

Proposal Summary

The voting system integrity problem is a paradigmatic hard Cyber Security problem, spanning the entire Cyber Trust program including trustworthy system architectures, security, integrity, privacy, anonymity, high assurance, and man-machine interfaces. Voting systems are an excellent example of the class of systems where any weak link may result in undetected accidents or enable malicious tampering.

Without exaggeration, voting systems are one of the pillars of our democracy. Voting systems allow the electorate to determine the course taken by our nation. As a result, voting systems face a wide variety of requirements and constraints. Voting systems must be secure against tampering, yet they must be easy to use for all voters. They must satisfy a variety of state and national standards, yet they must be affordable to purchase and maintain. They must help voters to correctly indicate their voting intent, even when the voter intends not to cast a vote! They must preserve a voter's privacy and anonymity, to reduce risks of voter coercion and bribery, yet they must be sufficiently auditable and transparent to allow for mistakes and errors to be identified and reconciled. They must be robust against corruption and malice among system developers and the officials who run the election, yet the systems must be safe enough to leave unattended overnight in a school cafeteria.

Engineering voting systems to satisfy these often contradictory constraints is difficult, requiring research into the full gamut of the problem, from the software and hardware design through the careful consideration of legal and administrative procedures. Human factors issues must be considered to make voting systems accessible to all eligible voters, regardless of disability. Likewise, the system must be comprehensible to poll workers and transparent to election observers. Ultimately, the election system is responsible not for naming the winner of a race, but for convincing the loser that he or she, indeed, lost the election. We will investigate software architectures, tamper-resistant hardware, and cryptographic protocols. We will look at the role paper should play in electronic voting systems. We will examine system usability and study how public policy and administrative procedure can better safeguard the system. Only by considering all possible aspects of these systems can we have any assurance, at the end of the day, that our elections will be fair and that the will of the electorate will be correctly reported.

Intellectual Merit To tackle the voting problem, the proposed research must answer many deep and difficult questions that are of great interest to a number of other types of systems. The most basic question is: How can we responsibly employ computer systems for tasks that require high levels of trustworthiness, when we know that those systems will not be totally reliable, bug-free, or totally secure, particularly when every human participant from the system designers to the end users is a potential adversary and when human errors are commonplace? Solving this problem requires thinking about the end-to-end behavior of a whole system, including software, hardware, procedures, law, and people. Perhaps most important, the research problem requires people from different areas of computer science, law, and human factors to combine their efforts in new and innovative ways.

Broader Impacts This proposal is motivated by a need to achieve greater integrity in our elections. This is a problem of burning public interest that has consumed an increasing amount of time for most of the PIs on this proposal — who have already been involved in the public dialog, whether through our existing research studies of voting systems or in our testimony and participation in government hearings and standards bodies. Furthermore, we have already integrated voting into many of our graduate and undergraduate courses, developing materials that other academics have begun to adopt. The proposed Center for Correct, Usable, Reliable, Auditable, and Transparent Elections (ACCURATE) will host annual public workshops for a broad spectrum of participants. Likewise, all of our materials, including videos of the workshops, will be made available on our Web site. Elections, and the technology underpinning them, directly impact our entire society. By improving our election systems, we can directly improve our democracy, itself. The expected results will also be relevant to other application areas.

1 Project Description

Elections are the defining institution in a democracy, and the integrity of the system of elections is essential to the integrity of any democratic nation. The rapid introduction of new election technology in the United States threatens the integrity of our democracy. Today, this technology is being developed, tested, and certified by agencies that are poorly prepared to judge questions about information security. In part, this is because elections pose extremely difficult information security challenges, problems that may be more difficult than the military security problems that have traditionally driven information security research.

Indeed, voting poses problems that go beyond the scope of traditional information security. Every participant in an election is a potential adversary, just as is the case in classical military security. However, military security models (e.g., confidential, secret, top secret) have no direct application when every voter, poll worker, election official, *and* software developer is a potential adversary. Likewise, while secrecy is an important factor in protecting voters from coercion and bribery, the integrity of votes against all forms of tampering is of paramount concern. Ultimately, the purpose of any election system is to provide sufficient evidence to convince the loser of an election that he or she has genuinely lost, even in the face of extraordinary threats. Election systems must be engineered to provide this level of evidence.

1.1 Overview of the Proposal

We propose to form *A Center for Correct, Usable, Reliable, Auditable, and Transparent Elections* (ACCURATE). In considering the voting problems as an end-to-end problem, we intend to adopt a defense-in-depth philosophy of security and not rely on any one line of defense, be it administrative or technical [11].

Broadly speaking, we can divide the proposed work into questions arising from existing voting technology and an exploration of the potential for new voting technology. We must consider techniques both to incrementally improve the security of existing voting technologies and to make more far reaching changes that could significantly alter the state of the art. Among the ideas proposed for this is the voter-verified paper audit trail, but there are many variants of this idea that have not been adequately evaluated and numerous alternatives (cryptographic and otherwise) for which strong claims have been made of their security.

Auditing methods, both those that allow absolute reconstruction of the election results from the original evidence and those that center on statistical evaluation must be explored along with the application of these methods to the canvassing process. Canvassing, the computation of vote totals over the distributed election system, needs to be carefully examined, and given the frequency of clerical errors in real elections, all the operations and procedures involved in the conduct of an election need to be studied with a goal of finding ways to add self-checking while keeping the system as simple as possible.

Looking at next-generation voting systems, we must explore several options for paperless (e.g., Direct Recording Electronic, or DRE) voting systems, including the use of trusted hardware and multiple independently developed components. Approaches to such systems must take an aggressive approach to their evaluation, viewing the system as a potential adversary and seeking out avenues of attack. We would prefer to find these flaws and design repairs or countermeasures before such systems are used in the field.

We intend to aggressively pursue design-for-verification principles, seeking ways to simplify the demonstration that an election is correct despite the presence of many components for which complete, mathematical proof is impossible. The use of voter-verified audit trails is but one model that allows this. We are interested in exploring others, using our knowledge of what can be verified to focus the designer's attention on the key system components where audit is difficult, and then bringing those components out into the open. In addition, we will explore the application of both hardware techniques, such as trusted computing platforms, and software techniques, such as proof-carrying code and assertion checkers to this problem.

Because elections are inherently distributed, we must explore the use of networking in elections, whether over the Internet or private networks, or via hand-carried data. Likewise, we will study the cases of remote

and absentee voting, where the threat of individual voters casting multiple votes is more significant than with traditional voting precincts.

All our work will be informed by studies of usability and accessibility, since election systems must be usable and accessible, not only to voters, but to the myriad of election workers who administer various parts of the system during the election cycle. In addition, because election systems are strictly governed by law, and because these laws make numerous requirements on the conduct of election officials and the design of their equipment, we must focus attention on the broad range of laws governing elections in this country.

While trustworthy elections are the driving application for the research of the center, there are other application areas to which the product or our research will apply. Here are three examples:

Secure auctions Many of the requirements of secure voting systems (e.g., auditability, transparency, usability, secrecy) overlap with the requirements of secure auctions. There is a need to maintain the secrecy of bids until the end, when they are revealed. It is important that the auction be transparent, and that post-auction audits be possible.

Spyware Spyware prevention, detection, and removal on end-user computers is an increasingly important problem. The research into securing voting platforms involves techniques for detecting malicious code and designing systems for verification. This same research will have applications to solving the problem of spyware.

Denial of service The research on remote voting and absentee balloting will involve studying network security problems, such as thwarting distributed denial-of-service attacks.

A major goal of the ACCURATE center will be to explore how the solutions we develop for secure voting systems can apply to other problems such as these.

1.2 Recent History

The problems in Florida in the year 2000 presidential election dramatically highlighted long-standing problems in U.S. elections [42, 44, 48, 9, 99, 77, 49] including high residual vote (under and overvoting) rates [112, 15, 70, 51], insufficient election monitoring [71], and correlations between voter education levels and problems they may have experienced [20, 66]. The Help America Vote Act of 2002 (HAVA) was passed to address these problems [10]. But election problems persist, and public concern appears to be increasing. Notably, in the 2004 Presidential election and in spite of the 3.5 million vote margin of victory, there were widespread claims of major irregularities or procedural complications [38, 21, 54, 111, 101, 102], and concerns around the discrepancy between exit polling and official tallies [104, 27]. More than 38,000 incident reports were collected from voters during and after the 2004 election, including about 900 related to voting technology [37].

Today there are approximately 20 vendors providing voting technology regulated by a patchwork quilt of federal, state and local rules [8, 89]. Many of the problems mentioned above spring from the limitations or defects of voting technologies, and mismatches between those technologies and the regulations and standards governing them. For example, where computers are used in elections, it would seem obvious that they must satisfy rigorous security and reliability standards comparable to those used in other “mission-critical” industries such as aircraft controls, medical devices, or military systems. Comparable standards do not currently exist for voting equipment.

Members of this research team have issued repeated warnings about the situation, since at least 1984 [83, 53]. Since the 2000 election, the broader community of computer scientists has raised an alarm about security issues in computerized voting, including the Resolution on Electronic Voting (which was written by a few members of the research team and endorsed by many of them) [109] and a recent resolution by the Association for Computing Machinery, the largest organization of computer professionals [16].

In the U.S., voting system certification is a state function. There are federal guidelines promulgated by the Federal Election Commission and the National Association of State Election Directors in 1990 [39] and revised in 2002 [40] (the FEC/NASED standards). More recently, HAVA has transferred the authority for updating these standards to the newly formed Election Assistance Commission, which is expected to announce new standards at some point in 2005.

The current standards are “voluntary,” meaning that states can ignore them. However, while some states have *no* certification process for voting equipment, most states require conformance to the FEC standards as a prerequisite for state-level certification. Systems conformance to the standards is decided by private, Federally certified Independent Testing Authorities (ITAs), who inspect the design of the equipment (including source code) and test the actual systems. The ITAs, paid by the vendors to perform their analyses, will only vouch that a system meets the minimum requirements of the FEC standard. An ITA’s written report is considered confidential and owned by the equipment vendor. Although these reports may be shared with a vendor’s customers, few such reports have ever been made public. In addition to the ITA’s certification, many (but not all) states do further testing to ensure that voting systems satisfy their own regulations. In some states, counties are allowed to purchase any system that is “approved” by the Secretary of State. Elsewhere, a single voting system is used uniformly statewide.

