CRITICAL NATIONAL NEED IDEA

Critical National Need Idea Title: High-Assurance Voting Systems

Submitting Organization:
Galois, Inc.
421 SW 6th Ave., Suite 300
Portland, OR 97204

Contact:
Jodee LeRoux
Direct phone: (503) 808-7209
Fax: 503-350-0833
e-mail: jodee@galois.com

Key words: voting, software, formal methods

High-Assurance Voting Systems

Aaron Tomb — Galois, Inc.

January 15, 2009

1 Introduction

Accurate and trustworthy elections are an essential component of an effective democracy. They provide key support for the *most* essential component: public trust in the process. When either malicious intent or accidental error may affect the outcome of an election, the foundation of our system of government is at risk.

Until recently, a practical method to efficiently achieve extremely high confidence in the correctness of election results has not existed. The most trustworthy process in current practice, a manual count by multiple independent parties, is time-consuming and expensive. The most efficient process in current practice, where voters enter indicate their choices directly on a computer that automatically keeps a running tally, has been shown to be extremely insecure. Researchers have demonstrated large numbers of serious weaknesses in the security mechanisms of a wide variety of touch-screen voting machines from the most popular manufacturers [4, 17].

1.1 A Critical Need

Public demand for immediate results, pressure from industry, and the need to lower costs are pushing us inexorably toward electronic voting systems. The documented weaknesses of existing attempts and the importance of accurate voting necessitate new approach to the design of voting machinery.

Manual voting technology — especially when aided by appropriate legislation — is simple and tangible and has a high degree of transparency and repeatability. The average citizen can understand how hand counting works, help in the process of a recount, if needed, and see the justification for the final tally.

On the other hand, electronic voting technology is not only complex and opaque, but extremely difficult to design correctly. While its efficiency benefits are tempting, its susceptibility to both malicious attacks and accidental flaws deserves serious attention. Nevertheless, it is likely inevitable that electronic systems will be increasingly involved in at least some segments of the voting process. It is therefore critical that we put careful attention into the design of those systems in order to prevent a decline in the integrity of our elections.

The trend toward electronic voting systems presents a significant societal challenge: failure to address the shortcomings of current electronic voting methods threatens to undermine public trust in the government process. In addition, the pressures affecting both industry and academia all but guarantee that neither group will come up with a solution on its own. A collaboration between industry and academia is essential for an effective solution.

1.2 A Transformational Result

A *disciplined* use of modern computer technology can result in a voting process with the dual benefits of highly *trustworthy* and highly *efficient* tabulation. These benefits have never existed simultaneously in the history of large-scale elections but are possible with the appropriate application of technology and careful design, including a rigorous approach to software development.

The research and practice of high-assurance software development has matured to a stage where we can feasibly create formal proofs demonstrating the correct implementation of the critical components in a software system. Though formal methods have reached the point of practical applicability, creating software amenable to formal analysis and constructing the proofs of its correctness are still large research endeavors, with many unknowns and significant risk. Only a few examples of formally verified software exist.

The rewards of a successful outcome, however, are substantial enough to offset the associated risk. A successful outcome would have the dual benefits of 1. providing a indisputable basis for trust in the accuracy of election results and 2. providing a clear demonstration of the feasibility of developing highly trustworthy software. The magnitude of the problem and the importance of overcoming the internal pressures and habits of both industry and academia suggest that such an outcome is unlikely to occur without governmental support. No organization is more appropriate to provide this support than the National Institute of Standards and Technology.

The next four sections of this white paper describe the desirable guarantees of a trustworthy voting process, how the state of the art in voting technology falls short of providing these guarantees, a brief description of the use of *formal methods* in software development, and some characteristics of an effective voting solution. The final section explicitly describes how this critical need and its potentials solutions align with the goals and requirements of TIP funding.

2 Desired Properties of Voting

A practical and trustworthy voting system must accurately count votes and preserve voter privacy, as required by law in many locations. We propose an additional desirable property: the ability of voters to check the validity of votes independently of election authorities. Here we describe each of these properties in more detail.

