The Bridge Archives | Subscribe

Voting as an Engineering Problem (Print This)

Michael Ian Shamos

Volume 37, Number 2 - Summer 2007

We have never encountered a security problem with electronic voting machines.

Electronic voting has been in the news almost constantly since 2000, when the punch-card debacle in Florida caused a rapid re-evaluation of elections in the United States and fomented a revolution in voting processes and technologies. In 2002, Congress passed the Help America Vote Act (HAVA), which allocated more than $3 billion to replace punch-card equipment, and states, eager to spend this bounty, rushed off to purchase a wide variety of electronic replacements.

A learning curve is to be expected after the introduction of unfamiliar technology, and election jurisdictions experienced a wide variety of problems with their new systems, including machine failures, inadequately trained poll workers, poor ballot setup, and voter confusion. In some cases, votes validly cast were not counted and were lost forever. In 2004, in Carteret County, North Carolina, in a now-classic failure, a voting machine with a capacity of about 3,000 votes was used to record more than 7,000, resulting in the loss of about 4,500 votes. As a result, a new election had to be held. The causes of incidents of this kind are generally well understood and can be traced to canonical engineering failures. In Carteret County, for example, if the voting machine had been engineered not to accept ballots once its memory was full, the problem would not have occurred.

One problem that was not encountered, and has never been encountered, is a security incident involving an electronic voting machine. There is no evidence that anyone in the 28-year history of direct-recording electronic voting machines (DREs) has ever breached security to alter the outcome of an election. Nevertheless, following the publication of a report suggesting that tampering might be possible (Kohno et al., 2004), public attention has focused almost exclusively on the issue of computer security. In recent months, unjustified fears caused the Maryland legislature to discard its statewide system, which cost the state more than $100 million, in favor of a completely different, antiquated technology, believed (wrongly) by law-makers to be more secure. The Florida legislature, reeling from embarrassment over yet another Florida voting incident in 2006, made the same decision, to outlaw DREs and replace them with optically scanned paper ballots. Neither legislature commissioned a scientific study comparing the old and new systems or received or evaluated relevant engineering data. In the end, the cries of activists not only trumped science, they pushed it off the table altogether.

Activists have pushed science aside in addressing problems with voting technologies.

History

Voting in the United States has always been subject to manipulation. A thorough but frightening treatment can be found in *Deliver the Vote* by Tracy Campbell (2005). Since the earliest days of paper ballots, a wide variety of techniques have been used to influence election results, including forging, altering, losing, substituting, and augmenting ballots, to say nothing of vote-buying and other coercive schemes. This very problem led to the invention in 1892 of the lever machine, which its creator, Jacob H. Myers, claimed could “protect mechanically the voter from rascaldom, and make the process of casting the ballot perfectly plain, simple and secret.”

As technology has advanced, mechanical machines have been replaced by computerized systems, including punch-card, optical scan, and DRE devices. The underlying reason for using new technologies is not only to streamline the voting process, but also to speed up the count so winners can be declared quickly.

During this technological development, the important distinction between document ballots (those marked by the voter on a physical medium) and non-document ballots (totals recorded on counters or ballot images stored electronically) was lost or ignored. A document ballot provides only one copy, namely, the one marked by the voter, and this copy must be handled to be tabulated. In the handling, a document ballot may be altered or substituted, and once a substitution is made, it is impossible to recover the original voter’s selections, which are gone forever. The movement from paper to lever machines to DREs was an effort to eliminate the inherent insecurity of document ballots.

By contrast, in a DRE system multiple copies of ballots are retained on different media in different physical locations. To affect the outcome of an election, redundant encrypted records would have to be altered. To date, no one has even suggested, let alone demonstrated, a way to do this.

Despite our lurid history of tampering with document ballots, repeated calls are being made to return to hand-counted paper ballots or optically scanned ballots, which are often used in other countries. However, most other countries use extremely simple ballots. In the more than 90 parliamentary and multiparty democracies in the world, voters typically select a party, rather than individual candidates. One race, one choice.