In 2003, concerns around the quality of voting equipment standards and certification were reaffirmed when one voting system vendor, Diebold Election Systems, accidentally disclosed the source code for the software used in its AccuVote-TS voting system to the public [43]. This story exploded into the press with the public release by several of us (Rubin, Wallach) of a report documenting serious security flaws in that software (the Hopkins/Rice report) [63]. Although the vendor has strenuously denied the significance of these flaws [35], subsequent reports commissioned by the state of Maryland from Science Applications International Corporation (SAIC) [100] and RABA Technologies [95], and by the state of Ohio from InfoSentry [56] and Compuware [33], substantially confirm all the major security flaws identified in the Hopkins/Rice report. The InfoSentry and Compuware reports additionally considered systems made by Election Systems and Software, Hart InterCivic, and Sequoia (which, with Diebold, collectively dominate the marketplace for voting equipment); every system had significant flaws.

The original ITA source-code audit for the system that would later become the Diebold AccuVote-TS system [115], available to one of us (Jones), indicated that the software was the best the ITA had ever examined and that its security was particularly impressive. In light of the security flaws that were missed by that report [53] (Jones testimony)—some of which were still present many years later—it is difficult to have any confidence in the present certification regime.

Of course, security is not solely an issue of equipment design. Appropriate procedures are at least as critical, whether the system is completely manual or highly automated. For example, an accepted security practice in the administration of general-purpose computers is to track and install the latest patches, both to fix bugs and to defeat security attacks. Likewise, many organizations place strict controls on what software versions are acceptable, to ensure smooth interoperation and predictable behavior. Unfortunately, elections administrators appear to be far more lax about such procedures. For example, an audit of the voting equipment used in 17 California counties determined that uncertified software versions were in use in *every one* of these counties in November 2003 [105].

Insufficient technical oversight and lax administrative procedures open the door for election fraud. Even in the absence of any genuine fraud, both parties in a tightly contested election can be counted upon to leave no stone unturned in their search for any “lost” votes that might affect the results. By improving engineering standards, the certification regime, and administrative procedures, we can hopefully provide convincing evidence to support the official tallies, even in the tightest of elections.

1.3 Requirements for Secure Voting Systems

To operate an election to the satisfaction of all parties concerned, a number of properties must be satisfied (e.g., [72, 81, 82]).

Correct capture is the property of recording each vote exactly as intended. This property can be compromised by voter error, possibly exacerbated by usability concerns. Likewise, hardware might fail, software might have bugs, and deliberate tampering might corrupt either hardware or software. *Correct counting* is the property that each vote is counted as it was originally captured. Counting can likewise be compromised by software errors and hardware failures, as well as through tampering with the transmission of vote records.

Secrecy is extremely important to voters. Ballot secrecy requirements usually go beyond privacy requirements in other domains, requiring that voters not be able to prove how they voted (even if the voters desire to do so), to defeat vote selling and coercion.

Voting systems must be *auditable*. It must be possible to reconstruct the results independently from original records of the votes, which requires that these records be kept secure from accidental or intentional modification until the audit occurs (and thereafter as well). At the very least, even if it is not possible to recover from all failures, it must be possible to detect failures. Also, no voting system should permit the possibility of undetected fraud.

Elections must be *transparent*, which requires that observers understand election technology and procedures well enough to be able to attest to the quality of the election. Elections must be trusted by the populace, and must be justifiably trustworthy. The legitimacy of election results must be so far beyond reproach that even the losers are convinced to accept the winners of the election.

Voting systems must be highly *available*. Nothing should prevent voters from casting their votes, including software bugs, hardware failures, or loss of power. Backup equipment or alternative (e.g., provisional) paper ballots should be available. All voters need election systems to be *accessible*, regardless of disability, language fluency, literacy, or other factors. Likewise, an election must be *administrable*, even for election officials with insufficient funds and poorly trained poll workers.

Of course, elections need to be conducted at reasonable *cost* while adequately achieving these other properties. To this end, and to accelerate the deployment of solutions to the many problems of elections, it is desirable to encourage *interoperability* among components of election systems through conformance to *widely agreed-upon standards*. For example, standard interfaces for voter-verifiable printers and standard data formats for ballots would directly lead to superior election systems through increased competition.

2 Research Plan

2.1 Fully Electronic Voting Systems

An important threat against any electronic voting system is *software tampering*, whereby an attacker might try to install some “Trojan horse” logic to cause the voting system to bias its results in some fashion. Even software testing, performed concurrently on the day of the election, cannot necessarily detect the presence of such tampering. To date, the only solution known to mitigate these risks is to have the voting machine print a voter-verifiable paper ballot. However, there may be other possible solutions. We will investigate novel architectures for paperless systems.

An intriguing possibility is to break a voting machine into separate parts, built by unrelated vendors, which must cooperate to produce the final tally [22]. For example, one part might interact with the voter to produce a ballot, and a second might ask the voter to verify his or her choices, while a third records the ballot for canvassing. We must assure that these are independent and do not collude; this will require research. We will study how such a system can be designed and implemented to be usable despite the large number of components with which the voter interacts.

Another intriguing possibility to consider is the application of recent *trusted hardware* concepts that allow a computer to *attest* to the software that it is running. We discuss these concepts in more detail in Section 2.4. Our research will consider whether cooperating machines or trusted hardware components may be able to increase the resistance of voting systems to tampering.

2.2 Design for Verification

Electronic voting machines consist of hardware and software. One of the most serious threats to the integrity of these machines is the possibility that an insider with access to the development environment might insert malicious code into the software that would undetectably alter the outcome of the election. A major challenge, then, is either to find ways to build voting machines that are verifiably correct, or to find system architectures that eliminate or reduce the need for verification. We will investigate both directions, and we expect that any solution will most likely need to rely on ideas from both approaches.

2.2.1 Building Software That Is Verifiably Correct

When election integrity relies on software to perform as expected, software must be deemed verifiably correct by independent observers. Unfortunately, automated reasoning about software is very difficult, and the state of the art in commercial software development is unable to support this goal. One promising research direction developed to date for dealing with this problem is *proof carrying code* (PCC) [79], where programs carry *proofs* with them that they do not violate some safety property. PCC is still the subject of active research and has not been applied to software on the scale of a voting system. Another promising direction is the use of tools supporting formal methods for building verifiable software. Examples of such tools include SRI's PVS [88] and SAL [34], Bell Labs' SPIN [52], Stanford's Murphi [36], CVCL [106], JML [23], and Spec# [18]. We will study how these may be applied to the design of voting systems.

Unfortunately, proving the absence of malicious code may be too difficult a problem to solve with PCC or theorem proving alone. Therefore, we propose a complementary research approach that may prove more tractable and that offers promise in avoiding maliciously installed software by insiders who develop electronic voting machines. The idea is to design software in such a way that it is easier to verify. Whereas the general problem of detecting malicious code is intractable, a constrained development environment might make it much more difficult to hide malicious code and avoid detection. The malicious code would have to conform to the design constraints, and this limits the flexibility of the attackers' design space. For instance, because randomized software is harder to test than deterministic software, we might restrict access to random-number generators and other sources of nondeterminism. Similarly, real-time software is harder to test, so we might restrict access to the real-time clock hardware [59]. Or, we might use programming environments that use logging and checkpointing so that all computation is replayable. The replayable record would be a strong deterrent to would-be attackers. These are just examples from a very large range of possible constraints; unfortunately, many of these violate fundamental constraints imposed on elections, so each can be applied to only part of the system. As part of this research, we propose to study and experiment with different ways one might constrain the development environment to maximize the potential for verification without violation of fundamental constraints or overly interfering with legitimate development. Similar constraints may also be applicable toward controlling spyware.

Another necessary component to an assurance argument is a secure configuration management system. It does little good to analyze code (at either the source code or object code level) if we cannot assure that the code being analyzed is actually the code being used. We also seek strong auditability of changes made to the system at all levels, from requirements to executable code. We absolutely must prevent the recent situation in which California counties using Diebold's DRE systems found themselves running uncertified code, contrary to state law [105]. We will explore the use of configuration management software on top of

operating systems that support mandatory integrity policies (e.g., SELinux). By effectively combining operating system integrity guarantees, along with the configuration management system's audit trail, we can gain additional assurance in our repository. We will investigate combining cryptographic integrity protection, as in OpenCM [103], with distributed configuration management, as found in many commercial products.

Finally, we will measure how successful we have been by attempting to break our own designs. In the security community, it is widely accepted that a system may be considered secure only after it has been subject to intense and continuing attempts to break it. This adversarial process has been applied successfully in the design of a variety of systems. While this process does not guarantee security, it is a useful way to gain increased confidence in the systems we build.

2.2.2 Tolerating Software That We Cannot Verify

We are concerned that reliance on verification tools and secure development systems greatly enlarges the trusted base, so are also interested in finding ways to minimize the need for trust. If we cannot verify software, we can attempt to exclude it from the trusted base of the system [60]. For example, if voters use a touchscreen machine to enter votes, and then examine a machine-printed paper ballot to confirm its accuracy before depositing the paper ballot in a ballot box, the entire vote-entry and ballot-printing machine is removed from the trusted base [58]. Rebecca Mercuri's voter-verified audit trail [72, 73] has a similar property. When a system is composed of a mixture of trusted and untrusted components, the interfaces between these components must be examined with great care. We are particularly interested in designing interfaces that guarantee the absence of covert channels [60].

An important principle here is that deterministic computations can easily be audited by external observers if both the inputs and outputs are published so that the observers can duplicate the computations to check the published output. This means that computations on public data can generally be removed from the trusted base; trusted software is required only where secrets must be guarded. Thus, a promising direction is to investigate architectures that can make public as much of the computation as possible.

