Democracy at Risk: The 2004 Election in Ohio

Section VII
Electronic Voting: Accuracy, Accessibility and Fraud

DEMOCRATIC NATIONAL COMMITTEE

VOTING RIGHTS INSTITUTE

Electronic Voting: Accuracy, Accessibility, and Fraud

Dan S. Wallach
Associate Professor
Department of Computer Science
Rice University

We have seen a number of studies of electronic voting systems over the past several years, roughly broken into two camps: computer scientists and statisticians. Among computer scientists, we have studies performed by academics and by a number of different testing organizations, many of which found significant flaws in the design and implementation of electronic voting systems. Among statisticians, we have studies of voting residual rates, turnout, and other important issues, many of which have concluded that new DRE voting systems are less accurate than more traditional optical scan ballots. This report considers many of the issues raised by these studies and some of our observations from the presidential election in Ohio in November 2004.

Incident reports and machine accuracy

A common feature can be observed in many "problem reports" from DRE voters. They will claim that they selected one candidate and then observed a "switch" of some kind to a different candidate. Inevitably, these problems are difficult or impossible to reproduce, and could be caused by problems with the engineering of voting systems, or could be exacerbated by a perception of machine inaccuracy. Unfortunately, *we have no baseline data* on how accurate DRE systems (or, really, any voting systems) are at capturing voter intent. Proper scientific studies would bring would-be voters into a controlled environment on a non-election day; they were asked to vote for their candidates and were videotaped while voting (no privacy being necessary for such an experiment because there would not be an actual election). The voters' input to the machine could be compared with a spoken survey after the fact, or otherwise corroborated with other factors. Such a study would determine a true, baseline *human error rate*. Most interestingly, such a study would help determine how many errors result from *calibration* errors[1], a common source of anxiety with current DRE systems. Today, the best we can measure are *residual vote rates*, that is, we can count how many ballots are cast with some races left blank ("undervotes") or with multiple selections on a given race ("overvotes"). Many studies of residual voting rates compared to voting technologies, including the DNC's study of Ohio, have shown that the lowest residual vote rates occur with *precinct-based optical scan systems*. In such systems, voters mark a plain paper ballot with a pen. A computerized scanner, mounted above the ballot box, will reject

[1] In typical commercial touch-screen systems, a layer of glass or plastic is placed above the actual screen to detect finger contact. Because some voters are taller and others shorter, every voter will have a different angle from their eye to the finger to the screen below. "Calibration" can be done for a "typical" voter height, but can never be perfect for all possible voters. Typically, buttons are drawn onscreen much larger than a human finger to minimize such errors.

overvoted ballots, eliminating a common mode of human error and giving voters a chance to restate their intent.

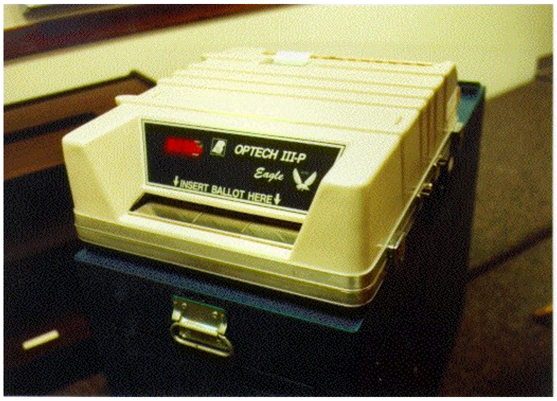

Figure 1: A Precinct-based optical scan system (the ES&S Optech Eagle)

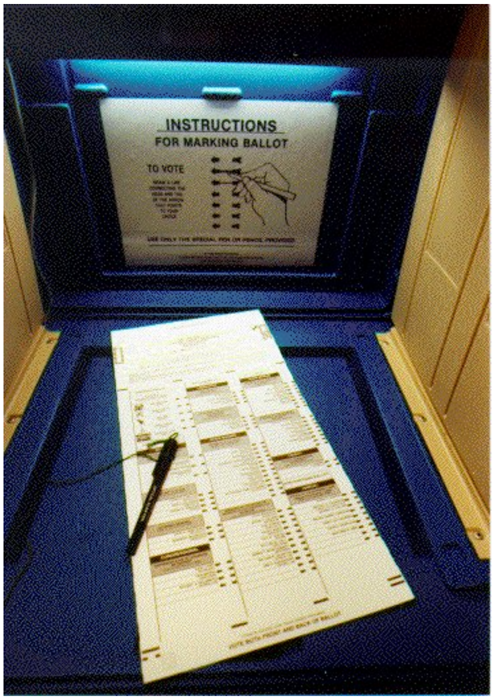

Figure 2: An optical scan ballot for the ES&S Eagle

In all of these studies, DRE systems are consistently shown to have higher residual vote rates than optical scan systems, even though all commercial DRE systems are engineered to simply prevent overvoting (when you select a second candidate for a race, the DRE will de-select the first candidate). This suggests that many voters are more capable of expressing their preferences accurately traditional optical scan systems than to newer DRE systems.