2.1 Accuracy

The most obviously desirable property of a voting system is *accuracy*: the tabulated results correctly reflect the choices indicated by individual voters. Both accident and malice can compromise accuracy. Accuracy requires the final tally to include every valid vote cast and to exclude every invalid one.

In a manual count, a tabulator can misread a ballot or intentionally misrecord it. Fortunately, multiple independent parties can examine each ballot, if desired, and reading a ballot does not require any special expertise (though ambiguous markings can occur).

A flaw in an electronic voting machine might cause it to record an incorrect value. Because the voter cannot see (much less interpret) the accumulation of charge in the computer's memory or the alignment of magnetic particles on a hard disk, this error, whether malicious or accidental, may easily go undetected.

2.2 Privacy

Another important property is that of *privacy*: the identity of an individual voter can in no way be linked to the content of a specific vote. Privacy helps prevent vote buying and voter intimidation: if a voter cannot prove, even voluntarily, what vote he or she cast, a malicious party intending to influence the election cannot be confident that payment or intimidation had the indented effect, and the voter can feel comfortable voting freely without fear of punishment.

Legislation and process ensure the privacy of paper ballots. Typically, all paper ballots are visually indistinguishable and can be rejected if a voter has made any distinguishing marks. Additionally, voters put paper ballots into secrecy sleeves for transportation between a polling station and a ballot box or through the mail.

The problems that can cause electronic voting machines to violate accuracy can also cause them to violate privacy. Just as a voter cannot be certain that a computer has correctly recorded a vote, a voter cannot be certain that a computer has *not* recorded his or her identity. To prevent malicious insertion of invalid votes, a voting machine should check the identity of each voter and allow each voter to cast only a single vote. However, it can be very difficult to guarantee that the machine does not record this identity information.

2.3 Verifiability

A voting system is *verifiable* if strong evidence or proof exists that the final count it produces is accurate. It is *voter-verifiable* if an individual voter (without any special authority) can determine whether an individual ballot was recorded correctly. No voting system in widespread use is voter-verifiable. The use of computers in some stages of the voting system can make voter verifiability possible, but only if we can trust that the overall voting process has certain mathematical properties and that electronic systems involved always function correctly.

3 Voting Technology

Current elections in the United States are conducted using a wide range of technology, some electronic, and some manual, including pen-marked paper ballots, punch cards, optical scanners, and Direct Recording Electronic (DRE) voting machines.

3.1 Optically-Scanned Ballots

In a system with optically scanned ballots, voters use pens to make dark marks on paper ballots to indicate their choices. Electronic scanners can then read these ballots, recognize the presence or absence of marks, and compute a tally from this information. Scanners help improve the efficiency of vote counting without removing a paper record usable for a manual recount.

3.2 Punch Card Systems

Punch cards are sometimes used as an alternative form of electronically-readable paper ballot. Instead of marking their ballots with pens, voters use punches to create holes in their ballots that indicate their choices. Electronic counters can then recognize the presence or absence of holes in each ballot and automatically create a tally from this information. Most states have preferred optically-scanned ballots to punch card systems in recent years, perhaps due in part to the controversy in Florida over "hanging chads" in the 2000 U.S. Presidential election [13].

3.3 Direct Recording Electronic Voting Systems

A Direct Recording Electronic (DRE) voting machine is a computer with a special user interface (perhaps a touch screen, dial, or set of buttons) that voters use to select their choices. The choices are then directly recorded to some electronic storage medium.

While DRE systems are very efficient — the final tally is typically available as soon as polls close — they have a serious weakness. In the process just described, the voter has absolutely no evidence that the information recorded to the storage device has any connection to the intended vote, or that personally identifying information is not associated with the vote.