A second conjecture is that statistical testing can sometimes be an cost-effective way of auditing results. For instance, California state law mandates a manual recount of a random sample of 1% of the votes as a way of checking the operation of optical scan vote-counting machines. We will investigate how random sampling and other methods can be used for *probabilistic audit*, we will seek architectures that maximize the utility of probabilistic audit, and we will study what this technique can and cannot achieve (for instance, its value in detecting malicious code or fraud).

2.3 Novel Cryptographic Techniques

We intend to study the broad applicability of cryptographic techniques to voting systems. Our focus will be on simplicity, that is, designing cryptographic techniques that can be understood by an intelligent lay person (as opposed to a crypto specialist). We mention only two promising directions here.

One relevant technique, called *mix nets* [29], can help ensure voter privacy while enabling public validation of the election. This property, known as *universal verifiability*, is appealing; however, a number of issues remain before it can be put to use. Most important is ease of use and simplicity. Is there a simple cryptographic mechanism that provides universal verifiability? Currently, the principles underlying mix nets are beyond most voters and election officials. The question is whether one can achieve universal verifiability by using a simpler mechanism that is easy to use, administer, and understand. Chaum [30] describes an idea in this direction by using visual cryptography. Neff [80] describes a complex scheme as well. We intend to pursue additional directions with the goal of simplicity in mind.

Another difficult problem is ensuring that a voting machine correctly records the voter's intent. Voters should be able to check that their votes were correctly recorded without the ability to prove to a third

party how they voted. These seemingly contradictory requirements can be addressed using cryptographic techniques. The challenge is to build sufficiently simple mechanisms that ask the voter to perform only simple tasks, such as picking a random element from a short list. Surprisingly, such simple steps give rise to challenge-response protocols that are sufficient to catch misbehaving voting machines. To obtain the simplest system we might rely on a back channel to the voter that is invisible to the voting machine (e.g., physical mail or printed paper). Neff recently presented a first step in this direction. We intend to explore other solutions to this important problem. It is worth noting that mechanisms that let users test that their intentions were correctly recorded could have applications beyond electronic voting. For example, some forms of sealed-bid auctions have similar security problems as voting: bidders need assurances that the auction was executed according to the rules, and bids need to remain secret after the auction is complete.

2.4 Trustworthy Hardware Platforms

Recently, IBM and HP, among other companies, have begun adding hardware conforming to the Trusted Computing Platform Alliance (TCPA) specification to their normal computers. There is a long history of research and development of tamper-resistant processors, recently exemplified by the IBM 4758. Recent research [69] suggests that we may be able to use the security guarantees of TCPA hardware for secure bootstrapping; thus making electronic voting systems more secure against malicious code and unauthorized tampering. We will examine both the role and impact of CPU tamper-resistance features in architectures for secure electronic voting, and whether commercial TCPA hardware can serve this purpose. A major issue we will study is trusted path: how can we assure that input (e.g., from a touchscreen) is not tampered with on the way into the tamper-resistant hardware? Furthermore, if we assume the presence of tamper-resistant hardware in every precinct, is it possible to leverage this hardware in support of canvassing activities? Also, we will consider the use of inexpensive very-thin-client trustworthy devices for vote casting.

The TCPA hardware can also be put to use for configuration management. While many are uneasy about the possibilities of widely deployed TCPA hardware [14], a low-level mechanism that will run only properly digitally signed code could be of great benefit to voting system integrity. A side effect of this mechanism is that “accidentally” running the wrong code will not be the result of a simple mistake.

While studying the trusted path problem for voting machines and how TCPA can be used for configuration management, we will also explore how our research can be used to help mitigate the problem of spyware. We believe that there are parallels between the problems of leveraging secure co-processors for voting machines and ensuring that there is no malicious software spying on a user’s activities.

2.5 Use of Networking in Voting

The last several years have seen a strong push toward electronic voting. It is only natural to consider the possibility of network-based or Internet voting; in the most extreme model, this allows citizens to vote from their personal computers at home. Internet voting is now a reality in Geneva, Switzerland, and in the United States, Internet voting has been used in several primaries. The first was the Arizona Democratic presidential primary, in March of 2000, in which approximately 85,000 votes were cast and counted. The Reform Party national primary was also conducted over the Internet that summer, as were various nonbinding Internet voting experiments in some counties of Washington, California, Arizona, and elsewhere. The use of the Internet for the Michigan Democratic Caucus in 2004 is also noteworthy, if only for the fact that it appeared to be an Internet-based election but without any of the protections of a secret ballot; in fact, all remote-site or absentee voting sacrifices many of the protections of ballot secrecy, for example, freedom from coercion, no matter what the technology.

There have been several important studies of Internet voting: The first was by the California Secretary of State’s Task Force on Internet Voting, whose January 2000 report [25] clearly articulated most of the

security issues regarding Internet voting.¹ Another study was conducted by the Internet Policy Institute. Its report [78], published in March 2001, stated “Remote Internet voting systems pose significant risk to the integrity of the voting process, and should not be fielded for use in public elections until substantial technical and social science issues are addressed.” Other reports have come to similar conclusions [26].

Last year, the Department of Defense considered fielding an Internet voting system called SERVE for overseas civilians and military personnel. The project was abandoned when a report,² co-authored by two of us (Rubin, Wagner) and two others, showed that the security concerns were too serious for SERVE to be used, even experimentally.

While much of the focus on the use of networks in voting has centered on the Internet and voting from the home, much of the criticism leveled against Internet voting also applies to the use of other network technologies, from wireless networks to the telephone network, and much of it applies to seemingly conservative uses of networks including the transmission of votes from precincts to a central tabulation facility or the transmission of election results from a tabulation facility to a public Web server and to the press.

We will pursue solutions to these problems, both technically sophisticated and as plainly simple as possible. For example, while networks could be used to make election results available for public consumption, they could also provide an avenue of attack for the public to reach the tabulation systems. Various firewall technologies could be used, but these add large components to the trusted base. Physical modification of the underlying communication channel is an attractive approach to avoid this. If we are careful, physically one-way communication channels can be made obvious even to relatively unsophisticated observers. For example, an LED transmitter could be paired with a photodetector to communicate data through the glass window of a secure tabulation facility [60].

Where general-purpose networking is necessary, such as with Internet voting, we must be robust against any possible threat. In addition to attacks against end-users’ machines, which might be addressable through emerging trusted hardware technologies (see Section 2.4), we must be concerned with distributed denial-of-service attacks (DDoS) against election servers. Likewise, we must be concerned with attacks against the BGP routing infrastructure of the Internet itself. Spammers, for example, have been caught using corrupt routers to literally reroute the Internet, giving their mail servers different and untraceable IP addresses while sending mail. Similar attacks could be expected against networked election systems.

Several members of this Center are active in designing robust network technologies that may be able to at least partially address these concerns. Wallach has been studying the security of *overlay networks* [28, 85, 86], a family of promising new techniques increasingly being used for content distribution and information sharing. Rubin has been studying the security of BGP routing, which remains one of the most vulnerable components of the Internet. Securing this critical protocol is an active area of research [62, 61, 67, 46, 91, 12, 116, 97, 113, 68, 55, 41]. The networking research Wallach and Rubin are already doing will complement the efforts of the Center. The resulting research from the synergies of these projects may lead to more options than are available today for secure remote voting, as many of these techniques will be well suited to election systems. For example, data replication techniques from the overlay networking community may be useful to preserve and archive election results, even in the face of a massive, coordinated campaign to disrupt an election. By leveraging and applying security research from the networking community, we may be able to achieve a level of robustness comparable to or exceeding traditional paper-based election management.

2.6 Remote and Absentee Voting

To preserve ballot secrecy and anonymity, it is clearly preferable to require voters to vote in polling places that offer the necessary privacy. However, voter turnout can be increased if voters have a way to vote early,

¹Jones’s critique of this report is available online at <http://www.cs.uiowa.edu/~jones/voting/california.html>.

²The report is available online at <http://servesecurityreport.org/>.

whether by mail or by visiting a designated polling place. Also, most jurisdictions make provisions for absentee voting by mail, to allow voters who are unable to appear at the polling place in person to vote. Mail-in ballots offer great convenience to voters, notably including soldiers serving abroad, who cannot return home to vote in their home precincts. However, mail-in ballots can also easily be sold by or coerced from voters. We would prefer a system that has the flexibility of mail-in ballots with the privacy guarantees of a secure polling place. Ideally, voters should be able to cast their ballots at *any* polling place in the state, but this will require replacing how voters are currently authenticated.

In many states, voters can “authenticate” themselves merely by stating their names and signing in a book. While it might seem natural to require voters to produce ID cards, single-use ticket stubs, or some other proof of their eligibility to vote, such measures might disenfranchise voters who cannot find their single-use tickets or might enable vote selling. Also, requiring the production of ID cards is seen to be intimidating by numerous minority groups. Traditional forms of biometric authentication may be unacceptable for the election setting, because many voters fear governmental collection of such data, and they may choose not to vote if they are required to use such systems.

Hybrid systems will be considered, where voters voting in their home precincts can validate their registration in the same fashion as they always have. However, voters wanting to vote remotely can request, in advance, suitable credentials that can prove their identity at any polling location. We have already investigated the use of visual cryptography in such an endeavor [90], and we believe that other cryptographic measures, perhaps borrowed from the digital cash literature [31], would allow each voter to cast one ballot anonymously, but would reveal the identity of any voter who attempted to vote more than once. In addition to normal cryptographic soundness proofs, we would need to investigate the usability and accessibility of our scheme as well as the additional cost and risks of sending such credentials to every voter, most likely through postal mail.

2.7 Operations and Procedures

We aim to avoid the trap of overemphasizing novel cryptographic techniques at the expense of good administrative procedures. In the same fashion that credit card numbers can be stolen from e-commerce Web sites (long after the “secure” Web connection has been decrypted), similar threats apply to voting systems. Cryptography may be necessary, particularly when data is moved over computer networks or phone lines, but it is only a small part of a big picture.

