Desirable Properties of Voting Systems

Svetlana Z. Lowry[1] and Poorvi L. Vora[2]

September 25, 2009

Abstract:

This paper provides definitions for some desirable properties of voting systems, including auditability, ballot secrecy, incoercibility, usability and accessibility. In the context of these desirable properties, it defines the class of end-to-end independently verifiable (E2E) voting systems that provide high levels of auditability, without requiring voters or election observers to trust polling machines or election officials. It provides examples of E2E systems, and compares their auditability properties to those of other voting system classes. Finally, it presents areas for further research in auditable voting systems.

Keywords: voting, end-to-end, auditability, verifiability, incoercibility, ballot secrecy, usability, accessibility, trust model

1. Introduction

The last few years have witnessed the emergence of end-to-end (E2E) voting systems, which enable voter-verification of election outcome. In a paradigm shift from that of the current election technology environment, the systems provide strong outcome-verifiability guarantees that do not require the voter, or the election observer, to trust the voting machine to count votes correctly or polling officials to maintain a secure chain of custody of ballots, or to count votes without error. Several proposed systems have been prototyped; some have been used in binding elections. As such, these systems demonstrate considerable promise. This white paper attempts to understand their strengths and weaknesses. In order to do so, however, it is important to first define the desirable properties of any type of voting system. Given the context of these definitions, one may then define the class of E2E voting systems, and then study their properties. It is hoped that this white paper will provide the background, motivation and framework for the upcoming NIST workshop on E2E Voting Systems.

E2E systems provide certain kinds of strong security guarantees; hence this paper begins with definitions of security properties, which should enable both a careful study of E2E systems and comparisons with existing systems. Additionally, it is very clear that the usability and accessibility properties of a voting system greatly influence its ability to capture voter intent. This, in turn, affects the integrity of the tally determined by the voting system. Hence this white paper addresses usability and accessibility and their incorporation into the (security) trust models implicit in security definitions. The white paper ends with a description of the current state of E2E voting systems, and briefly describes areas of future research.

The white paper is organized as follows. Section 2 presents definitions of security–auditability, ballot secrecy, incoercibility–as well as of usability and accessibility. Section 3 defines end-to-end systems. Section 4 addresses usability and accessibility issues in some more detail. Section 5 provides examples of trust models and common classes of voting systems, to illustrate the definitions of sections 2 and 3. Section 6 describes a general model of most E2E voting systems, with examples and properties. Section 7 presents directions for future work that would be of interest to NIST.

[1] Information Access Division, Information Technology Laboratory, NIST
[2] Computer Security Division, Information Technology Laboratory, NIST and Department of Computer Science, The George Washington University

2 Definitions

This section proposes simple security definitions: for auditability, privacy, incoercibility and usability. At a later stage, we will probably need to study measures related to these properties.

2.1 Auditability

Definition 1 (Auditable) A voting system is *auditable* if it provides verifiable information – about an election, to voters and the general public – that can be used to determine the correctness of the election outcome.

Note the following:

1. The focus is on the *election outcome*, not on the equipment. We do not expect to be able to determine if a particular voting system will work correctly forever, or for the next few elections. However, we require that the voting system be able to convince, after the election, that the outcome produced, for that particular election, is correct. Note that the outcome may be in several possible forms. For example, it could consist of a single winning candidate, a (possibly ordered) list of winning candidates, the round winners in an instant runoff election, or the vote tally for each candidate.

2. The information provided must be *verifiable*. That is, the system should be able to provide some support that the information it provides is accurate. For example, Voter-Verifiable Paper Audit Trail (VVPAT) rolls provide the possibility of manual audits of DRE records. VVPAT Rolls may be classified as verifiable information about the tally (and hence election outcome) because voters verify that they represent the votes. Note that VVPAT rolls do not provide any information to the general public for the purposes of verifiability, and that the public is required to trust election officials, who verify that the rolls have been securely retained, and correctly recounted.

3. We mention that it should be possible to determine the *correctness* of the outcome. That is, auditability should enable observers to determine both: when the system is cheating, as well as when it is *fraudulently* claimed that the system is cheating. Auditability should not come at the cost of a high rate of false negatives about election outcome correctness.

The VVPAT example above motivates two more definitions.

Definition 1.1 (Voter-Auditable) A voting system is *voter-auditable* if it provides verifiable information – about an election, to voters – that can be used to determine the correctness of the election outcome.

Definition 1.2 (Universally-Auditable) A voting system is *universally-auditable* if it provides verifiable information – about an election, to the public – that can be used to determine the correctness of the election outcome.

Note that VVPAT rolls in particular, and Independent Voter-Verifiable Records (IVVRs) in general, do not provide any information regarding an election that can be verified by the public. These systems are hence not publicly-auditable.

2.2 Other Desirable Properties

The purpose of an election is to determine the collective decision of a group. Clearly, an outcome reflects the collective decision of the group only if individual votes truly represent voter intent. There are two major potential hurdles to the capture of voter intent.

1. **Ballot Secrecy**: In an attempt to provide auditability, the voting system may reveal information on how a voter voted. This fact could influence how a voter votes, preventing the vote from reflecting true voter intent. Thus, while the determination of a verifiable election outcome is the primary goal of a voting system, some form of ballot secrecy has been deemed important to enable voters to reveal their intent in their votes. Ballot secrecy is discussed in the next two sections.

2. **Usability**: If the user interface of a DRE, or the ballot design in paper-based voting systems, is confusing for the average voter, or presents challenges to certain groups of voters, the system will not reliably capture voter intent. Similarly, if the system is difficult for election officials to administer, it will not serve its purpose. Usability and accessibility definitions are presented in section 2.6, and usability and accessibility considerations are addressed in section 4.

In most practical instances of voting systems, the provision of auditability has made the systems either less private or less usable or both. Additionally, the procedures and technology that enable auditability should be designed taking into consideration the model of the typical user and how likely he or she is to be able to follow procedures and use technology to achieve the desired outcome of auditability. For example, a DRE might subtly bias a voter towards making an error in casting the vote; such an error is not likely to be caught by a VVPAT or IVVR. Thus there appears to be significant interaction amongst the various desirable properties in most practical instances of voting systems. In the next few sections, we discuss the properties and their interactions.

2.3 Ballot Secrecy: A General Discussion

In this document, we do not go into further discussion of the merits of ballot secrecy as these subjects lie beyond our purview. We do, however, attempt to define the different types of ballot secrecy that can be provided, and indicate whether one definition is stricter than another. This effort would help us understand the levels and kinds of ballot secrecy possible, as well as the impact of ballot secrecy on auditability. Additionally, we identify the assumptions made about adversaries who attempt to determine information on individual votes. Once such assumptions are identified and categorized, a discussion on the appropriateness of the assumptions may follow.