3.4 Paper Trails

To address the key weakness of DRE voting machines, many states have added a requirement for a Voter-Verified Paper Record (VVPR). In this approach, as soon as a voter finishes entering choices into a DRE terminal, the machine prints out a receipt indicating the choices it will (supposedly) record. The voter can then either indicate that the receipt shows the intended choices, in which case it will be deposited in a secure compartment, or indicate that it contains a mistake, in which case it will be marked to show that it should be ignored. If the election is contested, auditors can manually examine the paper receipts, which voters have approved.

The use of a VVPR is an important step that improves the auditability of elections. However, the paper records come into play only when the results of an election are called into question. How do election officials know when to perform a manual audit? A small number of changed votes in a large number of precincts may be enough to change the outcome of an election without raising suspicions.

The use of the term *voter-verifiable* in this white paper refers to a property stronger than that guaranteed by a VVPR, or indeed by optical-scan or punch-card paper ballots. Our use of the term implies that an individual voter can check that an individual vote was correctly recorded and correctly included in the final tally. Errors in counting should be detectable *without* an expensive and time-consuming recount. A recount should be necessary only *if* errors are detected, and *if* the number of errors is high enough to affect the outcome.

4 Formal Methods and Software Verification

Researchers have devised a number of abstract protocols that can be used to conduct elections with a guarantee that certain properties, such as the previously-mentioned integrity, privacy, and voter verifiability, will necessarily hold, given certain basic assumptions. Many of these protocols have been developed with an aim to limit the number of components that must be trusted for the guarantees to hold, and particularly to limit the number of software components that must be implemented correctly.

However, a few software components still must be implemented flawlessly to obtain desirable guarantees. In particular, correct implementations of cryptographic algorithms are key to the success of most protocols, and many also require software that can reliably protect the secrecy of certain information.

Developing software that correctly performs the operation intended by its developers is notoriously difficult. Many techniques exist to address this problem and help increase confidence in software correctness. These techniques range from simple sanity checks demonstrating that typical conditions are handled correctly to complex analyses providing strong guarantees of conformance to specifications.

At the simple extreme, most software is *tested* to some degree. Testing involves running software as you would in practice, providing both expected and erroneous inputs, and observing the resulting outputs or behavior. If testers can examine a piece of software in response to *every* possible input and check that it operates as intended, they can confidently assert that it will always behave properly in the future.

However, most realistic software can accept an unlimited number of inputs, so it is impossible to test the real system under every possible condition. Instead, testers attempt to determine a certain set of characteristic inputs that test all of the essentially different modes of operation. This approach can ensure only a limited degree of correctness unless testers are willing to dedicate incredibly large amounts of time. In practice, the amount of testing that occurs is far too small to provide any significant guarantees.

It may seem at first that a voting system would have a finite number of possible inputs, and therefore be exhaustively testable. Though the number of explicit choices a voter may input is indeed finite, the system becomes much more complex when other factors are taken into account, such as different operating conditions, communication between systems, and error handling. Therefore, in practice, even a system as apparently simple as a voting machine has an infinite number of possible states.

An alternate approach to increasing confidence in software correctness is the use of *formal methods*. This approach gets its name from its use of constructing a formal mathematical model of a piece of software and using well-understood laws governing the model to construct a proof that the software has some desired property, even when the number of specific cases in the model is infinite. For instance, consider the statement that for every number n, the number $\frac{n}{2}$ is closer to zero than n. There are infinitely many numbers, but this statement can be proven to be true without examining every case individually.

In some cases, programmers create the source code of the program first, and extract the mathematical model from the concrete implementation. In other cases, the model comes first, and the concrete implementation is derived from it. In either case, assuming the connection between the concrete code and the abstract model preserves meaning sufficiently, proofs about the model apply to the running program.

Whichever approach is used, it can take significant work to create software that is amenable to formal modeling, and it requires substantial mathematical expertise. This effort and expertise are greater than that which are typically expended testing software but less than would be required to test a piece of software under every possible condition.