The design of reliable, secure, and trustworthy tallying procedures must involve an understanding of the audit criteria used to provide assurance that the results are correct. In the same fashion that banks will carefully design their procedures to prevent a solo bank teller from stealing money from that teller’s drawer without being detected, and likewise to catch common clerical errors that tellers will make, we need similarly crafted procedures for use in vote tallying.

These policies and procedures must satisfy a wide range of needs. To preserve the accuracy and integrity of an election, the chain of custody over any electronic or paper records must be carefully maintained, using a combination of physically tamper-evident seals and strong cryptography. Of course, to use cryptography correctly, it will be necessary for a county or state to distribute appropriate key materials to the voting terminals, whether using public key cryptography or symmetric key systems. We propose to study alternative models for secure key distribution and storage, for example, the use of internal smartcards to isolate the key from the rest of the voting system. We also intend to analyze the ability of “trusted hardware” designs (see Section 2.4) to securely manage these signing keys, among other tasks.

Voting systems, at their core, are distributed systems with a variety of protocols. These protocols include the smartcard-to-voting-terminal exchanges that occur in voting systems such as that of Diebold [63] as well as the spoken exchanges between voters and poll workers, between election administrators and their tallying equipment, and so forth. Tools from the theorem proving and model checking communities (see

Section 2.2.1) that have been used to examine a variety of networking and cryptographic protocols should be applicable to these exchanges between humans and machines during elections. We can model sealed ballot boxes with serially numbered tamper-evident seals in much the same fashion as we model encrypted messages with digital signatures and nonces. In the same way that formal tools have found subtle flaws even in widely deployed cryptographic protocols, we expect we can find similarly subtle flaws in the policies and procedures used to operate elections.

Where some election procedures require random sampling and recounting of ballots, we can use formal models of election procedures to measure the odds that certain amounts of fraud or error could escape detection. Furthermore, we will build formal tools to exhaustively simulate the extent to which a single faulty tabulation machine or corrupt election official could affect the reported election results. Such techniques will help identify weak spots in policies and procedures that need to be redesigned.

2.8 Usability and Accessibility

Usability by a broad public is particularly important in voting systems. No matter how secure and reliable a voting system is, if that system places demands on the voter such that he or she is unable to vote successfully, or is made uncomfortable with doing so, voters will be disenfranchised. Voting is a particularly challenging human factors problem (a prominent usability professional has recently termed it the "ultimate usability problem" [94]) because voting systems must be usable by citizens regardless of age, disability, education, socioeconomic status, history of computer use, literacy level, native language, and the like. A successful system must go beyond simple usability in terms of the voters' ability to accurately cast their votes, but also must produce confidence that their intent was accurately recorded and tallied. This problem will not simply be solved merely by the application of computer technology, as demonstrated by studies of the 2004 election in Florida [75], which showed that touchscreen voting machines generated 50 percent more undervotes than paper optical scan ballots.

Despite the breadth and depth of the problem, voting has received surprisingly little attention from the human factors community. The most important exception is a 2004 report of the NIST [64] which summarizes research results (such as Bederson et al. [19]) and sets out a detailed outline for future work in this domain. The report mentions that only a handful of studies have been conducted and that usability and accessibility standards and guidelines are nonexistent in this domain. The report goes on to recommend that research to support the development of such standards is particularly important in the three usability areas laid out in ISO 924-11: effectiveness (i.e., votes are for intended candidate, no errors), efficiency (i.e., voting takes reasonable time and effort), and satisfaction (i.e., the voting experience is not stressful, voter is confident). Further, these metrics apply to individuals with disabilities, so accessibility can be defined as usability for such individuals.

Our approach to this problem will consist of assessments of the three metrics by three different methodologies: laboratory usability testing, field usability testing, and usability analysis. These assessments will first be conducted on current voting systems (e.g., paper ballots, extant electronic systems) and then applied to new designs and technologies developed by the Center. One of the goals of the Center would be to communicate our results to bodies like NIST to inform the development of standards and guidelines.

Laboratory usability testing is empirical testing of voters (or potential voters). We intend to draw from two primary populations: Rice University undergraduates, representing in some sense the "best-case" scenario (i.e., highly educated, low rate of disability, high general visual acuity), and local Houston residents, recruited through newspaper advertisements. While this will not generate a completely representative sample, it should be substantially more diverse than the undergraduate sample. These participants will be brought into a laboratory environment and observed interacting with voting systems using quantitative objective techniques (i.e., performance measurement of time and accuracy), videotaped "think aloud" protocols, and

subjective satisfaction measurements, such as the QUIS [32] or SUMI³ instruments.

Field usability testing consists of similar measurements taken outside the laboratory. To collect the widest and most heterogeneous sample possible, participants will also have to be recruited and the voting systems assessed in participants' own neighborhoods. This should significantly increase the generality of the sample, particularly since the Houston area contains a strong diversity in terms of socioeconomic status, ethnicity, education, and so on. Testing in remote locations will by necessity rely less on video and objective measures, but they need not be entirely eliminated.

Efforts will be made in both the laboratory and field studies (especially the field studies) to recruit participants who are likely to have particular difficulty with voting, representing groups such as the visually disabled, the non-English-speaking, the elderly, and those with low socioeconomic status. We will furthermore attempt to determine if different voting technologies differentially impact such groups on any of the three measures described. This may have policy implications as well; for example, if a particular technology tends to lead to low effectiveness among the elderly, then use of that technology will differentially disenfranchise elderly voters, creating a clear policy concern.

Usability analysis will be used to augment these empirical methods. This includes informal methods such as heuristic evaluation [87] and more formal approaches such as Cognitive Walkthrough [93] or GOMS analysis [57]. In addition, computational/mathematical techniques such as Information Foraging analysis [92] or computational cognitive modeling [24] may be applied as well. Such quantitative analytic methods have proven to be valuable in other contexts (e.g., Gray et al. [47]) and are a particular specialty of the co-PI responsible for this portion of the project. Useful sources of information for such analyses are domain experts, such as local election officials. Thus, we intend to invite such officials for discussions of their experiences and to attempt to analyze problems they have encountered. Another useful source of information will be usability-oriented incident reports from the 2004 elections, including those available from the Usability Professionals' Association.⁴

Finally, migration of any techniques or technology developed by the Center to new domains such as auctions or spyware prevention will raise new human factors issues. For example, in online auctions, it is important to provide effective and efficient interfaces because the wide variety of auction types carry with them different security properties, and users must be clear about these properties. However, if the security measures generate significant usability costs, then such measures are unlikely to be successful in this domain; users may simply choose not to participate.

2.9 The Nexus of Policy and Technology

Historically, U.S. elections were relatively technically consistent and simple. Contests were decided by a show of hands, by depositing objects in containers, or by writing choices on slips of paper [13]. Over time, as populations grew larger, as ballots became more complex and with the introduction of complicated voting technology exacerbated by unevenly distributed resources, this early technical simplicity and consistency gradually eroded. At some point, the inconsistencies introduced by differences in technology choices and procedures become unconstitutional in the sense that there is no longer a guarantee that every person has the opportunity to vote and that each vote is counted as it was intended to be cast [1, 2, 3, 4, 5, 6, 108]. The result is "disenfranchisement by design" [98]. The disparities introduced into the system by technology and procedure are of increasing concern to election officials, candidates, legislators and election monitoring organizations (see Section 1.2).

Given that elections and other core government functions requiring consistency, reliability, transparency, and trust will be increasingly dependent upon technology, we must determine what mechanisms best ensure

³<http://sumi.ucc.ie/>

⁴See http://www.upassoc.org/upa_projects/voting_and_usability/voting_usability_problems.html for more details.

that the technology meets policy goals. A better understanding of the relationship between legally espoused goals and the process of setting standards, policies, and procedures that ultimately shape and choose technology to implement them is essential. In the election context mechanisms are needed to ensure that technology meets our commitments to equal protection and nondiscrimination (racial equality, multilingual access, disability access, interjurisdictional equality), privacy, and security; accurately captures voter intent; and is auditable [108, 9]. Through law, standards, certifications, and procedure we must guarantee that voters and votes are equally treated during the voting process (casting, counting, and recounting). With such a wide variety of voting technology currently in use or in development, there are different high-level and low-level needs for standards, testing, and certification. It is imperative that we identify mechanisms that can equalize the voting process given the diversity in election technology.

Our research aim is to consider the role of standards, certification, procedures, and procurement in conforming voting technology to public policy goals. As mentioned earlier, voting system usability needs additional research. The recent NIST report [64] identified usability as an area where standards are sorely needed. We, along with others, conducted a preliminary analysis of human factors issues in the 2004 general election [76] and found ample support for high- and low-level usability and human factors standards as part of the Voting Systems Standards [40]. Along with problems experienced in the field with technology, there have been several issues surrounding HAVA [10] implementation that reflect the complexity introduced by the need to translate legal rules into technical requirements. For example, the implementation of HAVA's audit requirement has been contentious, arguably misinterpreted to allow for paper records printed at the close of polling rather than contemporaneous with each vote cast [114]. The disagreement is in part a reflection of the difficulty posed by the need to translate values protected by the intrinsic properties of paper (manual recount) into standards that protect the value, yet leave room for different implementations. Similar comments apply to inconsistent application of the HAVA requirements for provisional ballots, which emerged in the 2004 election [50, 107]. While the need for standards that ensure consistent treatment of all individuals is apparent (e.g., across race, language, disability, and jurisdictional lines), the processes for identifying additional areas where standards could be helpful and for relating standard creation to policy goals and implementations are in need of further research.

As technology becomes more complex and elections become increasingly reliant upon technology it is necessary to consider what level of access, review, and openness of code is necessary to ensure that the standards, testing, and certification are capable of verifying an election technology's ability to support election policies. The increasing use of technology in elections throughout the 20th century has reduced their transparency [95, 74].⁵ The number of individuals who have the technical ability to evaluate the increasingly complex machines is decreasing and, more important, the limitations on access to code base of the systems placed on those who are capable of evaluating them undermines our ability to know whether they will perform as promised. This "enclosure of transparency" and the fundamental tension between openness and proprietary systems is a formidable barrier to those responsible for selecting technology, establishing procedures, and running elections. Equally important is exploring technical, legal, and regulatory tools for increasing election system reliance on open, verifiable, systems.

