

Electronic Voting Systems: the Good, the Bad, and the Stupid

Barbara Simons

“Those who cast the votes decide nothing; those who count the votes decide everything.”
Joseph Stalin

“We always pray for large margins.” Theresa LePore, designer of the “butterfly” ballot.

As a result of Florida 2000, some people concluded that paper ballots simply couldn't be counted¹, even though businesses, banks, racetracks, lottery systems, and other entities in our society count and deal with paper all the time. Instead, paperless computerized voting systems (Direct Recording Electronic or DREs) were touted as the solution to “the Florida problem.” Replacing hanging chads with 21st century technology, proponents claimed, would result in accurate election counts and machines that would be impossible to rig. Furthermore, with nothing to hand-count and no drawn-out recounts, computerized voting systems could report results shortly after the polls close. Many election officials loved the idea, believing the new machines would be cheaper and more reliable than the old systems. Also, the lack of recounts meant that they could go home early on Election Day. Vendor enthusiasm was enhanced by the almost \$4 billion of US government money that was promised in the Help America Vote Act (HAVA), passed in 2002. Yet now, two years after the passage of HAVA, voter verifiable paper trails are being demanded by numerous public interest groups, computing professionals, and members of Congress. Where did things go wrong?

Electronic voting machine software is proprietary, the certification testing process is both secret and incomplete, and the test results are secret. The tests check only for requirements in the Federal Election Commission (FEC) guidelines. To top things off, Commercial Off The Shelf software (COTS) contained in voting systems is not examined in any of the testing, simply because FEC guidelines don't require it.

For years, prominent computer security experts have been arguing that paperless DRE machines present major security problems, including buggy software and the risk of malicious code affecting the outcome of an election. But the warnings of experts such as Rebecca Mercuri (<http://www.notablesoftware.com/evote.html>) and Peter Neumann (<http://www.csl.sri.com/users/neumann/neumann.html#5>) were largely unheeded by election officials and the public until David Dill created a petition

¹ The most outspoken advocate of paperless DREs is Jim Dickson, Vice-President of the American Association of People with Disabilities. According to the NY Times, the AAPD received \$26,000 from vendors this year. (The National Federation for the Blind received a million dollars from Diebold in settlement of a lawsuit). The League of Women Voters also lobbied on behalf of paperless DREs. However, the national office retracted its support of DREs when the members revolted at the recent LWV convention.

(<http://www.verifiedvoting.org/index.asp>) calling for voter verifiable audit trails for voting systems. The core idea behind the Dill petition is that the voters should be able to verify that their ballots have been correctly recorded; also, it should be possible to conduct a meaningful recount.²

A few horror stories

Because of the secrecy surrounding almost every aspect of e-voting – along with a lack of public national incident reporting – independent computing technologists can provide only limited analyses of problems relating to hardware, software, testing, security, and human factors. Nonetheless, evidence of these problems is widespread and varied. A few representative examples follow.

In January 2004 a special election was held in Broward County, Florida. Only one contest was included on the ballot. Yet, of the 10,844 votes cast on ES&S (Election Systems & Software) paperless touch screen voting machines, 134 were ... for no one at all. Since the winning candidate won by only 12 votes, people understandably wondered what had become of those 134 votes; there was no way of telling if some had been lost by the computer. The mayor of Broward is now calling for paper ballots.

In November 2003 in Boone County, Indiana over 144,000 votes were cast even though Boone County contains fewer than 19,000 registered voters. And, of those, only 5,532 actually voted. The county clerk stated that the problem was caused by a “glitch in the software.” Updated results then were obtained that were consistent with the number of people who had actually voted, and the public was reassured that the new electronic tally was accurate. Still, because the county used paperless MicroVote (an Indiana company) DREs, it was impossible to verify independently that the updated results were indeed correct.