Despite their use in foreign countries, hand-counted paper ballots are completely impractical in the United States, because our ballots are far more complicated than any other country's. For example, the 2006 (off-year) general ballot in Marin County, California, included 98 candidates in 30 races and 30 ballot propositions. The resulting mark-sense ballot was six pages long. Or consider the still-disputed 2006 U.S. House race in Sarasota County, Florida, which presented voters with 21 pages of electronic ballot followed by four pages of review screens. Perhaps we shouldn't have been surprised that 15 percent of the voters were unable to cast votes for either congressional candidate in Sarasota.

The perception persists that we cannot control the security of isolated, non-networked machines that have no operating systems, but that problems with the production, storage, transportation, and counting of paper ballots were solved long ago. In fact, there are neither standards nor procedures nor any reference work on protecting, transporting, and storing document ballots in a secure way.

The Inevitability of Errors

The very notion that perfection is achievable in a voting system is a fantasy. All systems have error rates, which are easy to define but difficult to measure. In a given race, a system is "correct" if it records and tabulates votes for the people the voters intended to vote for. The error rate for that race is the percentage of recorded choices that do not match the voters' intent.

In any real election, there are errors, not necessarily because of machine failures or incorrect counting logic, but because of human interface issues. A voter may, for whatever reason, not understand how to cast a vote or may overlook a race or make a mistake. The error rate cannot be easily measured because there is no effective way to determine voter intent other than through the voting system itself. (If there were such a method, it would be a perfectly good voting mechanism by itself!)

In one of the first experiments of its kind ever conducted, subjects were given a sheet of paper telling them how to vote in a mock election; this sheet became the voter's intent. Each subject used a single voting machine, but several different types of machines were used in the study. In general, error rates were approximately 10 percent. That is, 10 percent of the ballots cast failed to conform in some way to the predetermined choices. When straight-party voting was allowed, the error rate went up to 20 percent.[1] And these rates were observed with fully functional equipment (Herrnson et al., 2006).

A 10-percent differential would have changed the outcome in 25 of the 45 presidential elections in which the popular vote was recorded.[2] However, even assuming that the error rate could be reduced to, say, 1 percent, elections inevitably will occur in which the margin of victory is below the error rate. Noise can exceed signal. In fact, this must happen because there is a lower bound on the error rate for a given system, but no lower bound on the winning margin.

Close elections occur frequently. The famous Bush-Gore margin in Florida in 2000 was 537 votes out of almost 6 million votes cast, a difference of less than 0.1 percent. How likely is such a close election? Assuming an equal distribution of Republicans and Democrats, if a sample of six million votes is taken from the population, the margin will be less than 537 votes about one-third of the time. The 2006 margin in Sarasota was 0.15 percent, far lower than the most optimistic estimate of system error rate for any voting method in existence.

The notion of perfection in a voting system is a fantasy.

An Engineering Formulation

Just because voting is fundamental to democracy does not remove it from the realm of science and engineering. A voting system must perform three functions:

- Present the correct set of possible choices to the voter.
- Capture and record the voter's actual choices.
- Produce an accurate tally of the actual choices.

It must satisfy the following conditions:

- Privacy—an individual voter's choice cannot be discoverable by any other person.

CONTACT

National Academy of Engineering

500 Fifth Street, NW
Washington, DC 20001
Tel. 202.334.3200
Fax. 202.334.2290
Staff Directory

Comments & Feedback

- Security—the system must not be susceptible to manipulation or corruption by a small set of individuals.[3]
- Auditability—independent, after-the-fact observers must be able to confirm that the correct person has been elected.
- "Receipt-freeness"—the voter must not be able to prove to any other person how he or she voted.
- Verifiability—the voter must be able to determine with confidence after the election that his or her vote was recorded and tabulated accurately, without violating the requirement of receipt-freeness.

Usability and reliability are relevant only to the extent that they affect the functions or conditions listed above. They are not separate requirements.

The design of a system that can achieve the required objectives is purely a matter of engineering. No one has proven that a system that satisfies the requirements either exists or cannot possibly exist. Arguments have been made that privacy, verifiability, and auditability are mutually inconsistent and that auditing and verification are impossible without being able to determine how a particular voter voted (or permitting that voter to prove how he or she voted). However, it has been demonstrated numerous times through the use of cryptography that the common intuition is incorrect (see, e.g., Acquisti, 2004; Neff, 2001).