Finally, in a complex system with a widely varied user base and disparate use conditions it is important to develop procedures ensuring that usage problems are fed back into the process of establishing rules and standards. For example, today there is no standard system for reporting technology problems during elections, nor is there a process for feeding such reports back into the standard and certification process. Election incident reports would be useful feedback into the standards-setting and testing process and would help to ensure that election incident knowledge is retained within the system. This research is relevant to a host of application areas where the performance of technology is critical to the attainment of public policy

⁵Note that bodies such as the EAC and California's Voting Systems and Procedures Panel are, in part, regulatory efforts to increase transparency.

objectives (see, e.g., Lessig [65], pages 135-136).

3 Education, Outreach, and Technology Transfer

3.1 Education

A significant fraction of the budget for the proposed Center will directly support graduate and undergraduate student involvement in the Center's research program, and some of us have a long record of involving undergraduate students in our research projects. It is noteworthy that Adam Stubblefield and David Price (now at Johns Hopkins and Stanford, respectively) both published research they did as undergraduates with one of us (Wallach); Stubblefield also had summer internship with two of us (Dean and Rubin, then at Xerox PARC and AT&T Research, respectively). Stubblefield's work earned him one of the two Computing Research Association Outstanding Undergraduate Awards in 2002.

Several of the principal investigators teach computer security courses to graduate and undergraduate populations that make extensive use of examples from election technology (Jones, Rubin, Wallach, Wagner), and ACCURATE will contribute to the development and improvement of these courses. One of us (Mulligan) directs the flagship legal clinic on technology policy where law, engineering, and computer science students engage in interdisciplinary research and advocacy on technology policy. Another of us (Jones) has taught a seminar on computers in elections that reached out to political science and business students as well as computer science students. Two of us (Wallach, Rubin) have taught courses where students built rigged voting systems that were then subject to security audits by other students [17]. We plan additional course projects based on this adversarial process and emphasizing *design for verification*. We will also design course modules to convey these concepts in a fashion suitable for integration into freshman- and sophomore-level introductory programming courses. Material produced by the principal investigators is already being heavily used in security and elections courses at Virginia, Carnegie-Mellon, Princeton, George Washington, and other universities.

All of the institutions participating in the proposed Center have organizations that support and encourage women in computer science or more broadly in engineering and the sciences. We expect that these organizations will be important contact points for involving students in the research program of the Center.

3.2 Outreach

CyberTrust issues, in general, and issues of election technology, in specific now attract broad-based public attention. The principal investigators have all been heavily involved in public outreach on these subjects, not only in academic settings but in the community through invited talks, panels, and media coverage. We will encourage students involved in the Center to join in these activities, training them for public relations and engaging them in our public activities.

Among our Center PIs are some who have long been involved in public elections, working as election observers and precinct election judges, and serving on county and state election boards, committees, and task forces. In the past, several of the principal investigators have involved students in some of these activities; for example, students at Johns Hopkins have served as election judges after taking a security course that covered all the recent studies of electronic voting [26, 25, 78, 100, 63].

The Center will host annual public workshops, bringing us together with election officials, technologists, vendors, and other stakeholders including political activists, civil rights watchdogs, and minority interest groups. For those who cannot attend our workshops, we will record all the talks and sessions, providing streaming video as well as text transcriptions through our project Web site. For organizations that are not yet online, we will produce DVDs of the same material. Grant funds will also be used to support travel expenses for workshop participants and honoraria for invited speakers. Meeting locations will change every year, but

the first meeting will be in Washington D.C. — the political center of the country and geographically near Johns Hopkins University, where the Center will be centered.

In addition to recording our public workshops, our Center will produce a variety of written and video materials, made available on the Center Web site. In collaboration with groups such as the Verified Voting Foundation, we will document best practices for election observers, and we will make recommendations to election administrators. All the results of our work, whether video, audio, written, or software, will be made available to the public, free of charge, on our Web site, and at minimal cost on DVD.

3.3 Technology Transfer Plan

The major vendors have participated with many of us on the IEEE voting standards development team, and have a vested interest in having improved accuracy, integrity, reliability, usability, auditability, and so on in their products. We expect that they will engage in ongoing discussion with our Center, and potentially offer products for testing and evaluation. Because of a variety of limitations (e.g., closed-source proprietary software, competitive vendors), our proposal explicitly does not include any of the existing commercial vendors as direct participants.

We will make a special effort to involve the Open Voting Consortium (OVC),⁶ a nonprofit group devoted to exploring the application of the open-source development process to the domain of voting and elections. We hope to be able to provide technical guidance to OVC volunteer developers, we plan to use the products developed by the OVC as test cases, and we may be able to use the efforts of OVC volunteer developers to implement and test results of our work. Furthermore, we see no problems with cooperating with any other voting system vendor willing to frame that cooperation in terms of an open-source development model. At least two expected contributors of hardware (particularly trusted computing platforms) and software have expressed their intent to interact technically with Center participants as well.

The FEC/NASED voting system standards define themselves as “working” standards, as does the IEEE standard now under development. At the same time that these standards are being used to certify particular voting systems, they are open to ongoing revision in response to changes in the technological landscape as well as changes in law and voting practices. The ACCURATE Center should be able to provide valuable resources to these standards efforts. Much of our proposed work is directly relevant, from the appropriate use of cryptography, criteria for evaluating the auditability of voting systems, and alternative models of voter verifiability, to studies of human factors and the legal context. Our work on adversarial testing should also contribute to the standards process, as should our studies of the relationships between system elements and our work on canvassing and recount procedures.

The IEEE has recently created a subgroup of its voting systems standards effort to study data transfer, hoping to define a protocol that can be used for cross-platform communication of ballot layouts. The OASIS Consortium⁷ is also interested in developing such protocols. We are particularly well qualified to evaluate and assist in the inclusion of security, reliability, and auditability features into such protocols and the development of protocols for the secure distribution of software updates to voting systems.

4 Results from Prior NSF Support

Drs. Dean and Neumann, and Professors Byrne, Jones, and Mulligan had no NSF support in the last 5 years.

Prior NSF support for Avi Rubin (a) Award number: G420-E46-2130-2000 Amount: \$616,923 Period of support: 10/1/03 - 9/30/06. (b) Title: Towards More Secure Inter-Domain Routing. (c) We evaluated

⁶See <http://www.openvotingconsortium.org/>.

⁷See <http://www.oasis-open.org/>.

historical BGP (Border Gateway Protocol) data to examine how security solutions would have performed under peak loads. We implemented several of the BGP security solutions in a simulator, and have an implementation of the IRV system. This work has been funded for just over one year, and we are in the process of writing a paper for publication. (d) Our first publication is in submission.

Prior NSF support for David Wagner (a) NSF CCR-0093337, \$268,000, 3/1/01 – 2/28/06. (b) Title: CAREER: Security in the Large: Gaining Assurance in Real-World Systems. (c) This grant supports research on model checking, lightweight formal methods, and domain-specific heuristics to detect security bugs in legacy systems. We have developed BOON, a program analysis tool that finds buffer overrun vulnerabilities in C code, and MOPS, a model checking tool that is used to find dozens of security bugs in C applications. We have shown how to use CQual, a type-inference tool, to find format string vulnerabilities in C programs. (d) This grant has resulted in more than a dozen publications on software security, cryptography, and related topics.

Prior NSF support for Dan Wallach (a) NSF-CCR-9985332, \$200,000, 4/1/00 – 3/30/04. (b) Title: CAREER: Security and Resource Management in Type-Safe Language Environments. (c) This project aims to add protection semantics, normally associated with operating systems, to language runtime systems to support the concurrent execution of multiple untrusted programs within the same runtime. We have developed a technique for rewriting programs to guarantee that they will terminate without destabilizing other programs using the same language runtime. We have developed memory accounting within a garbage-collected runtime. This grant has also partly supported other security-related work, including performance measurement of SSL systems, studies of copy protection systems, and the study of security issues in wireless networks. (d) This grant has directly supported ten publications and partly supported another five on topics in computer systems and security.

Prior NSF support for Dan Boneh 1(a) NSF grant CCR-9984259 (CAREER), \$225,000, 2/00 – 1/04. (b) Title: Security for Handheld Devices and the Web Environment (c) During this project we discovered the first usable Identity Based Encryption scheme. We also worked on message integrity in a multicast environment. We proposed a new encryption mode for the RSA and Rabin cryptosystems that provides a high level of security and is much simpler than previous constructions. Finally, we developed a digital signature scheme where the signatures are half the size of current popular digital signatures (e.g., DSA). Short signatures are important in environments where humans manually type in the signature. (d) This and the following NSF project resulted in more than twenty publications. 2(a) NSF grant CCR-9732754, \$160,000, 10/98 – 6/01. (b) Title: Hardness of Computing Fragments of Secret Keys in Diffie-Hellman and Related Schemes (c) We began by studying the feasibility of using the PalmPilot for digital payments. We built a digital wallet for the PalmPilot, devising new techniques for managing RSA keys that improve performance by as much as a factor of 5. Other publications that resulted from this project include (1) results on the strength of the RSA cryptosystem, (2) copyright protection, and (3) new anonymous authentication schemes. (d) See above.

Prior NSF support for David Dill (a) NSF ITR CCR-0121403, \$2,100,000, 10/1/01 – 9/30/05. (b) Title: ITR/SY: Computational Logic Tools for Research and Education. (c) This grant has supported research in computational logic, including automated decision procedures, formal verification tools for infinite state systems, programs, and cryptographic protocols, and educational software. (d) Research conducted under this grant by PI Dill's team (one of three PIs) has resulted in seven papers and two PhD theses to date. (e) CVCL is an efficient implementation of decision procedures for quantifier-free first-order logic that is being distributed in open source form over the Web.

5 Management Plan

Our project will have a two-tiered management structure. Avi Rubin at Johns Hopkins University will direct the Center and Dan Wallach at Rice University will serve as associate director. The two will serve as “second-level” managers, overseeing the entire effort and taking responsibility for administration and evaluation. Both have excellent assistants who will provide administrative support, funded through the Johns Hopkins and Rice components of the budget.

