, and the resulting action

- Methods by which voters are notified

6. $\mathbb{P}_{\text{maint}}$ denotes the ***maintenance policy***, which specifies how election officials ensure voter records are accurate and up-to-date.

   - Reasons for which voter registrations may be updated (change of address, etc), marked inactive (moved out of state, voter inactivity), or cancelled (death, crime, etc)[12]
   - Specific events or thresholds that trigger such maintenance actions (e.g. number of elections before voter is declared inactive)
   - Data sources used to inform maintenance

7. $\mathbb{P}_{\text{oversight}}$ denotes the ***oversight policy***, which specifies how third parties can review information in the VRDB.

   - Who can oversee the database
   - How oversight entities authenticate to the election official
   - Level of access given to oversight entities
   - Points at which oversight entities may review the VRDB (e.g. pre-election, post election, continuously)
   - How jurisdictions conduct internal audits, including security incident detection and response protocols

In summary, this section introduced the core elements of a voter registration system in the form of abstract functionality modules. We also described the general entities involved in the system, and defined security policies and parameters that enclose the fundamental differences across jurisdictions. In the subsequent sections, we will tie these elements together as we expand on the descriptions of these modules, which will be a function of the entities and policies.

[12]For a list of maintenance practices by U.S. states, see NCSL [56].

2.5 Detailed Model and Security Properties

This section provides a detailed modeling of each core functionality module (introduced in Section 2.4.2) of a voter registration system. First, we specify the categories of interacting entities, jurisdictional parameters, and communication patterns inherent to each module. Then, for each module, we enumerate security properties parametrized by jurisdictional security policies.

We model each core functionality module as a simple interactive protocol between entities: e.g., between a voter and the voter registration database (VRDB), possibly via intermediaries. Entities communicate with each other via *communication channels*: e.g., online, mail, or in-person communication. The VRDB can only be *directly* accessed by election infrastructure entities. Each protocol (i.e., module) is parametrized by relevant entities and communication channels, and takes as input voter data. For example, the registration module is parametrized by T, the entity through whom the voter is registering, and C the channel through which the voter communicates with T, and takes as input some voter data S.

2.5.1 $\mathsf{Registration}^{T,G}_{C,C',C'',C'''}(S)$

A member of the public, acting either directly by interacting with an election official or communicating via an intermediary, submits an application containing required information. The information is then reviewed by the election official, and if the voter is determined to be eligible and the submitted data determined to be accurate, the election official adds the voter's information to the VRDB. Information about the outcome of this process may then be communicated back to the applicant. Voters may only be permitted to register during certain time periods, as defined in the jurisdictional policy. In detail:

1. The voter sends some personal information S that contains a signature S' (determined by p_{auth}, $p_{\text{voter-info}}$, and $\mathbb{P}_{\text{access}}$) to an intermediary T (contained in $p_{\text{reg-methods}}$) via a communication channel C (e.g., in-person, mail, or the Internet, as determined by $p_{\text{reg-methods}}$). If registering in person at the election office

or via an official web portal, T is empty ($\perp$).

2. If $T \neq \perp$, then T forwards S and S' to an election infrastructure entity G via communication channels C' and C'', respectively. (If $T = \perp$, the voter is communicating their data directly to G.)

3. G verifies that the submitted data meets the criteria outlined in p_{elig}, and that the registration was submitted during an eligible timeframe, as defined in $p_{\text{freeze-reg}}$.

4. G then calls Maintenance (S), i.e., it triggers a subroutine to verify the registration information via third parties (if needed), following the list maintenance protocol defined in section 2.5.4 for the specific voter[13].

5. If all checks pass, G stores the voter's data in the VRDB, following the guidelines from $\mathbb{P}_{\text{data-chg}}$. Lastly, G sends a notification N to the voter through a communication channel C''', as outlined in $\mathbb{P}_{\text{notif}}$, confirming that the registration was successful (if unsuccessful, the verification subroutine from the prior step would send a notification to the voter).

The workflow of the Registration module is shown in Figure 2-1. Note that in practice, registrations may not sent directly, one at a time, from T to G: e.g., they might be sent in batches instead. Our model captures the basic information flow of the module and omits such implementation details, for clarity of presentation.

2.5.2 $\mathsf{UpdateRegistration}^{T,G}_{C,C',C'',C'''}(I, N)$

In order to update their record, the voter notifies the election official of a desired change, such as a change of address or name. Operating within the data change

[13]The process of verifying voter data during registration is very analogous to the list maintenance process, i.e., it interfacing with third-party entities. For example, verifying that a voter's address is correct and checking if a voter changed states may involve external communication with, e.g., USPS in both cases. Even though the specific information sent to these maintenance entities may change, the high-level behavior of these two processes is similar. Hence, for simplicity, we model the verification subroutine as a call to the Maintenance module.

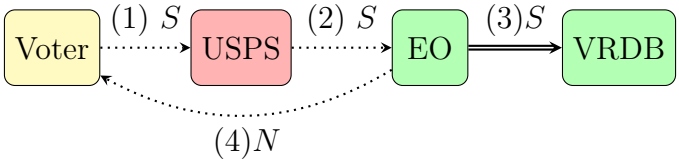

Figure 2-1: Example Registration flow. Here, T = DMV and G = EO (election officials). Dotted and double arrows indicate using mail and internal networks as communication channels, respectively. The call to Maintenance is left implicit.

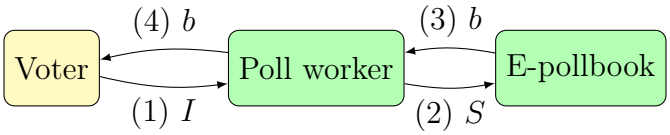

Figure 2-2: Example ProveRegistration flow. Here, G = poll worker and G' = e-pollbook. All communication channels are in person.