Security deficiencies can be designed out of a system.

The issue of security requires careful attention. Several systems that have been deployed in the field have significant vulnerabilities that have been remediated through administrative procedures rather than through redesign of the system. The most serious of these vulnerabilities would allow an attacker who gained private physical access to a voting machine to replace its software/firmware with a Trojan horse. Whether such an attack could be mounted successfully and undetectably in a way that would affect more than a small number of machines is purely speculative. However, security deficiencies should be, and can be, designed out of a system.

A demonstration that an attacker, under unrealistic conditions, might corrupt a voting machine, although cause for concern, does not justify discarding an entire technology. The United States did not outlaw bridges after the Tacoma Narrows Bridge collapsed due to uncontrolled wind-induced oscillation. Instead, bridges of similar design were reconfigured, and the drawbacks of the original design informed future designs. When the gas tank of the Ford Pinto caught fire and caused repeated highway deaths, we did not abandon cars, or even the Pinto. The offending models were recalled, and their tanks were replaced. Yet when computer security experts merely pointed out the possibility of security intrusions in voting machines, Congress and state legislatures moved to ban DRE machines entirely, one of the most extreme overreactions in the history of engineering.

Reliability

Voting machines are among the least reliable devices on this planet. It has been reported anecdotally that approximately 10 percent of DRE machines fail in some respect during the average 13 hours they are in use on election day.[4] In some cases, the percentage is much higher (Bishop et al., 2005). The percentage rises to 20 percent, for example, if the machines have paper-trail printers attached to them. If electric razors failed en masse within the first 13 hours, the Federal Trade Commission would institute action against the manufacturer. Why don't voting machines and their manufacturers face such consequences? The unfortunate answer is that these failure rates are permitted by applicable standards.

The development of voting standards is a cautionary tale too long to relate here, but the currently applicable standard is that a voting machine must exhibit a mean time between failure (MTBF) of 163 hours.[5] Under the exponential failure model, one would expect 10 percent of machines with this MTBF to fail within 17 hours. Given that machines are typically used for several hours during pre-election testing and that the ones that fail are removed from service, the observed failures are consistent with the standard. An ordinary personal computer has an MTBF of about 30,000 hours, and many voting machines are constructed from simplified PCs. So how can we explain the difference in reliability?

In fact, the standards for voting machines, which were heavily influenced by the manufacturers, are grossly inadequate. One reason is that they were developed ad hoc rather than from a set of principles applicable to voting systems. In addition, some standards are imprudent, to say the least. For example, there is a prohibition against interpreted code, which excludes some excellent programming languages, such as Java. Other standards are irrelevant or poorly formulated. But most important, necessary standards, such as computer security requirements, are missing altogether.

With the passage of HAVA, the National Institute of Standards and Technology (NIST) became responsible for formulating guidelines for voting systems. NIST has focused significant effort in this direction, but for a guideline to be adopted, it must pass first through two bodies, the Technical Guidelines Development Committee and, ultimately, the Election Assistance Commission (EAC).

The EAC, which was created by HAVA, is inherently political. By law, it is composed of four members, two

nominated by the majority party in Congress and two nominated by the minority party. EAC members, being presidential appointees, must then be confirmed by the Senate. This structure is supposedly designed to ensure that exacting technical standards are adopted.

Conclusion

It is no surprise that rampant reliability problems have led to widespread, largely justified mistrust of voting machines. The threshold question is whether we will implement and deploy reliable machines or turn to a technology with reliability properties that have not yet been studied. The first alternative is difficult because there is no accepted design for a reliable voting machine (e.g., 30,000-hour MTBF), and certainly none that satisfies the properties listed above.