An external Advisory Board will be established, with members knowledgeable in election procedure (e.g., county registrars and secretaries of state), election law, the needs of handicapped and minority voters, and computer security. The proposal explicitly avoids selection of the members of this board, in order not to contaminate the reviewer pool; however, one group that will be represented on the Advisory Board is the Open Voting Consortium (OVC).⁸

Professor Rubin will take primary responsibility for the distribution of software, educational materials, and articles produced by our project via the project Web site, which will be maintained at Johns Hopkins. Professor Rubin has led several projects that involved software distribution, such as Crowds [96], Publius [110], and Absent [45], while he worked at AT&T. He has also managed several projects at Johns Hopkins in his capacity as Technical Director of the Information Security Institute, including the electronic voting security research group and the RFID lab.

Professor Wallach will manage the annual meetings (see Section 3.2) over the course of the project. He is well qualified for this job, having organized several workshops and conferences, including serving as program chair for Usenix Security 2001 and co-organizing the South Central Information Security Symposium (SCISS). SCISS, in particular, brings together security researchers from academic institutions in Texas, Oklahoma, and surrounding areas, many of which serve underrepresented groups, to discuss their work in a less formal setting. Our Center’s annual workshops will serve an even more diverse community.

For each of our major technical themes, one or two senior project personnel will serve as a “first-level management team,” and Professors Rubin and Wallach will ensure that these teams are in place and on track. Each management team will set the 5-year agenda for this technical theme, adjust it as necessary, coordinate meetings and teleconferences of all involved students and senior personnel, and produce periodic progress reports. The management team will ensure that PhD students working on this theme get input from all relevant project participants and that instructional material is provided to all participants teaching relevant courses.

For example, the Design for Verification theme, managed by Professors Dill and Wagner, has the following tentative 5-year plan:

- **Year 1.** Problem formulation and initial development of basic techniques.
- **Year 2.** Development of techniques. Design and prototype implementation. Discussion with user communities, industry partners, and policy organizations to formulate challenges and requirements in specific voting scenarios. Examination of applicability of ideas to other security domains. Preliminary white papers and lecture notes.
- **Year 3.** Alpha release of an experimental platform for proof of concept. Further development of techniques. Design and implementation of design for verification system development tools.
- **Year 4.** Beta release and large-scale use of experimental platform. Further development of design for verification tools, and start of work with user communities, industry partners, and policy organizations.
- **Year 5.** Final release and documentation of experimental platform. Completion of technology-transfer activities with software-industry partners if appropriate. Second round of white papers and lecture notes.

Current plans for other first-level management teams can be found in the personnel matrix below.

⁸See <http://www.openvotingconsortium.org/>.

5.1 Team Coordination and Communication

The principal investigators will meet twice yearly, with an open workshop (see Section 3.2) and a private retreat. The workshop will rotate among our member institutions, and the retreat will be held each year at Johns Hopkins in conjunction with a meeting of the Advisory Board, invited guests, and selected graduate students. We will also use part of our domestic-travel budget for face-to-face meetings of subgroups that are working together on particular technical themes; these trips can be piggybacked with other activities requiring travel, for example, guest lectures in classes. Finally, we will hold periodic teleconferences and make extensive use of email. These are tried-and-true coordination and communication mechanisms that many of us have already used successfully in multi-institutional projects involving co-advising of graduate students, for example, in the PORTIA project [7], in which one of us (Boneh) is a PI.

To further facilitate collaboration with the computer industry, user communities, and the public policy community, PhD students and faculty members will do internships and sabbaticals at our partner institutions.

5.2 Personnel Matrix

In Table 1, L stands for “leader” (or first-level manager) and P for “participant.” In the top half, the first column lists the research leaders for the Center. Each principal researcher will lead one effort and participate in two or three others, and each area will have two leaders who will coordinate the research in that topic area. In addition to the Center participants, the Center will work with a group of unfunded *affiliates*, shown in the bottom half; these were strategically chosen to collaborate with the Center in areas of expertise that match the research thrusts of the ACCURATE Center. For example, Chris Edley, the Dean of the U.C. Berkeley Law School is known for his studies of the interplay between technology and policy. Several of us are participating in VSPR (Voting System Performance Rating, vspr.org), and one of us (Dill) is the founder of the Verified Voting Foundation, an affiliate of the Center. In choosing team members, we have paid careful attention to the professional achievements of the people involved, including prior collaborations. For example, Rubin and Wallach have worked together on source code analysis of electronic voting machines, and Dill and Jones have worked together on system level and verifiability issues. This assignment of roles will be revised as needed, with any changes explained in our NSF progress reports.

5.3 Evaluation Plan

The primary focus of the ACCURATE Center, through a multidisciplinary approach including technological, legal, and policy aspects, is to catalyze change in the way America votes, ensuring the continued future of fair elections in the nation. While available resources will constrain the full multidisciplinary approach to the election domain, we fully expect that technical research results will also apply to other application areas. The ACCURATE Center will achieve success by establishing itself as a leader in computer security for key areas such as election systems and electronic auctions.

Clearly, some traditional metrics for evaluating progress are applicable. One expected output of the Center will be new technical approaches to the computer security problems found in election systems. These results will be documented in papers submitted to top-tier computer security conferences (e.g., IEEE Symposium on Security and Privacy, ACM Computer and Communications Security, USENIX Security Symposium, ISOC Symposium on Network and Distributed Systems Security) as well as peer-reviewed academic journals (e.g., the *ACM Transactions on Information and Systems Security*). Traditional measures such as publication and citation counts can track the scientific progress of the Center and its researchers.

While publication and citation counts can track the ACCURATE Center’s technical progress, alone they will not measure whether the Center has achieved leadership outside the academic computer security community. One obvious approach is to measure the Center’s impact on legal changes supporting the switch

	System-level issues	Role of cryptography	Design for verification	Relating policy to technology	Usability and accessibility
Michael Byrne	P		P		L
Dan Boneh	P	L	P		
Drew Dean	P	L	P		P
David Dill	P		L	P	
Doug Jones	L	P	P		P
Deirdre Mulligan				L	P
Peter Neumann	P	P	P	L	
Avi Rubin	L	P	P		P
David Wagner	P	P	L		P
Dan Wallach	P	P	P		L
Kim Alexander				P	P
Josh Benaloh		P	P		
David Chaum		P		P	
Cindy Cohn				P	
Chris Edley				P	
David Jefferson	P		P		
Whitney Quesenbery					P
Verified Voting				P	P

Table 1: Personnel Matrix

to electronic voting, but this may be too narrow: good policy analysis can be held hostage to political forces. More meaningful metrics in the legal and policy arenas against which to evaluate ACCURATE include:

- Center participants testifying before government (e.g., Congress, state legislatures, county boards of supervisors, and municipal councils) and administrative bodies
- Center participants briefing studies in relevant areas such as those performed by the National Research Council
- Center participation (either as individuals or as an institution) in standards bodies (e.g., IEEE)
- Articles and editorials discussing the work of the Center
- Publication of scholarly articles in appropriate (nontechnical) journals and conferences

Taken together with the traditional scientific metrics, these metrics should measure the success and leadership of the Center.

The true technical success of the Center will be measured by having its technologies available and widely used in next-generation election systems, either from today's leading vendors, or from new entrants to the market. Unfortunately, this will most likely be impossible to measure before the end of the initial 5 years; ACCURATE-developed technologies will most likely become viable for commercial development only in the fourth or fifth year of the Center. One can then expect a further 2- to 3-year delay of commercialization when product cycles are taken into account. Technology transition to other application areas, such as online auctions or robust networking systems, may be more rapid, due to reduced certification requirements.

6 References

- [1] *Bush v. Gore*, 531 U.S. 98 (Dec. 2000).
- [2] *Common Cause v. Jones*, 213 F. Supp. 2d 1106 (Aug. 2001).
- [3] *Weber v. Shelley*, 347 F.3d 1101 (Oct. 2003).
- [4] *NAACP V. Harris*, No. 01-Civ-120 (S.D. Fla. Filed January 10, 2001).
- [5] *Black v. McGuffage*, 209 F. Supp. 2d 889 (N.D. Ill. 2002).
- [6] *Southwest Voter Registration Education Project V Shelley*, 344 F.3d 914 (9th Cir. 2003 (*En Banc*)).
- [7] Portia project. <http://crypto.stanford.edu/PORTIA>.
- [8] Known vendors of computerized vote tabulation systems. FEC Web site, Dec. 2000. <http://www.fec.gov/pages/vendors12-00.htm>.
- [9] Task Force on The Federal Election System, National Commission on Federal Election Reform, to Assure Pride and Confidence in the Electoral Process, 2001.
- [10] Help America Vote Act, 2002. United States Public Law 107-252.
- [11] Defense in depth. In *Security Recommendation Guides*. National Security Agency, 2003.
- [12] W. Aiello, J. Ioannidis, and P. McDaniel. Origin authentication in inter-domain routing. In *ACM Conference on Computer and Communications Security*, 2003.
- [13] S. D. Albright. *The American Ballot*. American Council on Public Affairs, Washington, D.C., 1942.
- [14] R. Anderson. Cryptography and competition policy — issues with ‘trusted computing’. In *Proceedings of the 2nd Workshop on Economics and Information Security*, College Park, MD, May 2003.
- [15] S. Ansolabehere and C. Stewart. Residual votes attributable to technology. *Journal of Politics*, 67(2), May 2005. <http://journalofpolitics.org/Contents/Vol67/arts672/stewart.pdf>.
- [16] Association for Computing Machinery. *ACM’s Position on Electronic Voting*. <http://www.acm.org/usacm/weblog/index.php?p=73>.
- [17] J. Bannet, D. W. Price, A. Rudys, J. Singer, and D. S. Wallach. Hack-a-Vote: Demonstrating security issues with electronic voting systems. *IEEE Security & Privacy Magazine*, 2(1):32–37, Mar. 2004.
- [18] M. Barnett, K. R. M. Leino, and W. Schulte. The Spec# programming system: An overview. In *Construction and Analysis of Safe, Secure and Interoperable Smart devices (CASSIS 2005)*. Springer-Verlag, 2005. To appear.
- [19] B. B. Bederson, B. Lee, R. M. Sherman, P. S. Herrnson, and R. G. Niemi. Electronic voting system usability issues. In *Human Factors in Computing Systems: Proceedings of CHI 2003*, pages 145–152. ACM, New York, 2003.
- [20] H. E. Brady, J. Buchler, M. Jarvis, and J. McNulty. *Counting All the Votes: The Performance of Voting Technology in the United States*. Department of Political Science, Survey Research Center, and Institute of Governmental Studies, University of California, Berkeley, Berkeley, CA, Sept. 2001. <http://macht.arts.cornell.edu/wrml/countingallthevotes.pdf>.