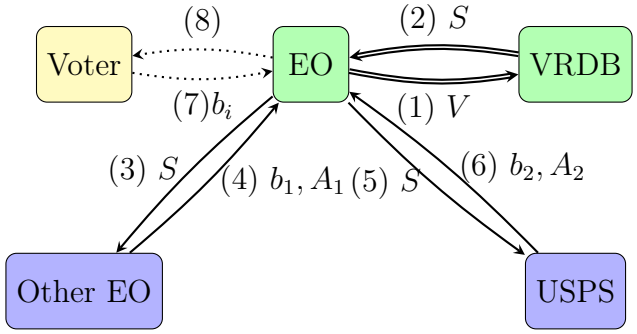

Figure 2-3: Example Maintenance flow. Here, G = election office, M_1 = other EO (i.e., election officials from another state), and M_2 = USPS. Dotted, double, and bold arrows indicate using mail, internal networks, and the Internet as a communication channels, respectively

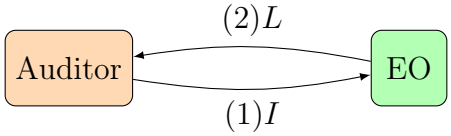

Figure 2-4: Example Oversight chain. In this case, A = auditor and G = EO (election officials). All communication channels are the Internet.

control policy $\mathbb{P}_{\text{data-chg}}$, the election official authenticates this change and updates the voter's record accordingly.

The workflow of updating a registration is much like that of the registration module defined above, with some small differences: instead of sending all their data in step (1), voters send an identifier I and just their new data N (e.g., a new address); T then uses I to authenticate the voter, and proceeds with the rest of the steps in the registration module. Given its similarity to the Registration module (Figure 2-1), the workflow of the UpdateRegistration module is not depicted separately.

2.5.3 $\mathsf{ProveRegistration}^{G,G'}_{C,C'}(I)$

A voter must prove that they are registered to vote in order to cast a ballot. The voter supplies information in accordance with the voter authentication criteria to the poll worker, who authenticates the voter and confirms that they are an active voter in the pollbook. In the case that an election official is unable to confirm a voter's registration, they may provide the voter with a provisional ballot,[14], in which case the voter's registration is validated after the ballot is provisionally submitted. In the case of absentee ballots, a voter remotely authenticates themselves to the election official, e.g. by providing a signature. In more detail:

1. The voter sends some identifying information I (determined by p_{auth} and $\mathbb{P}_{\text{access}}$) to election infrastructure entity G (e.g., a poll worker, election official or a web portal) via a communication channel C.

2. G forwards I (as indicated in $\mathbb{P}_{\text{access}}$) to another election infrastructure entity G' (such as an electronic pollbook or the election office), which verifies if I corresponds to a valid, eligible voter by interacting with the VRDB via a communication channel C' (either an internal network or the Internet) or doing a local check (in the case of a pollbook).

3. G then sends a bit b to the voter through the original channel C; if $b = 1$, the voter proceeds to vote (this is now a separate system, outside the scope of voter

[14]This is required by federal law in the United States [6].

registration). If $b = 0$, G provides the voter with a provisional ballot, and the voter's registration is validated after the ballot is cast.

Starting the voting process through mail or the Internet represents requesting an absentee ballot, while in-person represents physically going to the polling center. For the latter, information verification tends to happen with an (e-)pollbook, which either checks the information locally (if the VRDB is downloaded a priori) or contacts the VRDB via the Internet. The workflow of the ProveRegistration module is shown in Figure 2-2.

2.5.4 $\mathsf{Maintenance}^{G,\{M_i\}}_{C,\{C_i\},C'}(V)$

Election officials perform maintenance activities on their VRDB. In the United States, certain maintenance is required under the National Voter Registration Act [5]. Maintenance activities may include updating records of voters who have moved and removing ineligible or inactive voters, and often occur based on communication with external maintenance entities. Maintenance activities may be paired with notifications to voters, as determined by the voter notification policy. In more detail:

1. When indicated by $\mathbb{P}_{\text{maint}}$, an election infrastructure party G acquires some information S (following $\mathbb{P}_{\text{access}}$) for a specific voter (specified by an identifier V, e.g., an SSN, following p_{auth}) from the VRDB via a communication channel C (internal network). If the input to the module is the full voter data S itself (in the case of a voter verification subroutine), skip this step.

2. G sends S (following the guidelines of $\mathbb{P}_{\text{access}}$) to various maintenance authorities $M_1, ..., M_n$, defined in $\mathbb{P}_{\text{maint}}$, via communication channels $C_1, ..., C_n$. Each M_i, after doing local checks, replies with a bit b_i (which identifies if, for example, the voter is alive or still at their same address) and some auxiliary data A_i (e.g., the voter's new address).

3. If $b_i = 0$, G updates, marks inactive, or removes the voter from the VRDB via C (following $\mathbb{P}_{\text{data-chg}}$), and sends b_i to the voter via a communication channel

C' (i.e., upon seeing $b_i = 0$, the voter knows that their registration got deleted from the VRDB). If $b_i = 1$, G does not do anything. If $b_i = null$ then a *voter-confirmation subroutine* gets triggered in accordance with $\mathbb{P}_{\text{notif}}$: contact the voter some number of times to try to confirm registration info; fails if no response or bad response.

The Maintenance module's workflow is shown in Figure 2-3. In practice, the maintenance protocol may be non-interactive (e.g., the maintenance entity simply sends their data to G).

2.5.5 $\mathsf{Oversight}_C^{A,G}(I, L)$

Oversight entities (as defined by the oversight policy) may access voter data in accordance with the jurisdiction's oversight and access control policies. The oversight entities may assess voter records and identify discrepancies (such as voters who were incorrectly marked inactive), and inform the public and/or election officials of their findings. Election officials may accept the claims and issue corrective actions or refute the claims (ideally, with supporting evidence). Next, we describe this process in more detail for designated or official oversight entities (putting aside general oversight entities since they only access public information, as defined in Section 2.4.1):

1. An oversight entity A sends some identifying information I to an election infrastructure entity G via a communication channel C.

2. G checks if the request is coming from a valid oversight entity (as specified in $\mathbb{P}_{\text{oversight}}$), and verifies I. In addition, G also confirms that this oversight entity is allowed to review the database at this point in time, as specified in $\mathbb{P}_{\text{oversight}}$, too.

3. If all checks pass, G sends a subset L of the voter registration list (as permitted by $\mathbb{P}_{\text{access}}$ and $\mathbb{P}_{\text{oversight}}$) to A through the original channel C.

The workflow of the Oversight module is shown in Figure 2-4.

2.5.6 Security Properties

Next, we present security properties applicable to each of the core functionality modules, followed by some broader systemwide security properties. As usual, since jurisdictions differ in their voter registration policies, these are a function of the relevant jurisdiction's parameters and security policies.