When the polls opened in Hinds County, Mississippi in November 2003, voters arrived to find the WINvote DREs at the polls were down. Worse yet, there were no paper ballots available. By mid-morning, some machines were still down. Voters complained about waiting in long lines and of having to complete makeshift paper ballots – some being nothing more than scraps of paper – without adequate privacy. At 8 p.m., there were still voters standing in line. One report claimed the machines had overheated. Subsequently, the Mississippi State Senate declared the results in that district invalid and scheduled a new election. Had paper ballots been made available to voters, the machine related problems could have been bypassed.

Diebold – a case study in incompetence

² To avoid the risk that the machine prints the correct result but stores an incorrect result in computer memory, some number of paper ballots randomly selected should be manually recounted as a check on the machines.

Diebold, which has been manufacturing ATMs for years and is one of the major DRE vendors, has become the poster child of all that is wrong with DREs. Diebold's involvement with voting machines received significant national press when the CEO of Diebold, Walden O'Dell, stated in an August 14, 2003 letter to Central Ohio Republicans that he was "committed to helping Ohio deliver its electoral votes to the President next year."

However, the PR problem triggered by O'Dell's statement pales in comparison to the technical incompetence of Diebold uncovered when Bev Harris (<http://www.scoop.co.nz/mason/stories/HL0302/S00036.htm>) announced in February, 2003 that she had discovered Diebold voting machine software on an open FTP website. Computer science professors Avi Rubin and Dan Wallach, and their students Tadayoshi Kohno and Adam Stubblefield, subsequently analyzed some of that software and published a security analysis in a paper that is sometimes referred to as the "Hopkins paper" (<http://avirubin.com/vote/analysis/index.html>). One of the more shocking revelations was that Diebold used a single DES key (F2654hD4) to encrypt all of the data on a storage device. Consequently, an attacker with access to the source code would have the ability to modify voting and auditing records. Perhaps even more surprising, Diebold had been warned in 1997 about their sloppy key management by Douglas Jones, a professor of computer science at the University of Iowa and a member of the Iowa Board of Examiners for Voting Machines and Electronic Voting Equipment:

[N]either the technical staff nor salespeople at Global Election Systems [purchased by Diebold in 2001] understood cryptographic security. They were happy to assert that they used the Federally approved Data Encryption Standard, but nobody seemed to understand key management, in fact, the lead programmer to whom my question was forwarded, by cell-phone, found the phrase *key management* to be unfamiliar and he needed explanation. On continued questioning, it became apparent that there was only one key used, company wide, for all of their voting products. The implication was that this key was hard-coded into their source code!³

Because of the security issues raised in the Hopkins paper, the State of Maryland, which had just committed to purchasing Diebold DREs, commissioned a study of Diebold machines by Science Applications International Corporation (SAIC). The SAIC report (<http://www.dbm.maryland.gov/DBM%20Taxonomy/Technology/Policies%20&%20Publications/State%20Voting%20System%20Report/stateVotingSystemReport.html>) is a very fast read, since only about 1/3 of it was made public – the rest was redacted.⁴ But even the limited amount of information that was released in the report is quite damning. For example, the report states that the Diebold system is so complicated that even if all of

³ Doug Jones provides an excellent overview of the Diebold story at <http://www.cs.uiowa.edu/~jones/voting/dieboldftp.html>.

⁴ According to Frank Schugar, project manager for SAIC, the report was redacted by Maryland, not by SAIC. The Electronic Privacy Information Center has submitted a public records request to obtain the full unredacted version.

the problems were fixed, there still could be security risks because of poorly trained election officials.

Section 5 of the report, which “provides the risk assessment findings, including a discussion of the SBE security requirements, threats to the implementation of the AccuVote-TS, likelihood of exploitation of the threat, vulnerabilities, and mitigation strategies and recommendations for improving the security posture” is completely redacted.⁵

Even the name of the operating system being used is redacted (page 17): “The voting terminal is an embedded device running Microsoft Windows [redacted] as its operating system.” However, we know from internal Diebold emails that were posted on the web that Diebold was running Windows CE 3.0.

Why, one might ask, would anyone feel the need to redact the name of the Windows operating system being used by Diebold? A likely explanation is that Windows CE is a modular OS tool kit that allows different operating systems to be assembled for different embedded applications. Yet, the certification process treats Windows CE as being equivalent to a non-modified operating system, which means that the actual code is never examined.