If an election is not close (that is, there is a large difference in vote totals between the winner and the next-highest vote getter), it doesn't matter which voting system is used. Even a bad system can discriminate when the margin is large. The sad fact, which is not obvious, is that, if an election is very close, it also makes no difference which voting system is used, because even a good system cannot determine in a trust-worthy way who wins when the margin is very close (say, less than 0.1 percent). This is because of the large number of operational components and human beings par-ticipating in an election and because there is no system that can provide a post-election audit that will satisfy the losing candidate and his supporters. The situation might be different if universally verifiable systems were available.

Under the present circumstances, there is a great risk that the development of truly verifiable systems may be delayed for decades. (Paper-trail and optical-scan systems provide only instantaneous verification, which is insufficient to settle a close election.) When a legislature outlaws DRE voting, an inventor has no incentive to design a verifiable electronic system, which would then be illegal. If Congress adopts the bill introduced by Congressman Rush Holt (currently HR 811), which requires verifiable paper trails, there would be no point in improving current systems, or engineering a new one, because it could not be used in elections for federal office.

I am arguing for federal funding of an engineering project to develop an electronic voting system that can meet the requirements listed above. We know that government is capable of acting when it is properly motivated. After NASA's Challenger disaster of 1986, President Reagan created a select scientific panel, the Rogers Commission, to determine the cause of the accident and even to review the culture at NASA that may have contributed to it. The outcome was a striking success, and the space program continued. Nothing of the kind is in sight for electronic voting, however, and we appear destined to repeat the paper manipulations of the nineteenth century that led to the development of voting machines in the first place.

References

Acquisti, A. 2004. Receipt-Free Homomorphic Elections and Write-in Voter Verified Ballots. Carnegie Mellon University Technical Report CMU-ISRI-04-116 (April 2004).

Bishop, M., L. Guarino, D. Jefferson, and D. Wagner. 2005. Analysis of Volume Testing of the AccuVote TSx/AccuView. Report of the California Voting Systems Technology Assessment Advisory Board (October 11, 2005). Available online at http://www.ss.ca.gov/elections/voting_systems/vstaab_volume_test_report.pdf.

Campbell, T. 2005. Deliver the Vote: A History of Election Fraud, an American Political Tradition, 1742–2004. New York: Carroll & Graf.

Herrnson, P., R.G. Niemi, M.J. Hanmer, B.B. Bederson, F.G. Conrad, and M. Traugott. 2006. The Importance of Testing the Usability of Voting Systems. Paper presented at the Usenix/ACCURATE Electronic Voting Technology Workshop, Vancouver, British Columbia, August 1, 2006. Available online at http://www.usenix.org/events/evt06/tech/full_papers/herrnson/herrnson_html/.

Kohno, T., A. Stubblefield, A.D. Rubin, and D.S. Wallach. 2004. Analysis of an Electronic Voting System. Presented at the IEEE Symposium on Security and Privacy, Oakland, California, May 2004. The contents of the paper were widely circulated in 2003. Available online at http://avirubin.com/vote.pdf.

Neff, A. 2001. A Verifiable Secret Shuffle and Its Application to E-Voting. Pp. 116–125 in Proceedings of the ACM-CCS'01. New York: ACM Press. Available online at http://delivery.acm.org/10.1145/510000/502000/p116neff.pdf?key1=502000&key2=2841797711&coll=GUIDE&dl=GUIDE,ACM&CFID=11111111&CFTOKEN=2222222.

FOOTNOTES

1. "Straight-party voting" means that a virtual office, called the straight-party office, is on the ballot. Choosing a party for that office has the effect of casting a vote for all candidates of that party throughout the ballot. On multipage ballots this causes an unseen side effect because a voter cannot see the consequences of his or her choice without navigating the entire ballot.
2. This does not take into account the electoral vote.
3. Freedom from corruption cannot be guaranteed if arbitrary subsets of persons are allowed to collude.
4. In general, failures result in the loss of none, or possibly one, vote, namely, the vote that was in the process of being cast when the failure occurred. The problem is not loss of votes but the loss of voting capability, voter inconvenience, and the loss of voter confidence.
5. See the 2002 Federal Voluntary Voting System Guidelines, Sect. 4.3.5.

About the Author

Michael Ian Shamos is Distin-guished Career Professor in the School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania. He has been an official examiner of voting systems since 1980.