Cryptographers have considered the problem of secure voting systems for almost thirty years. The literature in the area of cryptographic voting systems and protocols hence contains several definitions on ballot secrecy, most very well thought out and addressing minute differences in the levels of ballot secrecy. Very consciously, we do not simply repeat these definitions. Instead, we initiate a discussion on a definition for the evaluation of practical voting systems, to be used in various types of communities for various types of elections, keeping in mind that vote-buying and voter-coercion might be of greater concern in some communities and countries than in others.

More specifically, we note that the following questions are worth discussing:

1. Should one consider different allowable levels of ballot secrecy in voting systems?

2. How powerful may one assume the adversary to be? For example, can the adversary communicate with the voter while he or she is voting? Provide instructions on a card such as "if the number you get is "3", do …"?

3. How should one define the ballot secrecy/incoercibility requirement for voting systems?

In the next few sections we present three definitions of different types of ballot secrecy. In section 2.4 we focus on ballot secrecy in a general sense, assuming any entity wishing to determine how a voter voted has access to information made publicly available by the voting system. We provide two definitions of different levels of ballot secrecy. In section 2.5, we provide a definition of incoercibility, which is a special type of ballot secrecy: one where an entity wishing to determine how a voter voted can intimidate or bribe a voter to provide verifiable information on his individual vote if such is provided by the voting system.

2.4 Ballot Secrecy: Definitions

In this section, we present two definitions of ballot secrecy, assuming the adversary has access to information made publicly available by the voting system. The first definition, Definition 2, is motivated by [10].

Definition 2 (Ballot secrecy, Informational) A voting system is *private* if it (and the procedures/process for using it) does not make available additional information on an individual voter's ballot choice(s).

We note the following:

1. By *additional information* is meant information provided by the voting system, beyond that contained in the tally, and beyond that contained in other external information available about the voter, individually or as part of a demographic.

2. The ballot secrecy requirement addresses information leakage additional to that contained in the *tally*. On the other hand, the auditability requirement is that the observer be convinced the *election outcome* is correct. The outcome can be determined from the tally, but the tally contains more information than does the outcome; the ballot secrecy requirement thus allows the leakage of additional information beyond that contained in the election outcome. This is intentional, for

two reasons. First, voters are used to knowing the tally, so we allow the leakage of tally information, even though it is the outcome that is being determined by the election, and audited. Secondly, the more relaxed requirement was consciously chosen for each definition. For the auditability definition, it is easier to prove the outcome is correct than to prove the tally is correct. Similarly, for the ballot secrecy definition, greater information about a single vote is allowed to be revealed, through the tally, than would have been revealed through the outcome.

3. Finally, note that the term *additional information* in the definition refers to both intentional information leakage for the purposes of auditability (such as in the example of voter names being published alongside votes), as well as unintentional information leakage, such as through covert channels.

4. Definitions in the literature on cryptographic protocols address not only the information made publicly available by the voting system, but also *information that could be coerced out of the voter*. It is perhaps important at this stage to determine what a general definition would be, and then to naturally extend the definition to the case of a coercive adversary.

5. Definitions in the literature on cryptographic protocols address not only information-theoretic leakage, but *whether this information can be used*. For example, there is information-theoretic leakage when the asymmetric-key encryption of a message is made available publicly; however, unless the cryptographic technique is broken, that information is not usable. In this document, the assumption of secure encryption as a means of ballot secrecy protection is assumed, unless otherwise specified. However, we call out secure encryption as an assumption because voting protocols do exist (Moran and Naor's everlasting privacy protocol [16]) that can provide ballot secrecy without assuming the security of an asymmetric-key cryptographic technique.

6. Definitions in the literature distinguish between whether the information leaked is focused on a few votes or is spread out evenly over all votes. For example, one might be able to determine accurately the votes of ten voters and have only a 50% chance of guessing correctly the votes of the 990 other voters in a thousand-voter two-candidate election. In this case, the information is focused on a few voters. On the other hand, one might be able to improve the chance of guessing correctly all votes to 50.5%. In both cases, the expected number of correctly guessed votes is 505, however the first case appears to violate ballot secrecy more strongly. The first case is referred to as a privacy breach in the data mining literature.

7. In most voting system designs, certain entities are trusted with information that could match votes with the order in which they were cast.

8. Just as it is not possible to prove the non-existence of bugs in software, it is *not possible to prove that there is no information leakage at all in an implemented voting system.* Additionally, tampered voting machines would typically reveal information about votes and the order in which they were cast. In this sense, the ballot secrecy property is very different from the accuracy/tally correctness property, which can be proven for a particular election. However, it might be possible to prove that there is no "intentional" information leakage; that is, that the declared protocol is privacy-preserving (for example, that the tally verification audit is computational zero-knowledge). Such a proof would require a model of the adversary: is the adversary computationally unbounded? Is there a channel between the adversary and the voter (the scratch-off card used in the Kelsey-Moran-Regenscheid attacks [15]). Is the voter colluding with the adversary? (voluntary vs. involuntary privacy, coercion resistance, receipt-freeness, etc.)

Another possibility is Definition 3, which is not as stringent as Definition 2, but provides some measure of plausible deniability to the voter.

Definition 3 (Ballot secrecy, Deniability) A voting system is *private* if, given all the additional information provided by it (and the procedures/process for using it), there are at least two ballot choices (of reasonable probability) associated with each voter.

Note the following:

1. The definition might not allow deniability on the nature of voter's choices (when two candidates have some similarity in profile). Consider, for example, two candidates supporting a park at a particular geographical location. If the system information implies that the voter voted for one of the two, it has provided the fact that the voter voted for a candidate supporting the park.

2. The definition does not distinguish between information (on the nature of voter choices) obtained from sources external to the voting system (such as party membership lists), and that obtained from the voting system itself, which is what should be evaluated.

2.5 Incoercibility

In situations where it is possible and/or likely that a coercive adversary will try to bribe or threaten a voter to vote in a certain manner, it is important that the voting system satisfy a more stringent definition of ballot secrecy; we provide an example which is an example extension of Definition 2. This definition includes information that a voter would provide; that is, the voter would be convinced to cooperate with the adversary and to provide information not provided by the voting systems. In systems that provide no information other than the tally, the voter's cooperation would not mean much, as the voter could say what he thinks the adversary would like to hear. However, in the case of auditable systems that do provide verifiable information, it might be difficult for a voter to lie about how he or she voted, as the adversary might be able to check this against the verifiable information provided.

Definition 4 (Incoercible) A voting system is *incoercible* if additional information provided by the voting system (and the procedures/process for using it), combined with any verifiable information provided by the voter, does not improve an adversary's guess on how the voter voted.