In spite of the fact that even the redacted version of the SAIC report was very critical of Diebold and supported the Hopkins report on most issues, both the State of Maryland and Diebold claimed that the SAIC report vindicated the purchase of Diebold machines.

In November 2003, the Maryland Department of Legislative Services commissioned yet another study of Diebold machines by RABA Technologies (http://www.raba.com/text/press/TA_Report_AccuVote.pdf). The *Trusted Agent* report, released in January 2004, based on a “red team” effort to hack Diebold voting systems, revealed physical security problems such as the use of identical keys on security panels covering PCMCIA and other sockets on the machines – as well as locks that could be picked in a few seconds.

Unfortunately, when DRE vendors discuss the virtues of DREs to election officials, they gloss over security issues related to short- and long-term storage of the machines, as well as machine access control before and after elections.

Meanwhile, the State of Ohio, which had been considering the purchase of Diebold DREs for the entire state⁶, hired Compuware to test hardware and software, and InfoSentry to conduct a security assessment. The Compuware study uncovered yet another hardwired password, this time involving the supervisor’s card, used to start up each voting machine on Election Day as well as to terminate the voting process at the end of the day. When

⁵ The description of Section 5 is on p. 2. It probably was supposed to have been redacted, since the title of Section 5 is redacted in the Table of Contents.

⁶ Diebold is headquartered in Ohio.

the card is inserted into the DRE, the election official must enter the same password or PIN⁷ that is hardwired into the card - but not into the voting software. Consequently, someone who is able to obtain a supervisor's card, or who manages to create a fake card with a different password, would be able to conduct a denial of service attack by prematurely halting the voting machines, thereby denying some voters the opportunity to vote.

ES&S – a software bug prevents audits

An intriguing link had existed for a long time between Diebold and ES&S, another major voting machine vendor, and had generated a great deal of criticism. Until very recently, Bob Urosevich was the CEO of Diebold Election Systems (O'Dell is the CEO of the parent Diebold company), and his brother Todd had been the vice-president of ES&S⁸. Together, DREs and optical scan voting systems manufactured by Diebold and ES&S will count somewhere between two-thirds and 80% of the ballots in the November election.⁹

There is also a connection between ES&S and Sen. Chuck Hagel (http://www.csd.cq.com/senate_mem/s0531.html). Until two weeks before he announced his candidacy for the Senate in 1996, Sen. Hagel had been the CEO of American Information Systems, Inc., a fact not mentioned in Hagel's 1996 campaign financial disclosure statements. AIS, founded by the Urosevich brothers, subsequently purchased another company and become ES&S. AIS was used to count many of the votes that elected Sen. Hagel to the Senate in 1996, the first Republican to have been elected from Nebraska in twenty-four years. Hagel's 2002 Democratic opponent, Charlie Matulka, claims that Hagel owned 35% of ES&S, when ES&S machines were used to count the votes in the 2002 Senate race.

More recently ES&S has been in the news, because a software bug had corrupted the audit log and vote image report in ES&S machines used in Miami-Dade and many other parts of the country.¹⁰ An internal memo written in June 2003 by Orlando Suarez, division manager of the county's enterprise Technology Services Department and obtained through a public records request made by the Miami-Dade Election Reform Coalition, describes a discrepancy in the internal auditing mechanism of the ES&S machines. Suarez stated that the software bug(s) make the audit reports "unusable for the purpose that we were considering (audit an election, recount an election and if necessary, use these reports to certify an election)." This information was not made public until it

⁷ The Compuware study discovered that the pin was 1111.

⁸ Whether they left their positions because of criticism from groups concerned about collusion or for some other reasons is not known to this author.

⁹ See the attachments in http://www.electiondataservices.com/EDSInc_DREoverview.pdf for a detailed breakdown by machine type.

¹⁰ For a detailed discussion of the ES&S bug, see <http://www.cs.uiowa.edu/~jones/voting/miami.pdf>

was announced by the Coalition in April 2004, almost a year after the initial memo was written.