- Registration
  - ***Completeness:*** An eligible voter possessing the requisite proof of eligibility must be able to register their accurate information in the VRDB only once and only during the periods in which new registrations are allowed, as determined by $\mathbb{P}_{\text{data-chg}}$ in accordance with p_{auth}, p_{elig}, and $p_{\text{freeze-reg}}$.
  - ***Soundness:*** Nobody must be able to register incorrect information or ineligible voters to the VRDB. This is governed by $\mathbb{P}_{\text{data-chg}}$, and assessed by voters (via $\mathbb{P}_{\text{notif}}$), election officials (via $\mathbb{P}_{\text{maint}}$) and by oversight entities (via $\mathbb{P}_{\text{oversight}}$).
  - ***Secrecy:*** Only entities authorized under $\mathbb{P}_{\text{access}}$ to access (specific types of)information submitted by applicants may learn such information during the registration process.
- UpdateRegistration
  - ***Completeness:*** Any eligible, registered voter must be able to update their existing registration with their correct information, and to delete their VRDB record, subject to $\mathbb{P}_{\text{data-chg}}$ in accordance with p_{auth}.
  - ***Soundness:*** Nobody must be able to (1) update a VRDB record that they are not authorized to update under $\mathbb{P}_{\text{data-chg}}$, or (2) edit any VRDB record to contain incorrect information. As with soundness of registration, this is governed by $\mathbb{P}_{\text{data-chg}}$, and assessed by voters (via $\mathbb{P}_{\text{notif}}$), election officials (via $\mathbb{P}_{\text{maint}}$) and by oversight entities (via $\mathbb{P}_{\text{oversight}}$).

- ***Secrecy:*** Only entities authorized under $\mathbb{P}_{\text{access}}$ to access (specific types of)information submitted by applicants for updates, and to access (specific types of) VRDB data, may learn such information during the update process.

- ProveRegistration

  - ***Completeness:*** Any registered voter should be given access to the casting process (to vote at most once), according to p_{auth}.
  - ***Soundness:*** No ineligible voter should be given access to the casting process, according to p_{auth}.
  - ***Secrecy:*** Only entities authorized under $\mathbb{P}_{\text{access}}$ to access (specific types of) VRDB data may learn such information during the process of proving registration.

- Maintenance

  - ***Completeness:*** After a list maintenance update,
    * any VRDB record that a maintainer entity flags as possibly containing incorrect information or corresponding to a person who is not eligible to vote should trigger a voter communication as specified in $\mathbb{P}_{\text{notif}}$;
    * any other VRDB record must remain unchanged in the VRDB; and
    * if a voter notification about a flagged record results in timely voter feedback that demonstrates (in accordance with p_{auth}) that the voter is still eligible, and either confirms the information in the record is correct or provides updated correct information, then the record must remain in the VRDB.
  - ***Soundness:*** After a list maintenance update,
    * any record that all maintainer entities flag as possibly incorrect or ineligible must be marked as such in the VRDB, and

    - any record flagged by a maintainer entity as possibly incorrect or ineligible, where the follow-up voter communication does *not* result in timely voter feedback that demonstrates eligibility and correct information,

    must be marked as such in the VRDB in accordance with $\mathbb{P}_{\text{maint}}$.

  - ***Secrecy:*** Only entities authorized under $\mathbb{P}_{\text{access}}$ to access (specific types of) VRDB data may learn such information during list maintenance.

- Oversight

  - ***Completeness:*** Any oversight entity must be able to learn the information that $\mathbb{P}_{\text{oversight}}$ authorizes it to access for oversight purposes. There should be an appeal process in case they cannot do so.

  - ***Soundness/secrecy:*** No oversight entity must be able to learn any information that it is not authorized to access under $\mathbb{P}_{\text{oversight}}$ and $\mathbb{P}_{\text{access}}$.

2.6 Policy implementations

In this section, we show how our model of voter registration systems, presented in Sections 2.4 and 2.5, can be instantiated with concrete jurisdictional parameters to represent a real-world system. We propose a structured format for jurisdictional information by presenting a series of tables (Tables 2.1–2.7), and provide a case study for the state of Colorado. The tables may be expanded and customized for different jurisdictions; we present just the core components needed to capture the jurisdictional parameters and policies described in Sections 2.4 and 2.5.

Up to this point, our definitions abstract away the core jurisdiction-specific details in general jurisdictional parameters and security policies. This approach is beneficial when modelling voter registration in general, as we can think of these policies and parameters as abstract objects and ignore the specifics of how these would look for a specific jurisdiction (e.g., when designing cryptographic protocols related to voter registration). However, when analyzing a specific jurisdiction, we can explicitly express

p_{elig}	U.S. Citizen, resident of Colorado for at least 22 days, at least 16 years old, and not serving felony sentence
$p_{\text{reg-acts}}$	None (for automatic voter registration), otherwise submit voter registration application
$p_{\text{reg-methods}}$	Online, email, fax, mail, in person
$p_{\text{voter-info}}$	See access control policy table
$p_{\text{freeze-reg}}$	8 days before election (mail/online), up to and on election day (in person). County election officials may choose to process registrations submitted later than 8 days.
$p_{\text{freeze-db}}$	N/A
$p_{\text{keep-logs}}$	At least 2 years
p_{auth}	Updating record: Date of birth/driver's license number or last 4 digits of social security number, signature. Looking up record online: Name, zip code, birthday Checking in at pollbook: 1 form of ID Vote by mail: signature, if first time may need to provide copy of ID

Table 2.1: **Jurisdictional parameters for Colorado.**

these policies and parameters as a function of jurisdictional parameters. As such, our definitions can be thought of as a *framework* that one can use to get an abstract representation of a jurisdiction's voter registration system, where the inputs to the framework are the jurisdictional parameters and security policies from Sections 2.4.3 and 2.4.4.

In order to validate our framework's applicability and realism, we conducted a case study based on Colorado's voter registration system: constructing policy tables for each of the seven security policies. In doing so, we consulted publicly available documents on Colorado's existing laws and policies. We then checked the resulting policy tables for accuracy with an official from Colorado's Department of State.

In constructing the tables, we consulted existing laws and policies in Colorado. Part 5 of Title 1, Article 2 in Colorado Revised Statutes governs voter registration [14]. As part of a rulemaking process, the Colorado Secretary of State publishes its election rules, of which Rule 2 governs voter registration [29]. Beyond these, we consulted Colorado's voter registration form and technical requirements of its voter registration database [28].

We found that, in the case of Colorado, we were able to complete most information in the policy tables with publicly available information. This suggests that, for future use, policy tables can either (ideally) be published by the jurisdiction itself, or, if not, be constructed based on public information.