The above definition provides a starting point for discussions on what should be captured in the definition of incoercibility. It can be made more specific in the manner of the definition of receipt-freeness due to Benaloh and Tuinstra [4] (the first to capture this concept), or that of incoercibility due to Canetti and Genarro (who define it for a general multiparty computation) [5]. Juels, Catalano and Jakobsson [14] extend the notion of incoercibility to take into consideration coerced abstention; Moran and Naor [16] extend it to allow the coercer to query the voter during the protocol.

In determining a definition, one may want to evaluate the risk of coercion, and the cost of a coercive or vote-buying attack. For example, a voting system where an adversary may actually successfully buy or coerce single votes, but one where the effort to do so is not easily replicated across votes and hence does not scale well for an election may bear a large cost for a coercive attack. Also related is the question of coerced abstention, or coerced vote randomization, which are not as effective in changing vote totals as the coercion of a voter to vote for a certain candidate.

2.6 Usability

In this section we address the desirable property of usability. We first address the general notion of usability, and then attempt to provide a definition for the more specific notion of the usability of a voting system

2.6.1 The General Notion of Usability

Usability is a quality attribute that assesses how easy user interfaces are to use. The word "usability" also refers to methods for improving ease-of-use during the design process. Usability is defined by five quality components:

•Learnability: How easy is it for users to accomplish basic tasks the first time they encounter the design?

•Efficiency: Once users have learned the design, how quickly can they perform tasks?

•Memorability: When users return to the design after a period of not using it, how easily can they reestablish proficiency?

•Errors: How many errors do users make, how severe are these errors, and how easily can they recover from the errors?

•Satisfaction: How pleasant is it to use the design?

There are many other important quality attributes. A key one is utility, which refers to the design's functionality: Does it do what users need? Usability and utility are equally important: It matters little that something is easy if it's not what you want. It's also no good if the system can hypothetically do what you want, but you can't make it happen because the user interface is too difficult. To study a design's utility, you can use the same user research methods that improve usability, see, for example, [18]. In the context of voting systems, perhaps auditability and ballot secrecy may be viewed as the utility properties of the voting system.

2.6.2 Voting System Usability

A voting system is usable if it meets the performance measures defined in the TGDC-Recommended VVSG [26]. Because there are many attributes that make a system usable, the TGDC defined both design requirements (e.g., minimum font sizes) and performance requirements for voting systems. The performance requirements set a baseline for measuring if voters successfully cast ballots (according to a script).

The TGDC-Recommended VVSG defines usability as a measure of the effectiveness, efficiency, and satisfaction achieved by a specified set of users with a given product in the performance of specified tasks. In the context of voting, the primary user is the voter (although the equipment is used by poll workers as well), the product is the voting system, and the primary task is the correct recording of the votes (although other tasks are associated with poll workers as users, e.g. system setup).

To objectively assess voter performance, high-level performance-based requirements were developed and tested. This include specific metrics for effectiveness (e.g., correct capture of voter selections), efficiency (e.g., time taken to vote), and satisfaction. The voting system is tested by having groups of people (representing voters) attempt to perform various typical voting tasks. The requirement is met only if those tasks are accomplished with a specified degree of success. The metrics are defined as follows:

Total Completion Score – the proportion of users who successfully cast a ballot (whether or not the ballot contains erroneous votes).

Perfect Ballot Index – the ratio of the number of cast ballots containing no erroneous votes to the number of cast ballots containing one or more errors (either a vote for an unintended choice, or a missing vote).

Voter Inclusion Index – a measure of both voting accuracy and consistency. It is based on mean accuracy and the associated standard deviation. Accuracy per voter depends on how many "voting opportunities" within each ballot are performed correctly. A low value for the standard deviation of these individual accuracy scores indicates higher consistency of performance across voters.

Preliminary research at the direction of the TGDC that included experimentation with a variety of voting systems has established following benchmark values that would allow better systems to pass the test, while preventing certification of weaker systems:

Total Completion Score: 98%

Perfect Ballot Index: 2.33

Voter Inclusion Index: 0.35

Thus one may, for example, propose the following definition of a *voter-usable* voting system (thus distinguishing the usability of a system by voters from usability by poll workers and election officials).

Definition 5 (Voter-Usable) A voting system is voter-usable if its total completion score is at least 98%, its perfect ballot index at least 2.33, and its voter inclusion index at least 0.35 computed based on VPP (Voter Performance Protocol) data.

Note the distinction between Definition 5 and Definitions 1-4. Definition 5 is based on measurements of a specific instance of a voting system, and not on voting system designs or algorithms. Thus, in trying to agree on a definition, the important issues would be what quantities should be measured, and what the ideal measurements should be. (Hence, of course, the actual numbers in the above definition would typically be debated). Note also that specific auditable system designs, such as end-to-end voting systems defined in section 3, may require the application of additional or different usability metrics and values, due to additional voter tasks.

2.6.3 Interplay between Usability and Auditability
We note that there exists an interesting interplay among the desirable properties of voting systems. In fact, poor usability of a system can weaken its auditability properties as well. Consider the following examples (it needs to be stated that the voting machines referred to below have not been certified yet) about systems with poor poll worker usability. Similar examples exist for voter-usability, especially with respect to the auditability components of end-to-end independently-verifiable systems (defined in section 3).

Most security problems that derive from poll worker usability issues have to do with the manner in which standard procedures, such as replacing the rolls of paper for paper audit trails, are carried out. National Public Radio reported that in Cuyahoga County, Ohio in at least one case in the 2006 mid-term election, a thermal paper roll had been installed backward, so nothing printed out onto it. In other locations, there were reports of paper jamming so that votes printed over one another [11]

The Washington Post reported in 2008 that data cartridges that store votes were unreadable at one precinct in Washington, DC. The voting system manufacturer suggested two possible causes: static discharge or election workers mishandling the cartridges (which would be resulted from usability issues). These situations that could lead to unintended security breaches but could also easily be exploited to compromise security [24].

The examples demonstrate that poll workers may break the chain of custody or otherwise compromise security if the usability of voting systems is not well-addressed. For example, exceptional situations that are difficult to train for and are probably not well covered in system documentation may lead to security exposures. Consider the installation of battery back-up units. If these are not properly installed and the power goes out, the security of the systems cannot always be completely assured.

2.7 Accessibility

Like usability, accessibility is measured of a specific instance of a voting system. The 2005 VVSG says the following:

The accessibility of a voting device consists of the measurable characteristics that indicate the degree to which a system is available to, and usable by, individuals with disabilities. The most common disabilities include those associated with vision, hearing and mobility, as well as cognitive disabilities

Based on HAVA 301(A) (3)(a) one may also say:
An accessible voting system provides the same opportunity for access and participation (including privacy and independence) to voters with disabilities as to other voters.