Ultimately, both extensive testing and formal methods are effective ways to determine the correctness of software. Ideally, the development of critical software would include significant application of both techniques. Unfortunately, taking either approach far enough to make confident guarantees is time consuming and, therefore, rarely practiced. Nevertheless, the importance of accurate voting to public trust in the democratic process justifies the required effort.

5 Characteristics of a Trustworthy Solution

To help ensure the properties of accuracy, privacy, and verifiability described earlier, and to avoid the problems of current approaches, a trustworthy voting system would benefit from a number of other design characteristics, summarized here. These properties allow for a range of solutions, but rank some as much more desirable than others.

Justification A trustworthy solution would come with strong evidence that it operates correctly. This evidence may consist of formal proofs of the correctness of the abstract protocols or the concrete software implementation. Some argue that such verification is impractical. We instead point to cases where it has already been achieved, and to the growing expertise in developing high assurance software, which, at one extreme, involves complete formal proofs of correctness.

Openness A trustworthy solution would benefit from an open, transparent design process and from the release of any source code and verification artifacts either into the public domain or under an open-source license approved by the Open Source Initiative. Openness alone will not guarantee a higher degree of correctness, but it will help prevent outright intentional flaws and will allow independent experts to convince themselves that the system operates as intended.

Simplicity Simplicity helps prevent accidental flaws and helps make intentional flaws easier to detect. It also makes it more tractable to formally verify that the system behaves as intended.

Cooperation Because the voting process is complex and depends on the products and expertise of multiple segments of society, a good solution must involve collaboration between separate groups. In particular, involvement of both commercial and academic organizations is crucial.

Software Independence and Auditability Even with the assurance techniques described so far, we can never be entirely sure that the complete system is free of flaws. While software is an important and useful tool, becoming dependent on it would be a mistake. A trustworthy voting system *must* allow for manual recounting in any situation, even if manual intervention is expected to happen extremely rarely.

6 Justification for TIP Funding

6.1 Map to Administration Guidance

Voting systems profoundly affect the fabric of our government and involve infrastructure spread over the entire country. As the systems become increasingly dependent on software components, we need a technological solution to the problem of ensuring that those components operate correctly.

Developing software that can be shown to work correctly under all circumstances is difficult. Techniques for formal verification of software are powerful but not yet widely used, so the risk of failure is significant. However, the potential reward for even a partial solution would also be substantial, and failure to address the problem could lead to steady decline in the public trust of our elections.

New approaches to software engineering and application of recent computer science research suggest a solution the problem. However, the difficulty of formal methods and the need for a combination of technical knowledge and practical focus have prevented significant, practical application of formal methods technology.

The issue of trust in elections has received considerable attention since the advent of electronic voting systems. Funding to investigate the problem and its possible solutions has come from the National Science Foundation [1] and the Election Assistance Commission [16].

6.2 Justification for Government Attention

The reliability of our election systems affects the entire government process, and the internal pressures of industry and academia all but guarantee that neither will come up with an adequate solution on its own.

A formally verified voting system would not only affect the security of our elections but also provide support for a fundamentally different and more trustworthy approach to software engineering, with wide-ranging consequences affecting our ability to compete in the global market.

Evidence of commitment A number of initiatives over the past few years demonstrate both government and academic commitment to the improvement of our election infrastructure.

On the government side, The Help America Vote Act (HAVA) of 2002 [5] gave the National Institute of Standards and Technology (NIST) a key role in the improvement of voting systems across the country. One outcome of this new role has been the creation of the Voluntary Voting System Guidelines (VVSG) [11], a set of recommended standards for the development of new voting systems that significantly improves on previous practice. HAVA also created the Election Assistance Commission (EAC), with duties including the certification of voting systems and the accreditation of voting system testing laboratories. In 2008, the EAC sent out a solicitation for proposals to perform a risk assessment of voting systems [16].

At a state level, California Secretary of State Debra Bowen has been a vocal critic of many electronic voting systems. In 2007, in concert with researchers from academia, she conducted a top-to-bottom review of the voting systems then in use in California, and collected a wide range of documents detailing their weaknesses [4, 17].