The event log contained results for some nonexistent machines, and it also failed to report all the results for the machines that were in operation. According to Doug Jones, there were actually two bugs. One - triggered by a low battery condition - caused corruption in the event log; the second caused the election management system to misread the machine's serial number in the face of this corruption. While the vote count was not impacted, the problems uncovered are symptomatic of the kinds of anomalies that are not tested for under the certification process, discussed below.¹¹

On July 27, 2004 the Miami-Dade Election Reform Coalition announced that audit data they had requested revealed that computer crashes had deleted all the election results from the September 2002 gubernatorial race in Miami-Dade, as well as from several more recent municipal elections. It appeared that no back-ups had been made, leading to speculation that the loss of the ballot images could be a violation of Florida law regarding the retention of ballots.¹² After spending a few embarrassing days trying to explain how election officials could have lost critical voting records, Miami-Dade County elections supervisor Constance Kaplan announced on July 30 that her secretary had located a computer disk containing the missing data in the conference room next to her office.¹³

In an interesting footnote to the Miami-Dade story, Florida Secretary of State defended the paperless touch screen voting machines against criticism that she likened to conspiracy theories by saying, "The touch-screen machines are not computers. You'd have to go machine by machine, all over the state [to rig an election]."

How did such flawed machines become certified?

The first FEC standard for electronic voting machines, issued in 1990, was replaced in 2002 (<http://www.fec.gov/pages/vssfinal/vss.html>). Many voting systems in use today were certified to the 1990 standards.

¹¹ Quoting Jones, "As of midsummer, the state of Florida has approved a fix to the two bugs that caused this problem, and in the pre-election testing conducted on August 13, the event records were extracted from compact flash cards showed correct reports of low battery conditions without any corruption of serial numbers. Curiously, it was a member of the Miami-Dade coalition who found this evidence as she went over printouts of the event logs generated from the compact flash cards."

¹² Amazingly, Miami-Dade officials chose to ignore a memo sent before the crashes occurred in which Cathy Jackson of the county's Audit and Management Services Department warned of the lack of back-up and suggested burning all data to CD ROMs after each election.

¹³ Quoting Jones, "The disk was a CD-R in a file folder. The county had only begun making archival CD-R copies of the data after the county Audit and Management Department suggested that they do so that summer. Apparently, although this was being done, there was as yet no institutional memory of where these disks were being put."

Machines are tested and certified by three private companies - Ciber, Wyle, and SysTest – which are referred to as Independent Testing Authorities (ITAs). The ITAs themselves are certified by the National Association of State Election Directors, but are not subjected to any government oversight. Individual states may have additional requirements that are certified by the ITAs. Vendors pay for all testing.

One of the bizarre aspects of the certification process is that distinguishes between “firmware” and “software”, with “firmware” being defined as the software that runs in the voting machines in the precinct, while “software” is used to refer to the code utilized by the election management system. Wyle certifies only firmware, and Cyber certifies only software. SysTest certifies the entire system.

Rather than checking the software for security flaws and attacking the software to see if it can be compromised, the ITAs limit their tests strictly to items specifically required by the FEC standards. Particularly prominent among these are control flow requirements, with Do-While (False) constructs and intentional exceptions used as GoTos being explicitly prohibited. The 2002 FEC standards also call for “effective password management,” but the phrase is not defined. We can infer from the Diebold results, however, that no one is checking to see if encryption keys have been hardwired into the code. The testing also fails to check for exceptions, and there are no provisions for the inspection of COTS code.

States typically are provided with only a one-page certificate saying that the software satisfied the FEC standards. By contrast, vendors are given detailed test results. Some states request the test results, but results have been provided only when the states or election officials sign non-disclosure agreements. Not only should test results all be made public, but there also should be a central data depository that collects all test results and problem incidents from voting machines - much as is done for airplanes - so that the government and election officials can check to make sure that all known problems have been rectified.

Then there is the matter of ballot definition files (BDF). These files contain the candidates and issues information for each election. Because BDFs tend to be difficult for election officials to write, they frequently are prepared by the vendors. Whether the BDFs are prepared by the vendor or by someone local, they can't be produced until the candidates and issues have all been decided.