Note that neither statement describes how a system might be made available to individuals. It is generally assumed that voters with disabilities would access a voting system through the use of specialized interfaces – for example, blind voters would use audio interfaces. While this is sufficient for a definition of accessibility for the voting process, it may not be sufficient when one considers auditability. Thus, for example, does the blind voter audit the specialized interface of the voting system (i.e. provided by the voting system provider)? Or does he or she simply trust it to represent him or her accurately to the system? One way to address this problem is to require an independent vendor or organization to provide the specialized interface. Another way is to ensure that a large enough number of sighted voters use the same interface, enabling the detection of error or intentional vote-changing on the part of the device. See the report of the HFP subcommittee of the TGDC [12] for a more detailed examination of this question.

2.8 Trust model

It is clear that Definitions 1–2 require a trust model. For example, when the auditor determines that the election outcome is correct, what does he trust? Does he trust that a paper audit trail has been kept in secure custody? Does he trust the computational system he uses to check cryptographic claims? If the auditor requires an assistive device, does he trust the software of that device? If the voting system is being evaluated for ballot secrecy, may one assume that cryptographic algorithms are secure? It has been shown that, unless one can make some assumptions about the limitations of an adversary, or about the properties of the voting device, ballot secrecy and verifiability cannot be simultaneously guaranteed [13]. The trust model is a formal articulation of the assumptions; a weaker voting system requires more assumptions. One could rank voting systems based on the assumptions. Example assumptions of the trust model are:

- Trusted Polling Officials: polling officials (or certain other privileged/pre-determined individuals such as election observers) follow procedures without error or malfeasance.

- Secure Chain-of-Custody: Ballot boxes and/or election equipment—such as DREs—and data—such as cryptographic keys—are kept securely, handled appropriately, and are not manipulated, destroyed or accessed by unauthorized individuals.

- Verified Software: the voting system (or independent verification device) software is error-free. (See [20, 21] for a detailed discussion of software independence for voting systems).

- Secure Hardware: the voting system (or independent verification device) hardware is secure and tamperproof.

- Secure Cryptographic Algorithms: an adversary trying to compromise the ballot secrecy of the vote or the accuracy of the tally cannot break the cryptographic algorithms used.

- Non-Colluding Participants: There is a certain set of *n* participants trusted not to collude to compromise the ballot secrecy or accuracy of the voting system. For example, consider the situation when a set of *n* participants is given a (distinct) key each, such that all keys are required to determine how an individual voter voted, and fewer than *n* keys reveal no information about the vote. In this case, the vote is private if the participants do not collude and not all cryptographic keys are compromised; that is, if even one of the *n* participants is honest.

- Trusted specialized user interfaces for voters with disabilities.

3 End-to-end Voting Systems

In this section we define end-to-end voting systems, a class of cryptographic voting systems that provides strong auditability and ballot secrecy guarantees with a minimal set of assumptions. In the early (remote) cryptographic voting systems, voters encrypted their own votes on their personal computers. These votes were then tallied in a manner that could be verified by anyone. Such systems provided very strong mathematical guarantees on the correctness of the election outcome, and allowed voters and observers to verify the correctness of an election outcome.

The model of these systems is, however, one of remote voting, which is highly susceptible to coercion attacks (an adversary can watch while the voter votes) as well as vote-casting attacks caused by malware. Because it was not clear how a voter could encrypt a vote using the untrusted voting machine in the polling booth, it was not possible to propose the use of cryptographic voting systems in public elections.

This changed in approximately 2003, with descriptions of voting systems by Chaum [6] (see [25] for more technical detail) and Neff [17] that enabled a voter to encrypt his or her vote without access to trusted computational power in the polling booth. Several more usable voting systems of this kind have followed, see sections 6.1 and 6.2. These voting systems have been described as being *end-to-end independently verifiable*, and possess two key properties:

1. The E2E systems do not require that a specified piece of hardware (such as a mechanical lever machine) or software (the software of a specific voting system) or a specified set of individuals (election officials) be trusted to count votes without error. Any hardware or software used by the voting system may be audited on its performance during the election by voters and observers.

2. To the extent that hardware or software is required to perform the audit (for example, to check the correctness of certain cryptographic operations), such hardware or software may be chosen by the individual performing the audit, and several individuals will typically perform independent audits.

Thus the class of end-to-end voting systems makes the fewest assumptions about specific entities that need to be trusted for system auditability. In particular, it does not require that voting system software be trusted, and is hence *software independent*.

Recall the definition of a software independent voting system [20, 21]:

A voting system is *software independent* if an (undetected) change or error in its software cannot cause an undetectable change or error in an election outcome.

We propose the following definition of E2E systems:

Definition 6 (End-to-end independently verifiable) A voting system is *end-to-end independently verifiable* if an independent, honest observer can determine—with virtual certainty—whether a declared election outcome correctly represents the votes cast by voters. To the extent that the observer is required to trust entities, software or hardware, he or she should be able to choose said entities, software or hardware.

We make note of the following:

1. The *independent observer*—and not the privileged observer, such as the polling official—should be able to determine if the election outcome is correct.

2. The use of the term *virtual certainty* is deliberate. There will typically be some (even if small) uncertainty in the mind of the independent observer regarding the correctness of the outcome. For example, cryptographic end-to-end systems prove that the tally is correct with overwhelming probability, not that the tally is certainly correct. We thus require that the independent observer should be almost certain that his conclusion (regarding tally correctness) is itself correct.

3. The fact that the independent observer determines *correctness of the election outcome* protects both the voting system and the voter. That is, it ensures that both attempts to rig the tally, as well as false charges of election fraud, will be detected with virtual certainty.

4. We require that observers not be forced to trust a certain piece of hardware or software, or a certain entity. However, verifying encryption and performing cryptographic audits would require access to software and hardware that would perform correctly the required computations. Hence we allow the observer to choose such hardware or software.

5. In all published E2E system designs, ballot stuffing is detected through the matching of the number of votes processed by the system with the number of votes cast as recorded by voter registration logs. Thus, in order to audit election correctness, if the process of voter validation is not verifiable, the independent observer will be required to trust the record of the number of votes cast.

6. Note that the definition is about an election outcome caused by *votes cast by voters*. That is, it does not address the election outcome that would have resulted from votes that voters might have intended to cast, (but did not cast, either because they were coerced into casting another vote, or because the ballot presented was so confusing, voters did not communicate the intended choice). Thus this definition does not specifically address the usability of the system, and whether it presented the voter with a confusing ballot. The usability of a voting system is treated as a desirable quality.

Clearly, the outcome of an election depends on the votes cast and, in an election where no one watches the voter fill up her ballot, only a voter can verify that her vote is represented correctly. Hence, correctness of an election outcome, in an E2E voting system, requires that:

1. The recording of the vote is *voter-verifiable*, and

2. The correctness of the processing of collected votes is *universally-verifiable*.

Recall that we defined voter-auditable and universally-auditable voting systems earlier in this paper (see Definitions 1.1 and 1.2). These systems provided verifiable information to voters and the public respectively. However, it was not required that these types of systems provided enough information for voters and the public to determine that the election outcome is correct, simply that they provided *some* verifiable information.