We hope that organizing jurisdictional information in the structured form that we propose, as demonstrated via the Colorado case study, may be helpful in order to:

- specify detailed jurisdictional-specific *threat models* for voter registration systems, which is helpful for security analyses and research;
- organize voter registration policy information for *convenient comparison* between jurisdictions, and learn about common and uncommon approaches;
- enhance *transparency* of voter registration systems, thereby promoting civic engagement and accountability;
- *identify strengths and weaknesses* of a particular jurisdiction's approach to voter registration security, which can inform where to focus resources for improvement;
- *identify underspecified aspects* of a particular jurisdiction's voter registration policies;
- *identify mismatches* between a jurisdiction's stated policies and its implementation of voter registration; and
- *encourage constructive dialogue* between election officials and the security research community regarding details of voter registration systems that are important to security analyses and research.

Category	Entity	Name	Home address, Mailing address	Birth year	Birth day	Phone	Email	Driver's License/ ID card number	SSN last four digits	Party, Affiliation Date, Gender	Sig.	Voting activity history
REGISTER/ UPDATE	Voter being registered	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	VRDB	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Online registration/update portal	✓	✓*	✓	✓	✓*	✓	✓	✓	✓	✓	✓
	NVRA agency (e.g., DMV)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	County clerk	✓	✓†	✓	✓	✓†	✓	✓	✓	✓	✓	✓
USE REG TO VOTE	County official (polling place)	✓	✓†	✓	✓	✓†	×	✓	×	✓	✓	×
	County official (mail-in ballots)	✓	✓†	✓	✓	✓†	×	✓	×	✓	✓	×
LIST MAINTENANCE	NCOA	×	×	×	×	×	×	×	×	×	×	×
	Department of Revenue	×	×	×	×	×	×	×	×	×	×	×
	ERIC	✓	✓	✓‡	✓‡	✓	✓	✓‡	✓‡	×	×	✓
TRANSPARENCY	The public	✓	✓*	✓	×	✓*	×	✓	×	✓	×	✓

Table 2.2: **Colorado Access Control Policy.** The access control policy determines which entities can access certain fields. We represent the access control policy as a table that maps entities to registration fields, with binary values in each cell denoting whether the entity in that row is allowed to view the data point in that column, for any voter.
*Hidden for address confidentiality program voters
† only accessible by designated address confidentiality program election staff
‡ hashed before sending to ERIC

For the **system change control policy**, Colorado does not publish information related to the system change control policy. See Table A.7 in the appendix for a template table. The system change control policy specifies all guidelines that must be followed when making meta changes to the voter registration system. We represent the system change control policy as a table that maps "types of changes" to stages of a change's lifecycle. Each cell specifies the directives that are in place at a particular stage of a (type of) change.

Oversight entity	Voter data	VRDB logs	VRDB code	Interactive access	Time periods
Nonprofit	Yes, as public in accordance with access control policy	No	No	No	Continuously
Political organization	Yes, as public in accordance with access control policy	No	No	No	Continuously
Third party pentester	No	Yes	Yes	Yes	Over 90 days before election
VoteShield	Yes, as public in accordance with access control policy	No	No	No	Continuously
Department of State	Yes	Yes	N/A	N/A	Continuously

Table 2.3: **Colorado oversight policy.** The oversight policy governs how third parties can review information in the VRDB. We represent the oversight policy as a table mapping oversight entities to the type of voter data and other information they can access, along with time periods for oversight.

Prohibited uses	Not specified.
Approved entities	Public
Information released	See access control policy table.
Opt out policy	ACP participants; confidential voters; pre-registrants.

Table 2.4: **Colorado voter data use policy.** The voter data use policy specifies limitations on how (and by whom) the data can be used. We represent the voter data use policy following the structure of [54].

Notification reasons	Notification protocol	Notification methods
Incomplete registration	Send notice via notification methods	Mail, Email (by county)
New registration	Send notice via notification methods. If returned as undeliverable, do not register. If not returned as undeliverable, register.	Mail, Email (by county)
Inactive registration	Send elector voter confirmation card at least 60 days before election via notification methods. If not returned or not marked undeliverable, and voter has not voted in two general elections, cancel registration.	Mail, Email (by county)
Address change	Send notice to new address.	Mail, Email (by county)
Cancelled registration	None	N/A

Table 2.5: **Colorado voter notification policy.** The voter notification policy governs how jurisidictions notify voters of various changes to their records. We represent the voter notification policies as a table mapping notification reasons to notification protocols and methods.

Reason	Data source	Threshold	Action
New driver's license or updated address	Department of Revenue	New driver's license or updated address	Register voter or update existing record
Moved in state	NCOA	Address changes in state	Update address
Moved out of state	NCOA / ERIC	Address changes out of state	Mark inactive – NCOA
Returned mail	County	Returned mail	Mark inactive – returned mail
Undeliverable ballot	County	Ballot could not be delivered	Mark inactive – undeliverable ballot
Voter inactivity	VRDB	Has not voted in past two elections	Mark inactive; cancel reg after two more inactive elections
Death	Dept. of Public Health and Env., Social Security Death Index	Voter dies	Cancel registration - deceased
Crime	Dept. of Corrections, CO U.S. Attorney's office	Voter currently incarcerated	Cancel registration - convicted felon

Table 2.6: **Colorado maintenance policy.** The maintenence policy governs how jurisdictions keep their VRDB accurate and up-to-date. We represent the voter maintenance policy as a table mapping maintenance reasons and their associated data sources to maintenance thresholds and actions.

Category	Entity	Type of Data
AUTHORIZATION	Voter	Personal data
	State and county election officials	Data from list maintenance update
TRIGGER	Online update portal	Data from voter who started update
	Mail	Data from voter who started update
	ERIC	Data of voters in other states
	Department of Revenue (DMV)	Data of new/updated license
	Other NVRA agency	Data of new voter
	NCOA	Data of voter move
	Department of Public Health and Environment, Social Security Death Index	Data of death
	Colorado Department of Corrections, Colorado U.S. Attorney's office	Voters who committed crime
EXECUTION	State election officials	
	County election officials	

Table 2.7: **Colorado data change control policy.** The data change control policy includes information about the entities involved in updating the VRDB or associated policies. We represent the data change control policy as a table that specifies the entities allowed to authorize/start updates, trigger updates (send updated data to election officials), and execute the update (directly modify the data inside the VRDB). In this table, we map these entities to the type of data they update, and if there is a notification involved in this type of update.