On the academic side, leading researchers at Stanford University, the University of California at Berkeley, Princeton University, the Massachusetts Institute of Technology, Johns Hopkins University, Rice University, and the University of Iowa have worked to analyze the weaknesses of existing voting systems, to create high-level protocols for voting that are guaranteed to have desirable properties, and to begin exploring what would be needed to bring these protocols into practice. Many of these researchers are experts in the use of formal methods for software development.

In industry, several organizations have experience bringing formal methods into practice. Galois, Inc. is a technology transition company with the business of applying cutting-edge research in formal methods and language design to information assurance problems. Galois has extensive expertise in high-assurance cryptography [6] and secure information sharing [7]. SRI International is a non-profit research and development organization with a wide range of focus areas. SRI has made significant contributions to formal methods [9] and computer security [10].

Finally, a number of independent groups, including A Center for Correct, Usable, Reliable, Auditable, and Transparent Elections (ACCURATE), the Verified Voting Foundation, the California Voter Foundation, and the Electronic Frontier Foundation have provided unofficial oversight, documenting weaknesses in existing voting processes and suggesting improved procedures.

Need for funding While the organizations listed above all have a an interest in developing more trustworthy voting systems and have the necessary skills to complete such a project, only a small number of sources have provided relevant funding in the past. NIST is the only organization of which we are aware that is currently providing funding with goals compatible to the development of trustworthy voting systems by a collaborative group of commercial and academic players. Past funding sources often have limited their support to universities or non-profit groups.

6.3 Details of Existing Work

This section provides additional detail about some of the more notable projects and organizations aimed at making voting systems more trustworthy.

California Voting System Review California Secretary of State Debra Bowen conducted a top-to-bottom review of existing voting systems during 2007 [12]. Along with David Wagner, from the University of California, Berkeley, and a number of other researchers, the study included an analysis of the software source code in several DRE voting machines. The study concluded that present DRE machines suffer from security flaws so severe that it was unsafe to disclose details [17].

ACCURATE One of the primary organizations involved in the research of reliable voting technology is A Center for Correct, Usable, Reliable, Auditable and Transparent Elections (ACCURATE) [1], created with funds from the National Science Foundation (NSF). ACCURATE performs research into the use of technology in the voting process, investigates tradeoffs between different voting procedures, and serves as a public resource about issues related to voting systems. It includes members from Johns Hopkins University, Rice University, Stanford University, the University of Iowa, SRI International, and the University of California at Berkeley, with an Advisory Board spanning a wide range of other independent, academic, and government organizations.

In one particularly relevant effort, ACCURATE has coordinated an initial investigation into the construction of a voting machine designed from the start with verification through formal methods in mind. The work is currently in its early stages.

The Scantegrity II Prototype David Chaum and Ronald Rivest, along with a group of other researchers, developed a prototype voting system called Scantegrity II [14]. In this system, voters mark paper ballots in the usual way, but with special pens. The oval next to each candidate contains a random code printed in ink that is invisible until it reacts with the ink from the pen.

A voter can (if desired) mark down the codes revealed next to each choice (which are different on every ballot), along with the unique serial number on the ballot. Then, the voter can look up this serial number on a public web site and check the codes recorded along with it. If the codes match those revealed in the chosen bubbles, the vote was accurately recorded.

The Scantegrity II system is the closest existing approximation of a high-assurance, voter-verifiable election technology. However, while the principles of the system are sound, the software component is only a proof of concept and is potentially vulnerable to programmer mistakes or malicious attacks. A trustworthy solution must explicitly include provisions to prevent or detect software errors.

6.4 Essentials for TIP Funding

Collaboration is essential for the development of trustworthy voting systems. Neither industry or academia has the capabilities solve the problem alone, and no other funding organization has demonstrated a commitment to a collaborative approach.