Voter-verifiable vote recording and universally-verifiable vote processing are properties of special cases of voter-auditable and universally-auditable voting systems respectively.

Definition 7.1 (Voter-Verifiable) A process of a voting system is *voter-verifiable* if an honest voter can determine—with virtual certainty—whether the process was correctly carried out. To the extent that the voter is required to trust entities, software or hardware, he or she should be able to choose said entities, software or hardware.

Definition 7.2 (Universally-Verifiable) A process of a voting system is *universally-verifiable* if an honest observer can determine—with virtual certainty—whether the process was correctly carried out. To the extent that the observer is required to trust entities, software or hardware, he or she should be able to choose said entities, software or hardware.

An observer wishing to determine election correctness – given a system that has voter-verifiable vote recording and universally-verifiable vote processing – can check that (a) a large enough number of voters have verified their votes are recorded correctly in the collection, and (b) the processing of collected votes is correct. The observer may then conclude that the election outcome is correct with large enough certainty. The extent of the certainty depends on the number of voters who have verified the recording of their votes, and the certainty that the processing is correct.

The class of E2E systems is compared with other classes in section 5.2, and examples of E2E systems are presented in section 6.

4. A Discussion on Usability and Accessibility

Traditionally, the security of a computer system has been addressed separately from its usability. This practice has been changing recently, (albeit slowly). The usability of a voting system is a particularly important property. If a voting system is not easily used by all voter communities, then it is not an effective recording agent of voter intent, and the integrity of its outcome—as representative of voter intent—is questionable. Further, ballot secrecy and auditability requirements might impose an additional usability burden on the user; this, in turn, might make the system less private or less auditable or less accurate, subverting the original purpose. This is particularly true in the case of voting systems, which are used infrequently by voters, who never really get familiar with the voting systems. For these reasons, we attempt to integrate discussion on the usability of a voting system into our discussions on auditability and ballot secrecy. Because the idea of an integrated discussion is relatively new, our aim is to simply identify problems worthy of further study:

1. Is there a general set of rules for the design of user-friendly voting systems? Such a set of rules would be very useful to the security researcher or to the E2E system designer, even in the form of some rules of thumb.

2. What are specific usability issues with auditable voting systems, and what is their impact on specific security (auditability and ballot secrecy) goals? For example, the usability of E2E system procedures by poll workers can affect the auditability and ballot secrecy properties that rely on the ability of poll workers to follow certain procedures.

3. Conversely, what is the impact of auditability and ballot secrecy requirements on usability goals? Is there a "most usable" system that satisfies auditability and ballot secrecy requirements? Does such a system possess enough usability? That is, can we design a system satisfying all three of our goals?

4. How usable are the existing end-to-end voting systems? In particular, does auditability require the voter to perform additional tasks that could be more complex than simply casting a ballot, and, if so, whether the auditability requirement makes the vote casting process less usable for voters.

5. As newer E2E systems are designed, what should system designers pay attention to? For example, what is reasonable to expect a voter to do? Is it reasonable to expect a voter to be able to remember a 6-digit string? Or to determine if two 6-digit strings are identical? Or to be able to use a ballot with indirection, such as the PunchScan ballot?

6. What are research issues in designing appropriate usability experiments for E2E systems? These experiments should address usability from the point of view of poll workers, the usability of the ballot casting part of the system, as well as the usability of the voter-verifiability part.

7. How do human factors affect user interaction with E2E systems? What physical (e.g., input device manipulation), behavioral (e.g., memory limitations), and demographic characteristics (e.g., age and familiarity with electronic devices, including computers) affect the usability of these systems? How usable are these for all relevant roles – voter, poll worker, auditor, amongst others. What are the main usability problems? Note that usability problems would lead to failure to vote as intended; failure to close voting systems correctly; failure to identify incorrect vote counts, when they are present. Additionally, such problems would result in threats to the auditability or ballot secrecy of the E2E system.

8. Related to the above point, what are the usability-related assumptions in the trust model? For example, do we assume that the ballot design is not intentionally biased to cause voter inaccuracy? Do we assume that voters vote accurately as they intended, that is, without error.

9. What properties of usability affect the utility of E2E voting systems (meaning the degree to which performing a task achieves the intended goal (casting a ballot for the intended candidate)? At what point does task complexity cause voters to fail to vote as intended or worse fail to vote at all and not be aware of it? At what point does a lack of feedback cause users to lose trust in the reliability of the device and/or confidence in their ability to "do the job right"; such that they perform incorrect actions or perform unnecessary actions that produce inappropriate system response? Can E2E systems be designed to tolerate some voter error?

10. Consider a voting system which uses paper ballots. Perhaps it uses an AutoMark device to enable blind voters to mark the ballots. Such a device presents audio input about the ballot to the voter, and enables the voter to mark the ballot. In this instance, how might a blind voter verify that her vote is recorded correctly? This question is particularly relevant when one considers auditability of the voting system – does the blind voter also audit the AutoMark device? How might he do that?

    Following immediately from the above question is the question of how one deals with accessibility in the trust model? Can we come up with a single model for how accessibility is allowed to impact the security trust model? Such a model might be: "special user interfaces are allowed, as long as independent ones are used for independent data streams". For example, a blind voter filling up a paper ballot might use a printer to print on the ballot and a scanner to determine that the printer printed correctly. Is it ok if the printer and scanner come from different vendors? From the same vendor? One source for a starting point is [12], which examines how voters with poor vision might use IVVRs in a software independent system. Note that an IVVR needs to be permanent (and hence not in electronic form, but typically in paper form).

5 Trust Assumptions of Some Voting Systems

In this section, we describe a few voting systems and name their trust assumptions, in order to determine the degree of auditability provided. Table 1 provides a summary of this section. Before we examine trust assumptions, we first describe a general model of an election.

5.1 An Election Model

For simplicity, we assume elections where a voter makes a single choice in each race, though the model trivially extends to other types of elections. The election consists of four loosely-defined stages:

- **Election Set-Up:** Election officials prepare ballots and may make public certain types of information, such as candidates or public keys.

- **Ballot Casting and Recording:** A voter is presented with several races, and several choices for each race. He or she votes for at most one choice in each race, reviews the ballot, and then casts it, at which point it becomes a recorded ballot, and part of the official ballot collection. The voting system may produce an IVVR (Independent Voter-Verifiable Record) which is a copy of the recorded ballot and may be used to audit the election.

- **Ballot Tallying:** The ballots in the ballot collection are tallied by the voting system to produce a tally and an election outcome.