2.7 Critical Questions

In this section, we propose questions that policymakers, election officials, security practitioners, and researchers may wish to ask to evaluate candidate systems. We categorize our questions with respect to our varying policies presented above. We stress that this list is incomplete; our goal these questions can help foster discussion and in-depth evaluation of proposed and already deployed systems.

- ***General questions:***
  - Main question: **is there a (security) mechanism enforcing each item of every security policy?**
  - Is the voter registration system compatible with (1) the jurisdiction's expressed policy choices? (2) the framework's rigorous security definitions?
  - Are there undefined or incomplete portions of any policies?
  - Do the policies completely encompass how the voter registration system should work?
  - Is the voter registration database regularly audited to ensure that the policies outlined are enforced programmatically?
  - Are there reliability mechanisms in place in case any of the security policies gets violated?
  - How might external developers, researchers, and government agencies help improve the system?
  - Are there security mechanisms in place to enforce the security properties of each module?
- ***Access control policy:***
  - Is the access control policy in compliance with laws regulating voter data access?

  - Are there fields of voter data that are made accessible to third parties even though they are not required to by law?
  - Does the access control policy follow the *principle of least privilege*?

- ***System change control policy:***
  - Are system changes regularly audited to ensure that no unauthorized changes have been made?
  - Is the system change control policy followed every time there is a system change?
  - Does the system change control policy follow the *principle of least privilege*?
  - Is the system (including security policies) constantly evaluated for potential updates?

- ***Data change control policy:***
  - Are changes to voter data regularly audited to ensure that no unauthorized changes have been made?
  - Is there sufficient logging (at the application, network, and operating system level) to determine who made a change in the case of an unauthorized change being detected?
  - Is the data change control policy followed every time data gets changed?
  - Are there enough reliable backups of the VRDB in case voter data gets tampered with?

- ***Voter data use policy:***
  - How are third parties assessed and held accountable for incorrect uses of voter data?
  - Is there a rigorous evaluation process to authorize external entities to use non-public voter data, if applicable?

  - Are voters given the option to opt-out of their data being used for specific purposes, especially if it would threaten their safety?

- ***Voter notification policy:***
  - Is there enough redundancy in the notifications sent to voters in case they missed the first one(s)?
  - Is the notification policy working in harmony with the data change control policy and the maintenance policy?
  - Are replies from voters processed in an efficient and timely manner?

- ***Maintenance policy:***
  - What transparency practices are in place to allow third parties to audit maintenance activities?
  - Is voter data verified in-depth when maintenance activities indicate it should be removed/updated?
  - Are state-of-the art maintenance technologies like ERIC being used?
  - Are maintenance entities (e.g. ERIC or other states) regularly assessed for correct behavior?

- ***Oversight policy:***
  - Is external oversight encouraged and advertised?
  - Is there a proper channel through which oversight entities can notify election officials of suspected irregularities?
  - Is there a timely and well-defined procedure to investigate and resolve potential irregularities found by oversight entities?

2.8 Conclusion

In this chapter, we provide the first systematic formalization of *voter registration systems* as they exist today. We defined the entities and core functionalities inherent

in most voter registration systems, the jurisdictional policies that constrain specific implementations, and key security properties. As a tool for mapping our general definitions to real-world implementations, we provided a series of tables organizing jurisdiction-specific policy information, illustrated with a case study of Colorado. Finally, we offer a critical question list.

Though voter registration is a fundamental part of secure elections, it is often comparatively understudied. One contributing factor may be the lack of detailed understanding of problem definitions, practical constraints, and security issues. It is our hope that our work can help promote the study, development, and adoption of new practical systems.

Chapter 3

Increasing Transparency in Voter Registration

The accuracy of electoral rolls is one of the core pillars of voter trust. That is, after a member of the public registers to vote, it is fundamental for them to feel certain that their data is correctly stored in the voter registration database (VRDB). Without this, their confidence in the entire system (and election outcomes) is tarnished. Unfortunately, simply registering to vote does not provide this certainty.

For example, VRDBs are a target for potential voter purges, i.e., the removal of voters for illegitimate reasons, potentially discriminating against particular groups [13]. According to a Brennan Center for Justice report, between 2014 and 2016 states purged nearly 16 million voters from registration lists [23]. Some of the states where this occurred do not allow for same-day voter registration, so voters may have been unable to vote.

As a second example, VRDBs are one of the most enticing targets for external attackers who may wish to interfere with an election. According to U.S. intelligence officials, around the time of the 2016 US elections Russian hackers targeted all 50 states' VRDBs, succeeding in at least two states [11]. In one of these, the hackers exfiltrated hundreds of thousands of records before being caught [11]. Besides accessing sensitive voter data, there is also the threat that hackers could modify voter records.

These two examples, and the value of government accountability to the public,

highlight the importance of greater transparency surrounding voter registration systems. Even in situations were these attacks do not alter the final outcome of the election, lack of transparency impacts confidence on the legitimacy of the results and the system as a whole.

Current registration systems have a gap that needs to be addressed: there is no reliable way for election officials to prove to voters that their data is in its desired state. Using tools from cryptography—most notably, from the transparency logs [30] literature—our ongoing work is in the process of designing a voter registration database that supports this functionality.

3.1 High-Level Design Ideas

In this section, we describe some of the main ideas of our design. The current draft of this project is more fleshed out than what is written here but, due to its current rough state, has been omitted from this thesis.

The central primitive of our design will be an append-only data structure [16, 30, 49, 69], where records can be added but not modified nor deleted. Furthermore, such a data structure supports membership proofs and "append-only" proofs, which can be verified by members of the public to ensure that a particular entry is present in the log and that the log is being used in an append-only manner. For example, many of the implementations of append-only data structures are based on Merkle Trees, which rely on the security of cryptographic hash functions to support these types of proofs.

Drawing inspiration from the key transparency literature [36, 49], the main idea is to store voter registration data on an auxiliary append-only log, subject to various jurisdictional policies concerning voter data. Thus, our design must be amenable to jurisdiction-specific adaptation: our data structure must flexibly support functionalities like obfuscating certain fields and not others, completely hiding a record for voters in a witness protection program [4] while still letting them verify their data, giving specific, custom-formatted data to third-party organizations, etc.