Although critical to elections, BDFs are never independently inspected by an ITA. While properly conducted pre-election testing should uncover errors in BDFs, such testing is not routine in many jurisdictions, where state laws merely require that the tests include casting at least one vote for each candidate in each race on the ballot, using each ballot style in use in the jurisdiction. In Miami-Dade County, for example, there were 222 distinct ballot styles in the August 2004 primary.¹⁴

¹⁴ Private communication with Doug Jones.

When errors in BDFs do occur – leading, for example, to votes for one candidate being credited to a different candidate – they can be detected with optical scan voting systems, because anomalous computer-reported results can be discovered through manually recounts of paper ballots.¹⁵ With paperless DREs, however, there is no way to perform such a recount.

Malicious code

While many obvious software bugs have been inferred or uncovered, to my knowledge no clearly malicious code has been detected in voting machine software, though some software bugs have behaved as if they were malicious. An obvious approach for dealing with buggy or malicious code is the use of open, or at least public, source software.

Making software public would expose it to more eyes, thereby increasing the likelihood of the bug detection. But there is still the risk that the software running on the voting machines may not be identical to the software that was made public. Further, as we know from Ken Thompson's Turing Award speech "Reflections on Trusting Trust" <http://www.acm.org/classics/sep95/>, it is possible to write a compiler that will insert malicious code into object code.¹⁶

Even open source code can be vulnerable. A recent attempt to insert a two-lines-of-code backdoor into Linux was caught by some observant programmers <http://kerneltrap.org/node/view/1584>. But, the fact that this particular backdoor attempt was stymied is no guarantee that some equally subtle future attempt will also be detected.

With inadequately tested secret code, one can only speculate about the likelihood that any malicious code, especially code that is cleverly designed to resemble a software bug (e.g. =, instead of ==), will go undetected.

Alternative models for voting design

Diebold, Sequoia, ES&S, and Hart InterCivic are the major manufacturers of paperless DREs. Most DREs use touch screens as inputs, though Hart InterCivic uses a dial for candidate selection. DREs also can be equipped with earphones and various devices, typically hand-held, that allow voters with vision impairments to vote independently.

¹⁵ See <http://www.votersunite.org/info/BallotProgramming.pdf> for a detailed discussion of BDFs.

¹⁶ Quoting Thompson, "You can't trust code that you did not totally create yourself. (Especially code from companies that employ people like me.) No amount of source-level verification or scrutiny will protect you from using untrusted code. In demonstrating the possibility of this kind of attack, I picked on the C compiler. I could have picked on any program-handling program such as an assembler, a loader, or even hardware microcode. As the level of program gets lower, these bugs will be harder and harder to detect. A well installed microcode bug will be almost impossible to detect."

DREs do not allow voters to select more candidates than allowed (overvotes) and alert voters to omitted votes (undervotes). They also allow voters to review their ballots before submitting them (second chance voting).

DREs that produce voter verifiable paper ballots. AccuPoll and Avante produce DRE voting systems that print out ballots that voters can check to ensure that an accurate paper record of their votes exists. Avante also manufactures an optical scan model that prints optical scan ballots that sighted voters can mark, as well as an “accessible” optical voting system that allows vision-impaired voters to print out optical scan ballots marked to reflect their choices.

Optical scan voting machines. Besides avoiding many of the security problems associated with paperless DREs, optical scan (or mark sense) systems are also less expensive. Typically these systems require the voter to mark his or her ballot, in much the same way that students taking standardized tests make computer-readable marks by using number 2 pencils to fill in ovals.

Precinct-based optical scan systems require the voter to “test” his or her ballot by submitting it to the scanner and having the scanner notify the voter if the ballot contains overvotes. The voter is also notified if the ballot is blank. Ideally, at the end of Election Day all the ballots are initially tallied in the precinct, and the ballots, together with the results, are sent to the tabulation center.¹⁷

The same vendors that produce the majority of DREs – ES&S, Sequoia, and Diebold – also produce the majority of optical scan voting systems.