- **Election Audit:** The outcome is audited, using the IVVRs and/or any other verifiable information provided by the voting system in any of the stages above. It may be necessary to audit other aspects of the election while it proceeds, on occasion even before votes are cast. For example, printed ballots may be audited by certain types of voting systems; DRE’s may be audited before use, the registration logs may be balanced with the number of cast and spoiled ballots, etc. Thus, the audit stage may overlap, chronologically, with all other stages. This stage also includes formally-required ballot recounts.

In this general election model note the following assumptions for auditability (there are also assumptions we make in order to obtain ballot secrecy, these are not noted below):

- **Secure Chain of Custody Assumption:** If the IVVRs are required for the audit, and are kept in the custody of the precinct/county/state between the election and the audit, the auditability of the system is based on the trust model assumption that a secure chain-of-custody is maintained for the IVVRs. If there are no IVVRs, a secure chain of custody needs to be maintained for the paper or electronic ballots (for example, for the DRE records).

- **Trusted Poll Worker Assumption:** If the audit involves a manual recount of IVVRs, or a manual count of ballots without IVVRS, the auditability of the system is based on the trust model assumption that the combined effort of officials performing the manual recount results in a good approximation of the true tally, and hence also results in an audit of the correct outcome. In particular, it means that if there is error or malfeasance on the part of those performing the recount, it is small enough to not have an effect on the audit outcome. If the recount is public, the trust model assumption is weaker, because efforts to make a larger difference in the tally are more likely to be detected.

- **Randomness Assumption:** In instances where random choices are required, such as choices of precincts for a manual recount, we assume we have access to randomness, and that these choices cannot be predicted by the adversary attempting to change the election outcome.

- **Usable and Human-Error-Resistant Auditability Assumption:** We assume that a voter, poll worker or election official performing a task related to auditability can perform it correctly, and that the technology used resists human error, to a degree sufficient for audit correctness. For example, we assume that it is easy enough to install VVPAT printers, and to check VVPAT records, that a large enough number of voters will have access to records and will check them, and that if VVPAT records are presented incorrectly to voters, a large enough number of voters will notice the error

We now look at common classes of voting systems

5.2 Example Voting System Classes and Trust Assumptions

1. Paper ballots with optical-scan counts and manual recounts

Voters fill out paper ballots. The ballots are fed into an optical scanner which computes a tally, from which the election outcome may be determined. The ballots (IVVRs) are kept in the custody of the precinct/county/state between the election and the audit, and are manually or machine counted for the audit.

2. Direct recording electronic devices (DREs)

Voters enter their choices into the DRE, which computes a tally, from which election outcome may be determined. This system provides no auditability.

3. DREs with voter-verifiable paper audit trail (VVPAT) or other independent voter-verifiable record (IVVR) (such as audio audit trail), or DREs with audit ports

Voters enter their choices into the DRE, which provides an IVVR. The DRE computes a tally, from which the election outcome may be determined. IVVRs are kept in the custody of the precinct/county/state between the election and the audit, and are manually or machine counted for the audit.

4. End-to-End Voting Systems

In a typical end-to-end voting system, voters cast encrypted votes which are tallied in a universally-verifiable and auditable manner.

Table 1 on the next page compares the systems and their trust assumptions for the auditability property. Note that the E2E voting systems do not require the two major auditability trust model assumptions of the other voting systems. They do not require the secure chain of custody assumption because a voter can check if her receipt is in the publicly-displayed vote collection, and need not trust another entity to maintain a secure chain of custody for it. They do not require that the entities computing the tally or performing the recount be trusted to do so correctly, as anyone may verify tally correctness. Some of the E2E voting systems (Scantegrity, for example) do require, however, that a secure chain of custody be maintained for ballots before the election.

System Type	Auditable	Publicly Auditable	Auditability Requires Trusted Poll Workers	Auditability Requires Secure Chain-of-Custody	Software Dependent
Paper + manual recount	√	×	Yes	Yes	No
DRE	×	×	Not Auditable		Yes
DRE + IVVR	√	×	Yes	Yes	No
E2E	√	√	No	No	No

Table 1: Auditability and Trust Assumptions: more “No”s implies a system with stronger auditability properties.

6 Example End-to-End Voting Systems

In this section we describe the general end-to-end voting system, and also describe two specific classes of proposed and prototyped voting systems that come close to satisfying the definitions we propose. In the typical end-to-end voting system, a vote is encrypted. The encrypted vote itself, or something closely linked to it, such as a hash, forms the receipt the voter can take home with him. The encryption of votes may be audited. Voters may take copies of the receipt out of the polling booth; because two identical encrypted votes do not correspond to identical unencrypted votes, this does not violate ballot secrecy. The receipts are posted in a publicly viewable location, such as on a notice board outside the polling booth or on a public website. Voters—or individuals representing them—can verify that their receipts match those on the website; these systems are hence voter-verifiable. Even a small number of voters checking in this manner results in a high certainty of fraud detection. The publicly-displayed vote collection may be tallied in a verifiable manner; the voting system provides information about the processing of encrypted votes that allows any observer to verify that votes were tallied correctly and has minimal impact on vote ballot secrecy. Note that, in the sense that it can be used for the purpose of voter-verifiability of the vote casting process, the receipt is similar to an IVVR. However, the voter-verifiability provided by the receipt is stronger than that provided by the typical IVVR, because it also enables the voter to determine that all votes were counted correctly.

Most end to end voting systems provide ballot secrecy given the following trust model assumptions:

(a) Secure Cryptography Assumption: It is not possible to determine the relationship between a vote and its encryption without other information;

(b) Non-Colluding Participants Assumption: Those involved in encrypting or decrypting the votes do not collude to determine the relationship between a vote and its encryption;

(c) Secure Chain of Custody Assumption: Those with access to printed (unfilled) ballots do not attempt to use the ballots to determine the relationship between a vote and its encryption.

(d) Procedural Assumption: Voters and polling officials follow a procedure that prevents vote-buying or selling attacks.

One end to end voting system (Moran and Naor [16]) provides ballot secrecy without recourse to the secure cryptography assumption; however, it requires the secure cryptography assumption for its auditability property.

In the next sections we describe two classes of systems where the voter may encrypt her vote without trusting a computer in the polling booth. The first class of system is based on paper ballots, while the second uses electronic ballots.

6.1 E2E Paper Ballot Systems

These systems use specially-designed two-part paper ballots, presented so that, by filling up these ballots, the voter encrypts her vote. The ballots are designed so that one part of a filled-in ballot forms the encrypted ballot, and can be taken out of the polling booth by the voter. This part also forms the cast ballot and the IVVR, and is typically scanned at the polling place. The encrypted ballot also bears information which is used to tally the ballots. Both the vote encryption and the tally computation can be audited. In order to check that the information used for tally computation is correct, ballots can be filled in, then spoiled and audited, to determine what the encrypted vote would decrypt to. If the spoiling is done in a manner that cannot be predicted by the voting system, a voting system that cheats on the information printed on a ballot will be caught by these audits. Further, the processing by the tellers can also be audited with minimal impact on vote ballot secrecy. Note that paper ballots immediately present accessibility issues for those who have difficulties filling up paper ballots independently; some of these problems are addressable with the use of accessibility devices, some are not. In the illustrative examples below, all ballots may be scanned, allowing the voting system to warn the voter of over or undervotes; this feature is not present in traditional paper ballot systems that do not use any automation.