When a voter first registers, they get assigned a random, unique identifier. Then, their data gets stored in the log, using the identifier as a key, with private fields obfuscated (e.g., hashed or encrypted). Lastly, voters are given a paper receipt with their unique identifier and the data they registered. Subsequent modifications to their data will result in completely new entries in the log, which get stored under the same key (i.e., voter identifier). In particular, deregistering results in a new entry with a special symbol denoting that said voter is not currently registered. Thus, we have a historical log of transactions on the voters' registration data.

At any point in time, voters can contact election officials and request a membership proof for their data. Even if election officials are not trusted, the security of the cryptographic data structure guarantees that voters will be able to detect if their data has been tampered with, as forging these proofs requires breaking the security of primitives such as cryptographic hash functions. Similarly, external auditors can request append-only proofs and ensure that the log has been used in an honest manner, and that entries have not been pruned from it.

Various external organizations (e.g., ERIC [2]) may require data for different purposes, most notably for list maintenance. Thus, for each approved party, we store a copy of the data in the custom format that they expect, alongside the voter's main data. Thus, by relying on voters verifying their data, external organizations can have greater confidence that the data they receive from states is accurate and complete.

3.2 Future Work

The main steps towards completing this project are to finish formalizing our protocol design, implement it, and run various experiments to get a sense of how our design scales to jurisdictions with millions of voters. The formal write-up is in a very rough state, but the protocol is almost complete; the biggest work left to do concerns interfacing with external parties, as there are various edge cases that we are still thinking through (e.g., how can we efficiently support adding new third parties?). Lastly, one additional extension we are considering is to provide a means for the public to verify

that list maintenance was performed properly; even if the election officials nor the deduplication authorities are trusted, we want to verify that no voters were illegally (un)pruned. This is a more speculative and preliminary direction than the rest of the work, but we are drawing inspiration from private set intersection protocols [34, 63], and are thinking of ways to design a malicious-secure protocol that interfaces with our new data structure.

Chapter 4

Conclusion and Future Work

In this thesis, we present the first formal treatment of voter registration by providing a framework that abstractly models these systems. This framework can be useful to study voter registration systems in a general sense, or to draw insights about specific implementations of these. Furthermore, we hope it provides clarity on how these complex systems work, and that it motivates further academical work in this are.

In addition, we outline a second project, currently a work in progress, that aims to increase transparency in voter registration systems. By providing a means for voters to verify their data without requiring trust in any other parties, we provide a detection mechanism for illegitimate data changes, and thus increase voter confidence.

These lines of work target two important gaps in the study of voter registration systems, the former from a definitional side and the latter from a technical one. The meta-conclusion from these projects is that there is a lot of opportunity to use tools from cybersecurity and cryptography to strengthen voter registration systems; we hope that this thesis brings more academic attention to a fundamental link in the electoral process, and inspires more research work in this field.

Appendix A

Additional Tables

We present template security policy tables that researchers can complete for jurisdictions in the subsequent pages.

Category	Entity	Type of Data
AUTHORIZATION	Voter Election officials	
TRIGGER	Online update portal Mail State agencies for update: (e.g., DMV) [enter] Other NVRA agency List maintenance mechanism: (e.g., ERIC): [enter] Other: [enter]	
EXECUTION	State election officials County election officials	

Table A.1: Template data change control policy.

Prohibited uses	
Approved entities	
Information released	See access control policy table.
Opt out policy	

Table A.2: Template voter data use policy.

Notification reasons	Notification protocol	Notification methods
Incomplete registration		
New registration		
Inactive registration		
Address change		
Cancelled registration		

Table A.3: Template voter notification policy.

Reason	Data source	Threshold	Action
New driver's license or updated address	Department of Revenue		
Moved in state	NCOA		
Moved out of state	NCOA / ERIC		
Returned mail	County		
Undeliverable ballot	County		
Voter inactivity	VRDB		
Death	Registrar of Vital Statistics Social Security Death Index		
Crime	Dept. of Corrections		

Table A.4: Template maintenance policy.

Oversight entity	Voter data	VRDB logs	VRDB code	Interactive access	Time periods
Nonprofit					
Political organization					
Third party pentester					
Other: [enter]					

Table A.5: Template oversight policy.

Category	Entity	Name	Home address, Mailing address	Birth year	Birth day	Phone	Email	Driver's License/ ID card number	SSN last four digits	Party, Affiliation Date, Gender	Sig.	Voting activity history
REGISTER/ UPDATE	Voter being registered VRDB Online registration/update portal NVRA agency (e.g., DMV) County clerk Other: [enter]											
USE REG TO VOTE	County official (polling place) County official (mail-in ballots) Other: [enter]											
LIST MAINTENANCE	State agencies for maintenance: [enter] Other third parties for maintenance: [enter]											
TRANSPARENCY	The public Other: [enter]											

Table A.6: Template access control policy.

Type	Description	Planification	Evaluation	Review	Authorization	Execution	Communication	Logging	How often system is evaluated for updates
STANDARD									
MINOR									
MAJOR									
SIGNIFICANT									
EMERGENCY									
SECURITY POLICY CHANGE									

Table A.7: Template system change control policy.

Bibliography

[1] 52 USC 20701: Retention and preservation of records and papers by officers of elections; deposit with custodian; penalty for violation. URL: https://uscode.house.gov/view.xhtml?req=granuleid:USC-prelim-title52-section20701&num=0&edition=prelim.

[2] Electronic Registration Information Center. URL: https://ericstates.org.

[3] VoteShield. URL: https://voteshield.us.

[4] Witness Security Program. URL: https://www.usmarshals.gov/witsec/.

[5] National Voter Registration Act of 1993. URL: https://www.congress.gov/bill/103rd-congress/house-bill/2, 1993.

[6] The Help America Vote Act of 2002. URL: https://www.congress.gov/107/plaws/publ252/PLAW-107publ252.pdf, 2002.

[7] Software glitch discloses Wokingham edited electoral register. URL: https://www.bbc.com/news/uk-england-berkshire-27304885, May 2014.

[8] Three Welsh councils' electoral roll data breaches probed. URL: https://www.bbc.com/news/uk-wales-south-east-wales-27159648, April 2014.

[9] Statement by Secretary Jeh Johnson on the designation of election infrastructure as a critical infrastructure subsector. URL: https://www.dhs.gov/news/2017/01/06/statement-secretary-johnson-designation-election-infrastructure-critical, Jan 2017.