- [21] H. E. Brady, G.-U. Charles, B. Highton, M. Kropf, W. R. Mebane, and M. Traugott. *Interim Report on Alleged Irregularities in the United States Presidential Election of 2 November 2004*. National Research Commission on Elections and Voting, Dec. 2004. <http://election04.ssrc.org/research/InterimReport122204.pdf>.
- [22] S. Bruck, D. Jefferson, and R. L. Rivest. A modular voting architecture (“Frogs”). In *Workshop on Trustworthy Elections*, Tomales Bay, CA, August 2001. <http://www.vote.caltech.edu/wote01/pdfs/amva.pdf>.
- [23] L. Burdy, Y. Cheon, D. R. Cok, M. D. Ernst, J. R. Kiniry, G. T. Leavens, K. R. M. Leino, and E. Poll. An overview of JML tools and applications. To appear in *International Journal on Software Tools for Technology Transfer*.
- [24] M. D. Byrne. Cognitive architecture. In J. A. Jacko and A. Sears, editors, *The human-computer interaction handbook: Fundamentals, evolving technologies and emerging applications*, pages 97–117. Lawrence Erlbaum, Mahwah, NJ, 2003.
- [25] California Internet Voting Task Force. *A report on the feasibility of Internet voting*, Jan. 2000. <http://www.ss.ca.gov/executive/ivote/>.
- [26] CalTech-MIT/Voting Technology Project. *Voting: What Is; What Could Be*, July 2001. <http://www.vote.caltech.edu/Reports/>.
- [27] Caltech/MIT Voting Technology Project. *Voting Machines and the Underestimate of the Bush Vote*, Nov. 2004. <http://www.vote.caltech.edu/Reports/VotingMachines3.pdf>.
- [28] M. Castro, P. Druschel, A. Ganesh, A. Rowstron, and D. S. Wallach. Security for structured peer-to-peer overlay networks. In *Proc. OSDI’02*, Boston, MA, Dec. 2002.
- [29] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Comm. ACM*, 24(2):84–88, 1981.
- [30] D. Chaum. Secret-ballot receipts: True voter-verifiable elections. *IEEE Security & Privacy Magazine*, 2(1):38–47, Mar. 2004.
- [31] D. Chaum, A. Fiat, and M. Naor. Untraceable electronic cash. In *Advances in Cryptology - CRYPTO ’88, 8th Annual International Cryptology Conference*, pages 319–327, Santa Barbara, CA, Aug. 1988.
- [32] J. P. Chin, V. A. Diehl, and K. L. Norman. Development of an instrument measuring user satisfaction of the human-computer interface. In *Proceedings of SIGCHI ’88*, pages 213–218. ACM, New York, 1988.
- [33] Compuware Corporation. *Direct Recording Electronic (DRE) Technical Security Assessment Report*, Nov. 2003. <http://www.sos.state.oh.us/sos/hava/files/compuware.pdf>.
- [34] L. de Moura, S. Owre, H. Ruess, J. Rushby, N. Shankar, M. Sorea, and A. Tiwari. SAL 2. In *16th International Conference on Computer Aided Verification*, Boston, MA, July 2004. Available at <http://www.csl.sri.com/~rushby/abstracts/sal-tool>.
- [35] Diebold Election Systems. *Checks and Balances in Elections Equipment and Procedures Prevent Alleged Fraud Scenarios*, July 2003. <http://www2.diebold.com/checksandbalances.pdf>.

- [36] D. L. Dill, A. J. Drexler, A. J. Hu, and C. H. Yang. Protocol verification as a hardware design aid. In *1992 IEEE International Conference on Computer Design: VLSI in Computers and Processors*, pages 522–525. IEEE Computer Society, 1992. Cambridge, MA, October 11-14.
- [37] Election Protection Coalition. *Reports from the Election Incident Reporting System*. <http://voteprotect.org>.
- [38] electionline.org. *Election Reform Briefing: The 2004 Election*, Dec. 2004. <http://www.electionline.org/site/docs/pdf/ERIP%20Brief9%20Final.pdf>.
- [39] Federal Election Commission. *Performance and Test Standards for Punchcard, Marksense, and Direct Recording Electronic Voting Systems*, Apr. 1990.
- [40] Federal Election Commission. *Voting System Standards*, Apr. 2002. <http://www.fec.gov/pages/vssfinal/vss.html>.
- [41] A. Feldmann, O. Maennel, Z. Mao, A. Berger, and B. Maggs. Locating Internet routing instabilities. In *SIGCOMM*, 2004.
- [42] E. A. Fischer. Voting technologies in the United States: Overview and issues for Congress. Technical Report RL30733, Congressional Research Service, Mar. 2001. <http://www.ncseonline.org/NLE/CRSreports/Risk/rsk-55.cfm>.
- [43] E. A. Fischer. Election reform and electronic voting systems (DREs): Analysis of security issues. Technical report, Nov. 2003. <http://www.epic.org/privacy/voting/crsreport.pdf>.
- [44] M. S. Frankel, T. Jacobovits, and A. Kroepsch. *Making Each Vote Count: A Research Agenda for Electronic Voting*. American Association for the Advancement of Science, Oct. 2004. <http://www.aaas.org/spp/sfml/evoting/report2.pdf>.
- [45] C. Gilmore, D. Kormann, and A. D. Rubin. Secure remote access to an internal web server. *Proceedings of the ISOC Symposium on Network and Distributed System Security*, pages 23–34, February 1999.
- [46] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin. Working around BGP: An incremental approach to improving security and accuracy of interdomain routing. In *Proceedings of the IEEE Network and Distributed System Security Symposium*, February 2003.
- [47] W. D. Gray, B. E. John, and M. E. Atwood. Project Ernestine: A validation of GOMS for prediction and explanation of real-world task performance. *Human-Computer Interaction*, 8:237–309, 1993.
- [48] B. Harris. *Black Box Voting: Vote Tampering in the 21st Century*. Elon House/Plan Nine, July 2003.
- [49] R. L. Hasen. Bush v. Gore and the future of equal protection law in elections. 29 Fla. St. U. L. Rev. 377, 2001.
- [50] R. L. Hasen. Time to fix election system. Law.com Commentary, Nov. 2004. <http://www.law.com/jsp/article.jsp?id=1099217144168>.
- [51] M. C. Herron and J. S. Sekhon. Overvoting and representation: An examination of overvoted presidential ballots in Broward and Miami-Dade counties. *Electoral Studies*, 22(1):21–47, Mar. 2003.
- [52] G. J. Holzmann. The model checker SPIN. *Software Engineering*, 23(5):279–295, 1997.

- [53] House Committee on Science, 107th Congress. *Improving Voting Technology*, May 2001. USGPO Serial No. 107-20.
- [54] House Judiciary Committee Democratic Staff. *Preserving Democracy: What Went Wrong in Ohio*, Jan. 2005. http://election04.ssrc.org/research/preserving_democracy.pdf.
- [55] Y. Hu, A. Perrig, and M. Sirbu. SPV: Secure Path Vector Routing for Securing BGP. In *SIGCOMM*, 2004.
- [56] InfoSENTRY. *DRE Security Assessment, Volume 1, Computerized Voting Systems, Summary of Findings and Recommendations*, Nov. 2003. <http://www.sos.state.oh.us/sos/hava/files/InfoSentry1.pdf>.
- [57] B. E. John and D. E. Kieras. Using GOMS for user interface design and evaluation: Which technique? *ACM Transactions on Computer-Human Interaction*, 3:287–319, 1996.
- [58] D. W. Jones. Auditing elections. *Communications of the ACM*, 47(10):46–50, Oct. 2004. <http://www.cs.uiowa.edu/~jones/voting/cacm2004.shtml>.
- [59] D. W. Jones. *The European 2004 Draft E-Voting Standard - Some critical comments*, Oct. 2004. <http://www.cs.uiowa.edu/~jones/voting/coe2004.shtml>.
- [60] D. W. Jones. *Minimizing the Trusted Base*, Dec. 2004. Presentation for the CSTB Framework for Understanding Electronic Voting, <http://www.cs.uiowa.edu/~jones/voting/nas-cstb2004a.shtml>.
- [61] S. Kent, C. Lynn, J. Mikkelsen, and K. Seo. Secure Border Gateway Protocol (S-BGP) — Real World Performance and Deployment Issues. In *Proceedings of Network and Distributed Systems Security 2000*, February 2000.
- [62] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (Secure-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4):582–592, April 2000.
- [63] T. Kohno, A. Stubblefield, A. D. Rubin, and D. S. Wallach. Analysis of an electronic voting system. In *Proceedings of the 2004 IEEE Computer Society Symposium on Research in Security and Privacy*, May 2004.
- [64] S. J. Laskowski, M. Autry, J. Cugini, W. Killam, and J. Yen. Improving the usability and accessibility of voting systems and products. Technical Report 500-256, National Institute of Standards and Technology, 2004.
- [65] L. Lessig. *Code and Other Laws of Cyberspace*. Basic Books, New York, 2000.
- [66] A. J. Lichtman. Report on the racial impact of the rejection of ballots cast in the 2000 presidential election in the state of Florida. Appendix To United States Commission on Civil Rights, Voting Irregularities During The 2000 Election, 2001.
- [67] R. Mahajan, D. Wetherall, and T. Anderson. Understanding BGP Misconfiguration. In *Proceedings of ACM SIGCOMM '02*, pages 3–16. ACM, September 2002.
- [68] R. Mahajan, D. Wetherall, and T. Anderson. Understanding BGP misconfiguration. In *SIGCOMM*, 2002.