The following are illustrative examples of the paper-based end-to-end systems. All the following examples rely on the security of cryptographic schemes in order to provide vote ballot secrecy. Some also require security of cryptographic schemes for tally accuracy.

1. **Prêt à Voter** [22, 9]: The candidates are listed in a random order on the left half of the ballot. The voter marks her choice on the right half of the ballot, which also bears the serial number of the ballot. The two ballot halves are separated before the right side of the ballot is cast. The position of the voter's mark, on the right ballot half, is her encrypted vote; without the left half, it is not possible for others to determine how the voter voted. Some jurisdictions do not allow the presentation of a random ordering of candidates, which requires voters to pay special attention not required in traditional paper ballots. Further, voters are given scanned copies of the right ballot half, and the requirement that voters check that the copy accurately reflects the original makes the voting process somewhat more complicated than that of filling a paper ballot. Similarly, the random candidate ordering might also be confusing to voters.

2. **PunchScan** [19]: The two-part ballot consists of a top and a bottom half, overlaid while the ballot is being filled (see figure). On the top ballot half is a random association between candidates and a dummy variable for each candidate Candidates are listed in a fixed canonical order, such as alphabetical order. Also on the top half of the ballot is a set of holes through which the voter can see the dummy variables listed on the bottom half, in random order. To vote for a candidate, the voter marks the hole revealing the corresponding

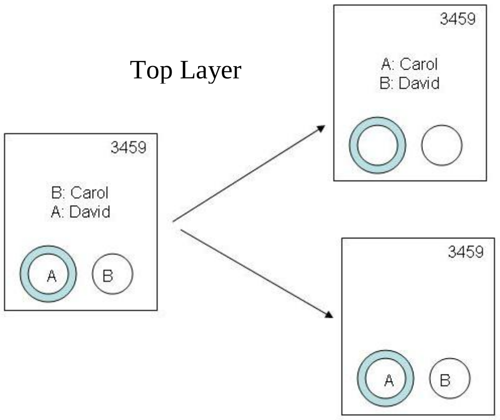

dummy variable with a bingo dauber, so that the mark is made on both ballot halves. In the figure, the vote is for David. The voter chooses one half for the receipt. (For technical reasons, to prevent coercion, the voter must make this choice before seeing the ballot). Each half taken by itself contains no information about the vote. However, each half contains a serial number, which may be used by the system to tally the votes.

This system has been used in two (small) binding elections: University of Ottawa Graduate Student Association (GSAED) election in March 2007 and the election of the Computer Professionals for Social Responsibility (CPSR) in August 2007. It won the Grand prize for the best election system at the university voting systems competition, VoComp 2007. The indirection in its ba Bottom Layer egatively impacted its usability.

3. **Scratch & Vote** [2]: The ballot is similar to the Prêt à Voter or PunchScan ballot, however, the manner in which votes are tallied is different. A detachable chit bears a scratch-off surface that covers a set of numbers. These numbers are needed to determine if the ballot was correctly formed, but are not needed to tally the votes. Because the numbers can be used to determine the vote, the scratch-off surface must be intact for a ballot that is to be cast. The surface would be scratched-off if the ballot were spoiled and audited. The chit is detached and destroyed when the vote is cast. This system inherits any problems of the Prêt à Voter or PunchScan ballots it uses.

4. **Scantegrity** [8, 7]: Ballots are visibly similar to optical scan ballots. By using a special pen to mark the oval of her choice, the voter exposes a code originally written in the oval with invisible ink. The code is the encryption of her vote. The voter may obtain a digitally signed paper receipt bearing this code, or may be required to copy the code onto a piece of paper, which forms the receipt. The receipt also bears a serial number, which is used by the system to tally the votes. This system was used in April 2009 for a (small) mock election by the City of Takoma Park, MD. It will be used in November 2009 for the City Council election, where about 1200 voters are expected to participate. The requirement that voters either write down the code or check the printed receipt to ensure it matches the revealed code, complicate the voting process.

6.2 E2E Electronic Ballot Systems

This section describes Simple Verifiable Voting [3] which uses electronic ballots. It works in a very straightforward fashion. It relies on two distinct types of machines, vote-casting machines and vote-encryption machines, neither of which is trusted to be error-free. Both may be located in polling places. The vote-encryption machine would encrypt a vote and provide the voter a printout of it on paper. Once the vote was encrypted, the voter would be asked if she wished to ready it for casting. If she replied in the affirmative, the encryption would be signed. If not, she would choose to audit the encryption, and the machine would provide her with the information required to determine if the encryption was correct; the voter may later check this information using software of her choice.

Once the voter is in possession of an encrypted vote that she wishes to cast, she authenticates herself at a vote-casting machine, where she casts an electronic ballot digitally signed by the valid vote-encryption machine that encrypted her vote. The paper copies of cast votes can be retained as IVVRs, and the encrypted vote collection can be tallied in a universally-verifiable manner. Examples of voting systems that use this approach include VoteBox [23], and Helios [1], a remote voting system used in March 2009 for a 4000-voter election of the Recteur of the Université Catholique de Louvain, Belgium.

Electronic ballots can be made much more accessible to voters with difficulty marking paper ballots.

6.3 E2E: Paper Ballot Voting Systems vs. Electronic Ballot Voting Systems

All precinct-based E2E systems today use paper for ballots or receipts. Simple Verifiable Voting uses electronic ballots and provides a paper receipt, while Prêt à Voter**,** PunchScan and Scantegrity use paper ballots and provide paper receipts. In general, the more paper involved in an election, the more difficult it is to administer, hence electronic ballot systems are easier to manage. These systems also provide an accessibility-related advantage, because not all voters will be able to handle and mark paper.

On the other hand, the advantage of paper ballots is in the write-once property of paper, which protects both the voting system and the voter. If a scanner or machine misreads a ballot, or decrypts a spoiled ballot incorrectly, this is easily proven by the voter, whose vote is irrefutably recorded on the paper. For the same reason, a voter cannot falsely claim that a spoiled ballot was decrypted incorrectly, or that a scanned ballot was recorded incorrectly.