[10] Commercial voter files and the study of U.S. politics. URL: https://www.pewresearch.org/methods/2018/02/15/commercial-voter-files-and-the-study-of-u-s-politics/, Feb 2018.

[11] Russian active measures campaigns and interference in the 2016 u.s. election. URL: https://www.intelligence.senate.gov/sites/default/files/documents/Report_Volume1.pdf, 2019.

[12] From voter registration to mail-in ballots, how do countries around the world run their elections? URL: https://www.pewresearch.org/fact-tank/2020/10/30/from-voter-registration-to-mail-in-ballots-how-do-countries-around-the-world-run-their-elections/, 2020.

[13] Block the vote: How politicians are trying to block voters from the ballot box. URL: https://www.aclu.org/news/civil-liberties/block-the-vote-voter-suppression-in-2020, August 2021.

[14] Colorado Revised Statutes 2021 Title 1. URL: https://www.sos.state.co.us/pubs/info_center/laws/Title1/Title1.pdf, 2021.

[15] Código Electoral de Panamá. URL: https://www.tribunal-electoral.gob.pa/publicaciones/codigo-electoral/, 2022.

[16] Ben Laurie Adam Eijdenberg and Al Cutter. Verifiable data structures. URL: https://github.com/google/trillian/blob/master/docs/papers/VerifiableDataStructures.pdf.

[17] Ben Adida. Helios: Web-based open-audit voting. In Paul C. van Oorschot, editor, *Proceedings of the 17th USENIX Security Symposium, July 28-August 1, 2008, San Jose, CA, USA*, pages 335–348. USENIX Association, 2008.

[18] Association for Computing Machinery. Statewide databases of registered voters: Study of accuracy, privacy, usability, security, and reliability issues commissioned by the U.S. public policy committee of the association for computing machinery, 2006.

[19] Belfer Center for Science and International Affairs at Harvard Kennedy School. The state and local election cybersecurity playbook. URL: https://www.belfercenter.org/sites/default/files/files/publication/StateLocalPlaybook%201.1.pdf, 2018.

[20] Susan Bell, Josh Benaloh, Michael D. Byrne, Dana Debeauvoir, Bryce Eakin, Philip Kortum, Neal McBurnett, Olivier Pereira, Philip B. Stark, Dan S. Wallach, Gail Fisher, Julian Montoya, Michelle Parker, and Michael Winn. STAR-Vote: A secure, transparent, auditable, and reliable voting system. In *2013 Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (EVT/WOTE 13)*, Washington, D.C., August 2013. USENIX Association.

[21] Josh Benaloh, Douglas W. Jones, Eric Lazarus, Mark Lindeman, and Philip B. Stark. SOBA: secrecy-preserving observable ballot-level audit. In Hovav Shacham and Vanessa Teague, editors, *2011 Electronic Voting Technology Workshop / Workshop on Trustworthy Elections, EVT/WOTE '11, San Francisco, CA, USA, August 8-9, 2011*. USENIX Association, 2011.

[22] 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.

[23] Jonathan Brater, Kevin Morris, Myrna Pérez, and Christopher Deluzio. Purges: A growing threat to the right to vote. URL: https://www.brennancenter.org/sites/default/files/2019-08/Report_Purges_Growing_Threat.pdf, July 2018.

[24] Jian Cao, Seo young Silvia Kim, and R. Michael Alvarez. Bayesian analysis of state voter registration database integrity. *Statistics, Politics and Policy*, 13(1):19–40, 2022.

[25] Carter Casey, Johann Thairu, Susie Heilman, Susan Prince, Brett Pleasant, and Marc Schneider. Recommended security controls for voter registration systems. URL: `https://www.mitre.org/sites/default/files/publications/pr-19-3594-recommended-security-controls-for-voter-registration-systems.pdf`, December 2019. Report, MITRE Corporation.

[26] Center for Election Innovation and Research. Voter registration database security. URL: `https://electioninnovation.org/wp-content/uploads/2020/08/2020_VRDB_Security_Report.pdf`, Aug 2020.

[27] Center for Internet Security. A handbook for elections infrastructure security. URL: `https://www.cisecurity.org/wp-content/uploads/2018/02/CIS-Elections-eBook-15-Feb.pdf`, Feb 2018.

[28] Colorado Department of State. Statewide voter registration system requirements. URL: `https://www.sos.state.co.us/pubs/elections/SCORE/files/systemrequirementsv10.doc`, 2006.

[29] Colorado Department of State. Election Rules [8 ccr 1505-1]. URL: `https://www.sos.state.co.us/pubs/rule_making/CurrentRules/8CCR1505-1/Rule2.pdf`, 2021.

[30] Scott A. Crosby and Dan S. Wallach. Efficient data structures for tamper-evident logging. In *Proceedings of the 18th Conference on USENIX Security Symposium*, SSYM'09, page 317–334, USA, 2009. USENIX Association.

[31] Election Assistance Commission. Voluntary guidance on implementation of statewide voter registration lists. URL: `https://www.eac.gov/sites/default/files/eac_assets/1/1/Implementing%20Statewide%20Voter%20Registration%20Lists.pdf`, July 2005.

[32] Election Assistance Commission. Statewide voter registration systems. URL: `https://www.eac.gov/statewide-voter-registration-systems`, Aug 2017.

[33] Electoral Knowledge Network. Voter registration.

[34] Gayathri Garimella, Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. Oblivious key-value stores and amplification for private set intersection. Springer-Verlag, 2021.

[35] Sharad Goel, Marc Meredith, Michael Morse, David Rothschild, and Houshmand Shirani-Mehr. One person, one vote: Estimating the prevalence of double voting in U.S. presidential elections. *American Political Science Review*, 114(2):456–469, 2020.

[36] Google. Key transparency. URL: https://github.com/google/keytransparency.

[37] GOV.UK. Register to vote. URL: https://www.gov.uk/register-to-vote.

[38] Douglas W. Jones and Barbara Simons. *Broken Ballots*. Center for the Study of Language and Information, 2012. pg. 243.

[39] Donald C Latham. Department of Defense trusted computer system evaluation criteria. *Department of Defense*, 1986.

[40] Ben Laurie. Certificate transparency. *Communications of the ACM*, 57(10):40–46, 2014.

[41] Law Commission of England and Wales and Scottish Law Commission. Electoral law: A joint final report, 2020.