Hybrid models. Ballot marking systems are a cross between DREs and optical scan systems. One, made by Vogue Election Systems (VES) and currently marketed by ES&S, offers a touch screen like a DRE. The voter inserts a blank optical scan ballot into the machine and then proceeds as he or she would if interacting with a DRE. Once the voter has entered all of his or her choices, the machine marks the optical scan ballot accordingly, avoiding overvotes and raising alerts to undervotes in the process. This also serves to eliminate any stray pencil marks that could otherwise confuse the scanner. Attached headphones provide an option that allow blind voters to vote without assistance.

Another system, produced by Populex, includes a screen that operates with an attached stylus. The system prints out a completed ballot once the voter has entered all of his or her choices. For human perusal, the ballot uses numbers to represent voter choices, along with a corresponding bar code for the optical scanner’s benefit. The system has attached headphones that allow blind voters to vote independently, and, like the Vogue system, it also avoids overvotes and warns about undervotes. For both systems headphones

¹⁷ The chance that ballot boxes or tabulation sheets will be illegally manipulated are reduced if local results are posted locally.

attached to the scanner would make it possible for vision-impaired voters, as well as the sighted, to verify their ballots.¹⁸

Because paperless DREs provide no audit trail, it is imperative that DRE software be free of malicious code and potentially damaging bugs. By contrast with paperless DREs, DREs that produce voter verifiable paper ballots, optical scan systems, and hybrid systems do not have the hidden expense of a huge testing and security overhead.

Cryptographic voting systems. Both VoteHere (<http://www.votehere.net/>) and David Chaum (<http://www.seas.gwu.edu/~poorvi/Chaum/chaum.pdf>) have developed voting systems that provide an encrypted receipt that voters can use to verify that their ballots has been accurately counted. Chaum's system is not currently being manufactured. A problem common to both systems is that they offer no way to conduct a recount should it be determined that a ballot tabulation problem has occurred, although individual ballots can be corrected. Also, neither scheme is particularly easy for voters to understand.

Open source. The Open Voting Consortium (OVC) (<http://www.openvotingconsortium.org/>) is a non-profit group of software engineers and computer scientists working to build an open source voting system that will run on PC hardware and produce a voter-verifiable paper ballot. They also hope to provide a general for interoperable open source voting software. Their system is currently under development.

Prudent precautionary measures for DREs

Because paperless DREs provide no audit trail, it's imperative that they be extensively tested before, during, and after each election. DREs must also be securely stored *between* elections, as well as at polling sites before and during Election Day.

DREs should be extensively tested before, during, and after every election. Similarly, all ballot definition files should be scrupulously tested — with all test results (not just results from BDF tests) not only made public but also archived in a central repository. There should also be a national repository of DRE problems, just as is done with aircraft.

Finally, paper ballots should be made available at every polling location that uses DREs, both as backup in the case of failures of the DREs and to provide voters with the option of voter-verifiable paper ballots,

None of these steps can ensure that DRE software is free of malicious code and potentially damaging bugs. The best we can do is to attempt to reduce the risks associated with these machines.

Conclusion

¹⁸ This option is not currently available.

The issue of e-voting should have been primarily a technological issue - one involving computer security, human factors, reliability, and efficiency. Instead, because of the vast sums of money involved, e-voting has been heavily politicized.

Election officials were told that DREs in the long run would be cheaper than alternative voting systems. They were told that DREs had been extensively tested and that the certification process guaranteed that the machines were reliable and secure. No mention was made of the significant costs of testing and of secure storage of DREs; no mention was made of the inadequacy of the testing and certification processes, to say nothing of the difficulty of creating bug-free software.

Technologists are attempting to educate election officials, policy makers, and the public about the risks of paperless DREs. It is critical for the continued existence of democracy throughout the world that we succeed.

Acknowledgments.

Thanks to Dan Wallach, Tracy Volz, Laura Gould, Lynn Landes, Ellen Theisen, Rebecca Mercuri, and Doug Jones for their very useful comments.