In the case of electronic ballots, consider the problem of vote-recording. When the voter communicates her vote to the voting machine or the vote encryption machine, she has no proof of what she communicated unless the communication, originating from both voter and machine, is recorded on a write-once medium that cannot be edited or over-written. Consider the following interaction. Voting Machine: Would you like to vote for Bob or Alice? Voter: Bob. Voting Machine: Thank you for your vote for Alice. Similarly, the machine may ask the voter to inspect the receipt and confirm if it is correct, the voter may press the "incorrect" button, and the machine responds "thank you for confirming your vote" and casts the ballot. In these cases, while the voter may know that the machine made an error, she has no way of demonstrating this. If one allows a voter to make a legitimate complaint without proof, so that a large enough number of such complaints can call the election into question, even a small number of dishonest voters can bring down an honest election.

The vote-recording problem is typically solved with the use of a write-once communication medium, such as paper, or, if the voter has access to a trusted encryption device through the use of blind signatures, or, if the voter has access to a printed code-book, through the use of encrypted votes. An interesting electronic approach that attempts to mimic the write-once property of paper is that of VoteBox, which broadcasts a Simple Verifiable Voting receipt to several other machines. However, because these other machines are part of the voting system, it is not clear that they could be trusted.

The paper-based protocols can probably be made electronic with the use of an electronic write-once medium that can be destroyed reliably when the protocol requires, and that the voter can examine independently (with the use of a trusted device or her own senses).

7 Conclusions and Future Research

The subject of electronic E2E voting systems (design, prototyping and usability testing) is perhaps the most important area for future research. While electronic voting systems can be made far more user-friendly than paper-based voting systems, to what extent can electronic voting systems provide the auditability and ballot secrecy properties of the paper-based end-to-end voting systems? Another important research area is that of usability of end to end systems, and, in particular, the interplay between usability and security properties. Finally, the development of definitions, metrics and a taxonomy of voting systems would aid in their standardization.

Acknowledgements

The authors thank Sharon Laskowski for discussions, numerous readings of the paper and very useful feedback. They also thank Barbara Guttman, Nelson Hastings, John Kelsey, Rene Peralta, Stefan Popoveniuc and Andrew Regenscheid for comments and discussions.

References

[1] B. Adida. Helios: Web-based open-audit voting. In Proceedings of the 17th Usenix Security Symposium, June 2008.

[2] B. Adida and R. L. Rivest. Scratch & Vote: self-contained paper-based cryptographic voting. In WPES ’06: Proceedings of the 5th ACM Workshop on Privacy in the Electronic Society, pages 29–40, 2006.

[3] J. Benaloh. Simple verifable elections. In EVT’06: Proceedings of the USENIX/Accurate Electronic Voting Technology Workshop, 2006.

[4] J. Benaloh and D. Tuinstra. Receipt-free secret-ballot elections (extended abstract). In STOC ’94: Proceedings of the 26th ACM Symposium on Theory of Computing, pages 544–553, 1994.

[5] R. Canetti and R. Gennaro. Incoercible multiparty computation. In FOCS 96: Proceedings, 37th Annual Symposium on Foundations of Computer Science, pages 504–513, Oct 1996.

[6] D. Chaum. Secret-ballot receipts: True voter-verifiable elections. IEEE Security and Privacy, 2(1):38–47, January/February 2004.

[7] D. Chaum, R. Carback, J. Clark, A. Essex, S. Popoveniuc, R. L. Rivest, P. Y. A. Ryan, E. Shen, and A. T. Sherman. Scantegrity ii: End-to-end verifiability for optical scan election systems using invisible ink con.rmation codes. In EVT’07: Proceedings of the USENIX/Accurate Elec-tronic Voting Technology Workshop, 2008.

[8] D. Chaum, A. Essex, R. Carback, J. Clark, S. Popoveniuc, A. T. Sherman, and P. Vora. Scantegrity: End-to-end voter verifiable optical-scan voting. IEEE Security and Privacy, 6(3):40–46, May/June 2008.

[9] D. Chaum, P. Y. A. Ryan, and S. Schneider. A practical voter-verifiable election scheme. In ESORICS 2005: Proceedings of the 10th European Symposium On Research In Computer Security, pages 118–139. Springer, 2005.

[10] L. Coney, J. L. Hall, P. L. Vora, and D. Wagner. Towards a privacy measurement criterion for voting systems. In National Conference on Digital Government Research, May 2005.

[11] Fessler, Pamela. 2006. All Things Considered, National Public Radio, Problems Found in Ohio Computer Voting.

[12] HFP subcommittee of the TGDC. Four approaches to SI and accessibility. http://vote.nist.gov/meeting-03222007/SI-n-access-031207.pdf, March 2007.

[13] B. Hosp and P. L. Vora. An information-theoretic model of voting systems. Mathematical and Computer Modelling. Special issue on: Mathematical Modeling of Voting Systems and Elections: Theory and Applications, 48(9-10):1628–1645, Nov. 2008.

[14] A. Juels, D. Catalano, and M. Jakobsson. Coercion-resistant electronic elections. In WPES ’05: Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society, pages 61–70, 2005.

[15] John Kelsey, Tal Moran and Andrew Regenscheid. Attacking Paper-Based E2E Voting Systems. *Best of WOTE.* To appear.

[16] T. Moran and M. Naor. Receipt-free universally-verifiable voting with everlasting privacy. In CRYPTO 2006, pages 373–392, September 2006.

[17] C. A. Neff. Practical high certainty intent veri.cation for encrypted votes. http://www.votehere.net/old/vhti/documentation/vsv-2.0.3638.pdf, October 2004.

[18] Jakob Nielson, http://www.useit.com/alertbox/20030825.html

[19] S. Popoveniuc and B. Hosp. An introduction to PunchScan. In WOTE 2006: IAVoSS Workshop On Trustworthy Elections, June 2006.

[20] R. L. Rivest. On the notion of “software independence” in voting systems. Phil. Trans. Royal Society A, 366(1881):3759–3767, 2008.

[21] R. L. Rivest and J. P. Wack. On the notion of “software independence” in voting systems. http://vote.nist.gov/SI-in-voting.pdf, 2006.

[22] P. Y. A. Ryan. A variant of the Chaum voter-verifiable scheme. Technical Report CS-TR-864, University of Newcastle upon Tyne, School of Computing Science, October 2004.

[23] D. Sandler, K. Derr, and D. S. Wallach. Votebox: a tamper-evident, verifiable electronic voting system. In Proceedings of the 17th USENIX Security Symposium, pages 349–364, 2008.

[24] Stewart, Nikita. 2008. Voter Database Is Fine, Firm Says: User Cited As Possibility in D.C. Vote Foul-Up, The Washington Post.

[25] P. L. Vora. David Chaum’s Voter Verification using Encrypted Paper Receipts. http://dimacs.rutgers.edu/Workshops/Voting/slides/vora.pdf, also available as IACR eprint archive, no. 2005/050, 2004.

[26] Voluntary Voting System Guidelines Recommendations to the Election Assistance Commission. Technical Guidelines Development Committee. August 2007. http://www.eac.gov/files/vvsg/Final-TGDC-VVSG-08312007.pdf