[42] Le ministre de l'intérieur. Instruction relative à la tenue des listes électorales complémentaires. URL: https://www.eure.gouv.fr/content/download/29070/193726/file/circulaire%20minist%C3%A9rielle%20du%2021%20novembre%202018.pdf, 2018.

[43] Justin Levitt, Wendy R. Weiser, and Ana Muñoz. Making the list: Database matching and verification processes for voter registration. URL: https://www.brennancenter.org/our-work/research-reports/making-list-database-matching-and-verification-processes-voter, 3 2006. Report, Brennan Center for Justice at NYU Law.

[44] Mark Lindeman and Philip B. Stark. A gentle introduction to risk-limiting audits. *IEEE Secur. Priv.*, 10(5):42–49, 2012.

[45] Mark Lindeman, Philip B. Stark, and Vincent S. Yates. BRAVO: ballot-polling risk-limiting audits to verify outcomes. In J. Alex Halderman and Olivier Pereira, editors, *2012 Electronic Voting Technology Workshop / Workshop on Trustworthy Elections, EVT/WOTE '12, Bellevue, WA, USA, August 6-7, 2012*. USENIX Association, 2012.

[46] Maine Department of the Secretary of State. Voter registration data, election data and online forms. URL: https://www.maine.gov/sos/cec/elec/data/index.html.

[47] Holly Maluk, Myrna Pérez, and Lucy Zhou. Voter registration in a digital age: 2015 update. URL: https://www.brennancenter.org/our-work/research-reports/voter-registration-digital-age-2015-update, 2015. Report, Brennan Center for Justice at NYU Law.

[48] Ivan Martin. It firm C-Planet fined €65,000 over massive voter data breach. URL: https://timesofmalta.com/articles/view/it-firm-c-planet-fined-65000-over-massive-voter-data-breach.92848, Jan 2022.

[49] Marcela S. Melara, Aaron Blankstein, Joseph Bonneau, Edward W. Felten, and Michael J. Freedman. CONIKS: Bringing key transparency to end users. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 383–398, Washington, D.C., August 2015. USENIX Association.

[50] Louis-Henri Merino, Simone Colombo, Jeff Allen, Vero Estrada-Galiñanes, and Bryan Ford. TRIP: trustless coercion-resistant in-person voter registration. *CoRR*, abs/2202.06692, 2022.

[51] Lorraine Minnite. Election day registration: A study of voter fraud allegations and findings on voter roll security, 2007.

[52] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review*, page 21260, 2008.

[53] National Academies of Sciences, Engineering, and Medicine. Securing the vote: Protecting American democracy, 2018. Consensus Report.

[54] National Conference of State Legislatures. Access to and use of voter registration lists. URL: `https://www.ncsl.org/research/elections-and-campaigns/access-to-and-use-of-voter-registration-lists.aspx`, Aug 2019.

[55] National Conference of State Legislatures. Same day voter registration. URL: `https://www.ncsl.org/research/elections-and-campaigns/same-day-registration.aspx`, Sept 2021.

[56] National Conference of State Legislatures. Voter registration list maintenance. URL: `https://www.ncsl.org/research/elections-and-campaigns/voter-list-accuracy.aspx`, Oct 2021.

[57] National Conference of State Legislatures. Automatic voter registration. URL: `https://www.ncsl.org/research/elections-and-campaigns/automatic-voter-registration.aspx`, Jan 2022.

[58] National Research Council of the National Academies of Sciences, Engineering, and Medicine. State voter registration databases: Immediate actions and future improvements. URL: `https://www.eac.gov/sites/default/files/document_library/files/State_Voter_Registration_Databases_-_Interim_Report.pdf`, 2008. Interim Report.

[59] North Carolina State Board of Elections. Voter registration data. URL: `https://www.ncsbe.gov/results-data/voter-registration-data`.

[60] Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *2014 USENIX Annual Technical Conference (Usenix ATC 14)*, pages 305–319, 2014.

[61] Sunoo Park, Michael A. Specter, Neha Narula, and Ronald L. Rivest. Going from bad to worse: from internet voting to blockchain voting. *Journal of Cybersecurity*, 7(1), 2021.

[62] Christopher Ponoroff. Voter registration in a digital age. URL: `https://www.brennancenter.org/sites/default/files/2019-08/Report_Voter-Registration-Digital-Age.pdf`, 2010. Report, Brennan Center for Justice at NYU Law.

[63] Peter Rindal and Mike Rosulek. Malicious-secure private set intersection via dual execution. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 2017.

[64] Ronald L. Rivest. On the notion of 'software independence' in voting systems. *Philosophical Transactions of the Royal Society*, 366:3759–3767, 2008.

[65] Jennifer S. Rosenberg and Margaret Chen. Expanding democracy: Voter registration around the world. URL: `https://www.brennancenter.org/sites/default/files/legacy/publications/Expanding.Democracy.pdf`, 2009. Report, Brennan Center for Justice at NYU Law.

[66] Service-Public.fr. Listes électorales : nouvelle inscription. URL: `https://www.service-public.fr/particuliers/vosdroits/F136`, 2022.

[67] Barbara Simons. Why internet voting is dangerous. *Georgetown Law Technology Review*, 4:543–563, 2020.

[68] James Temperton. The Philippines election hack is 'freaking huge'. URL: `https://www.wired.co.uk/article/philippines-data-breach-fingerprint-data`, April 2016.

[69] Alin Tomescu, Vivek Bhupatiraju, Dimitrios Papadopoulos, Charalampos Papamanthou, Nikos Triandopoulos, and Srinivas Devadas. Transparency logs via append-only authenticated dictionaries. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, CCS '19, page 1299–1316, New York, NY, USA, 2019. Association for Computing Machinery.

[70] Sam van der Staak and Peter Wolf. Cybersecurity in elections: Models of interagency collaboration. URL: `https://www.idea.int/sites/default/files/publications/cybersecurity-in-elections-models-of-interagency-collaboration.pdf`, 2019. International Institute for Democracy and Electoral Assistance.

[71] Dustin Volz. Iran probed state election websites since September, U.S. says. URL: `https://www.wsj.com/articles/iran-probed-state-election-websites-since-september-u-s-says-11604104649`, 2020.

[72] Wendy Weiser, Michael Waldman, and Renée Paradis. Voter registration modernization: Policy summary. URL: https://www.brennancenter.org/sites/default/files/2019-08/Report_Voter-Registration-Modernization.pdf, 2009. Report, Brennan Center for Justice at NYU Law.