In addition, The United States has fallen behind in the pursuit of software quality. Much of the research in formal methods and most of the successful application have taken place in Europe. Examples include the MÉTÉOR system used by Paris Metro Line 14 [3], a driverless shuttle for the Paris-Roissy airport [2], the SACEM system for train speed control [8], and the Airbus control systems [15].

Though verification is indeed very difficult, it is feasible. The use of formal methods for a trustworthy voting system would position the U.S. as a significant player in the field of quality software. The successful development of verified voting system software would be a significant step toward showing that such techniques are ready for practical application. It would also help reinforce our position as the world's leading democracy by demonstrating that we are serious about safeguarding the the key components of the democratic process.

The Help America Vote Act (HAVA) [5] assigned the National Institute of Standards and Technologies (NIST) with the task of recommending voluntary voting system guidelines. Therefore, NIST has a fundamental role in the development of voting standards. NIST also has an interest in fostering technological competitiveness. No other organization is more appropriate to support an effort to improve the quality of our nation's election infrastructure.

References

[1] A Center for Correct, Usable, Reliable, Auditable, and Transparent Elections. Available from World Wide Web: `http://accurate-voting.org/`.

[2] Frédéric Badeau1 and Arnaud Amelot. Using B as a high level programming language in an industrial project: Roissy val. *ZB 2005: Formal Specification and Development in Z and B*, 3455/2005:334–354, 2005.

[3] Patrick Behm, Pierre Desforges, and Jean-Marc Meynadier. MÉTÉOR : An industrial success in formal development. *B'98: Recent Advances in the Development and Use of the B Method*, 1393/1998, 1998.

[4] Matt Bishop. Overview of red team reports. Available from World Wide Web: `http://www.sos.ca.gov/elections/voting_systems/ttbr/red_overview.pdf`.

[5] Federal Election Commission. Help America Vote Act of 2002. Available from World Wide Web: `http://www.fec.gov/hava/hava.htm`. Public Law 107-252.

[6] Galois, Inc. Communications security. Available from World Wide Web: `http://galois.com/technology/communications_security`.

[7] Galois, Inc. Cross-domain solutions. Available from World Wide Web: `http://galois.com/technology/cross-domain_solutions`.

[8] Claude Hennebert and Gérard D. Guiho. SACEM: A fault tolerant system for train speed control. In *Proceedings of The Twenty-Third Annual International Symposium on Fault-Tolerant Computing*, pages 624–628, 1993.

[9] SRI International. Formal methods and dependable systems. Available from World Wide Web: `http://www.csl.sri.com/programs/formalmethods/`.

[10] SRI International. Secure systems. Available from World Wide Web: `http://www.csl.sri.com/programs/security/`.

[11] National Institute of Standards and Technology Technical Guidelines Development Committee. Next Version Voluntary Voting System Guidelines (VVSG), August 2007. Available from World Wide Web: `http://vote.nist.gov/vvsg-report.htm`.

[12] California Secretary of State. Voting systems review. Available from World Wide Web: `http://www.sos.ca.gov/elections/elections_vsr.htm`.

[13] Julian M. Pleasants. *Hanging Chads: The Inside Story of the 2000 Presidential Recount in Florida*. Palgrave Macmillan, August 2004.

[14] The Scantegity Team. Scantegrity II. Available from World Wide Web: `http://www.scantegrity.org/`.

[15] Pascal Traverse, Isabelle Lacaze, and Jean Souyris. Airbus fly-by-wire: A total approach to dependability. In *Building the Information Society*, volume 156/2004, pages 191–212. Springer Boston, 2004.

[16] U.S. Election Assistance Commission. Voting systems risk assessment. Available from World Wide Web: `https://www.fbo.gov/index?s=opportunity&mode=form&tab=core&id=6dcd89aa66e1a4e12777deaf0891d9a6&_cview=0`. Solicitation Number: EAC-RDV08-R-001.

[17] David Wagner. Principal investigator's statement on protection of security-sensitive information. Available from World Wide Web: `http://www.sos.ca.gov/elections/voting_systems/ttbr/State_of_protect(DW).pdf`.