- [69] J. Marchesini, S. W. Smith, O. Wild, and R. MacDonald. Experimenting with TCPA/TCG hardware, or: How I learned to stop worrying and love the bear. Technical Report TR2003-476, Department of Computer Science, Dartmouth College, Dec. 2003.
- [70] W. R. Mebane. The wrong man is President! Overvotes in the 2000 Presidential election in Florida. *Perspectives on Politics*, 2(3):525–535, Sept. 2004.
- [71] W. R. Mebane, J. S. Sekhon, and J. Wand. *Detecting and Correcting Election Irregularities*, Oct. 2003. Preliminary draft, <http://wand.stanford.edu/research/detecting.pdf>.
- [72] R. Mercuri. *Electronic Vote Tabulation Checks and Balances*. PhD thesis, Department of Computer and Information Systems, University of Pennsylvania, 2001. <http://www.notablesoftware.com/evote.html>.
- [73] R. Mercuri. A better ballot box: New electronic voting systems pose risks as well as solutions. *IEEE Spectrum*, 39(10):46–50, October 2002. <http://www.spectrum.ieee.org/WEBONLY/publicfeature/oct02/evot.html>.
- [74] R. Mercuri and L. J. Camp. The code of elections. *Communications of the ACM*, 47(10):52–57, October 2004. *Inside Risks* column.
- [75] J. Milarsky. Touch screens more likely to be flawed, analysis finds, January 16 2005. http://www.sun-sentinel.com/news/local/broward/sfl-cundervotes16jan16_1,5678995.story.
- [76] D. K. Mulligan and J. L. Hall. *Preliminary Analysis Of E-Voting Problems Highlights Need For Heightened Standards And Testing*. National Research Council, Committee on Electronic Voting, Dec. 2004. http://www7.nationalacademies.org/cstb/project_evoting_mulligan.pdf.
- [77] S. J. Mulroy. Lemonade from lemons: Can advocates convert Bush v. Gore into a vehicle for reform? 9 Geo J. Poverty Law And Pol’y 357, 2002.
- [78] National Science Foundation. *Report on the National Workshop on Internet Voting: Issues and Research Agenda*, Mar. 2001. <http://news.findlaw.com/cnn/docs/voting/nsfe-voterprt.pdf>.
- [79] G. Necula and P. Lee. Safe, untrusted agents using proof-carrying code. In *Mobile Agents and Security*, volume 1419 of *Lecture Notes in Computer Science*, pages 61–91. Springer-Verlag, 1998.
- [80] A. Neff. A verifiable secret shuffle and its application to e-voting. In *Proceedings of ACM CCS’01*, pages 116–125. ACM Press, 2001.
- [81] P. G. Neumann. Risks in computerized elections. *Communications of the ACM*, 33(11):170, November 1990. *Inside Risks* column.
- [82] P. G. Neumann. Security criteria for electronic voting. In *Proceedings of the Sixteenth National Computer Security Conference*, pages 478–482, Baltimore, Maryland, 20–23 September 1993.
- [83] P. G. Neumann. Illustrative risks to the public in the use of computer systems and related technology, index to RISKS cases. Technical report, Computer Science Laboratory, SRI International, Menlo Park, California, 2004. <http://www.csl.sri.com/neumann/illustrative.html>, click on “Election Problems”.

- [84] P. G. Neumann. Principled assuredly trustworthy composable architectures. Technical report, Computer Science Laboratory, SRI International, Menlo Park, California, December 2004. Final report, SRI Project 11459, www.csl.sri.com/neumann/chats4.html; also `chats4.ps` and `chats4.pdf`.
- [85] T.-W. J. Ngan, D. S. Wallach, and P. Druschel. Enforcing fair sharing of peer-to-peer resources. In *Proc. IPTPS'03*, Berkeley, CA, Feb. 2003.
- [86] T.-W. J. Ngan, D. S. Wallach, and P. Druschel. Incentives-compatible peer-to-peer multicast. In *2nd Workshop on the Economics of Peer-to-Peer Systems*, Cambridge, MA, June 2004.
- [87] J. Nielsen. Finding usability problems through heuristic evaluation. In *Human factors in computing systems: Proceedings of CHI 92*, pages 373–380. ACM, New York, 1992.
- [88] S. Owre, J. Rushby, N. Shankar, and F. von Henke. Formal verification for fault-tolerant architectures: Prolegomena to the design of PVS. *IEEE Transactions on Software Engineering*, 21(2):107–125, Feb. 1995.
- [89] R. Pastor. Improving the U.S. electoral system: Lessons from Canada and Mexico. 3 *Election L.J.* 584, 2004.
- [90] N. Paul, D. Evans, A. D. Rubin, and D. S. Wallach. Authentication for remote voting. In *Workshop on Human-Computer Interaction and Security Systems*, Fort Lauderdale, FL, Apr. 2003.
- [91] D. Pei, X. Zhao, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang. Improving BGP convergence through consistency assertions. *IEEE INFOCOM*, June 2002.
- [92] P. Pirolli and S. Card. Information foraging. *Psychological Review*, 106:643–675, 1999.
- [93] P. G. Polson, C. Lewis, J. Reiman, and C. Wharton. Cognitive walkthroughs: A method for theory-based evaluation of user interfaces. *International Journal of Man-machine Studies*, 36:741–773, 1992.
- [94] W. Quesenbery. Voting and usability: Lessons from the 2000 presidential election, 2001.
- [95] RABA Technologies. *Trusted Agent Report—Diebold AccuVote-TS Voting System*, Jan. 2004. http://www.raba.com/text/press/TA_Report_AccuVote.pdf.
- [96] M. K. Reiter and A. D. Rubin. Crowds: Anonymity for Web transactions. *ACM Transactions on Information System Security*, 1(1), April 1998.
- [97] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang. BGP routing stability of popular destinations. In *Internet Measurement Workshop*, 2002.
- [98] S. K. Roth. Disenfranchised by design: voting systems and the election process. *Information Design Journal*, 9(1), 1998.
- [99] P. M. Schwartz. Voting technology and democracy. 77 *NYU L. Rev.* 625, Sept. 2002.
- [100] Science Applications International Corporation. *Risk Assessment Report: Diebold AccuVote-TS Voting System and Processes*, Sept. 2003. <http://www.dbm.maryland.gov/SBE>.
- [101] J. S. Sekhon. *The 2004 Florida Optical Voting Machine Controversy: A Causal Analysis Using Matching*. Department of Government, Harvard University, Nov. 2004. Preliminary draft, <http://jsekhon.fas.harvard.edu/papers/SekhonOpticalMatch.pdf>.

- [102] T. Selker. *Processes can Improve Electronic Voting: A Case Study of an Election*. Caltech/MIT Voting Technology Project, Oct. 2004. http://www.vote.caltech.edu/Reports/vtp_wp17.pdf.
- [103] J. Shapiro and J. Venderburgh. Access and integrity control in a public-access, high-assurance configuration management system. In *Proceedings of the 11th USENIX Security Symposium*, pages 109–120, San Francisco, CA, Aug. 2002.
- [104] J. D. Simon and R. P. Baiman. *The 2004 Presidential Election: Who Won the Popular Vote? An Examination of the Comparative Validity of Exit Poll and Vote Count Data*. Free Press, Jan. 2005. http://freepress.org/images/departments/PopularVotePaper181_1.pdf.
- [105] State of California Secretary of State Voting Systems Panel. *Minutes*, Dec. 2003. http://www.ss.ca.gov/elections/vsp_min_121603.pdf.
- [106] A. Stump, C. Barrett, and D. Dill. CVC: a Cooperating Validity Checker. In *14th International Conference on Computer-Aided Verification*, 2002.
- [107] D. Tokaji. The 2008 election: Could it be a repeat of 2000? Findlaw’s Legal Commentary, Nov. 2004. http://writ.news.findlaw.com/commentary/20041130_tokaji.html.
- [108] D. Tokaji. The paperless chase: Electronic voting and democratic values. Public Law and Legal Theory Working Paper Series, Sept. 2004. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=594444.
- [109] VerifiedVoting.org. *Resolution on Electronic Voting*. <http://verifiedvoting.org/article.php?id=5028>.
- [110] M. Waldman, A. D. Rubin, and L. F. Cranor. Publius: A robust, tamper-evident, censorship-resistant web publishing system. *USENIX Security Conference IX*, pages 59–72, August 2000.
- [111] J. Wand. *Evaluating the Impact of Voting Technology on the Tabulation of Voter Preferences: The 2004 Presidential Election in Florida*. Department of Political Science, Stanford University, Nov. 2004. <http://wand.stanford.edu/elections/us/FL2004/WandFlorida2004.pdf>.
- [112] J. Wand, K. Shotts, J. Sekhon, W. Mebane, M. Herron, and H. Brady. The Butterfly did it: The aberrant vote for Buchanan in Palm Beach County, Florida. *American Political Science Review*, 95(4):793–810, Dec. 2001. <http://wand.stanford.edu/research/apsr2001.pdf>.
- [113] F. Wang and L. Gao. Inferring and characterizing Internet routing policies. In *Internet Measurement Conference*, 2003.
- [114] D. R. Wold. The HAVA requirement for a voter verified paper record. A whitepaper authored by former FEC Chairman Wold, July 2003. http://www.verifiedvoting.org/downloads/resources/documents/HAVA_Requirement_for_VVP_Record.pdf.
- [115] Wyle Laboratories. *Qualification Testing of the I-Mark Electronic Ballot Station, Report No 45450-01*, Sept. 1996. Confidential report, content disclosed here was discussed in open meetings of the Iowa Board of Examiners for Voting Machines and Electronic Voting Systems.
- [116] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang. An analysis of BGP Multiple Origin AS (MOAS) conflicts. In *Internet Measurement Workshop*, 2001.