Accessibility

Based on findings like this, an obvious recommendation would be to eliminate DRE systems and go strictly with precinct-based optical scan systems. They're more accurate, significantly cheaper, and offer significant benefits in terms of election transparency and resistance to wholesale election fraud (more on that later). Unfortunately, they're not accessible to several different populations of voters. Voters with low vision may be unable to read the small type that is often necessary to list all of the candidates on a relatively small piece of paper. Voters with zero vision (i.e., blind voters) cannot use optical scan systems whatsoever without assistance, either from an electronic system or a human assistant. Voters with low motor control might have difficulty using the pen to mark the paper ballot and to deposit their marked ballots into the ballot box. And, voters may be illiterate or may not be fluent in English.

DRE systems are often touted as the solution to accessibility needs in the polling place. HAVA requires all U.S. voting precincts offer “accessible” for elections subsequent to January 2006. Today’s DRE systems satisfy these accessibility concerns with a variety of add-on devices, including touch-pads and headphones as well as “sip and puff” input devices.

Election Fraud

A primary concern of any election system, whether done by hand, via computer, or any other mechanism is that *it must provide sufficient evidence to convince the losing candidate that he or she actually lost*. Naming the winner is the easy part. When we talk about *evidence*, however, we bring up all the same issues that might occur in a criminal investigation, including tampering (either by insiders or outsiders) and maintenance of a proper chain of custody over the evidence.

Vote by Mail

A simple system to first consider is voting by mail. Virtually all ballots in Oregon are cast by mail, and a significant number are cast in many other states. Mail-in votes are trivially subject to *bribery* or *coercion* (either “I’ll pay you $10 for your vote” or “I’ll break your kneecaps if you don’t give me your vote”) at the level of individual voters. This would become expensive to perform at a large scale, particularly without knowledge of the fraud becoming public. To perform such fraud at a *wholesale* level, where a small number of people might attempt to damage the system is far more difficult. A corrupt mail courier could only tamper with the ballots that he or she personally handled, and tamper-resistant features on the ballot or envelope might make such tampering hard to disguise. Once the ballots arrive at the central tabulation facility, fewer people would need to be involved, but hopefully stronger security measures are in place to prevent such fraud. If, for example, ballot envelopes are counted before even being opened, then those counts could be compared, in batches, to the tallies after the batches are scanned and processed. Such measures are comparable to *separation of duty* techniques common in the banking industry, where no one employee can ever embezzle funds without another employee discovering the missing funds as part of their job.

Precinct-based optical scan

Precinct-based optical scan systems compare favorably to vote-by-mail systems. Because the voter must vote privately in a (hopefully) well-controlled polling place, coercion and bribery don’t work. The precinct ballot scanner catches overvoting and allows the voter to try again, a feature not possible with mail ballots. The scanner also keeps its own tally of the votes, which can be rapidly transmitted over a modem or spoken over a telephone. Printouts can be physically signed by precinct-level voting officials, and independently tabulated by interest groups that are willing to send representatives to each precinct. This provides an important hedge against the risk of ballot box tampering, particularly while the ballot boxes are in transit from the local precinct to some form of central storage (probably the single greatest vulnerability in any paper-based election system). However, a significant risk remains. What if the *software* inside the scanner incorrectly tabulated the ballots? No election observer would be able to independently count the ballots themselves. Likewise, precinct-level election officials generally do not (and certainly should not) handle ballots after they are cast. The risk of software error might result from software bugs, or could possible be the result of *fraudulent programming* (sometimes referred to as a *Trojan horse*). Today's certification and "logic and accuracy testing" are completely insufficient to detect such problems[2]. However, so long as the paper ballots are handled properly, they will remain, after the election, allowing for a meaningful recount. *The ability to perform such a recount provides a critical hedge against the risk of scanner failures.*

DRE voting systems

Direct Recording Electronic (DRE) voting systems offer a number of benefits relative to precinct-based optical scan systems. They also introduce significant new complexity, new risks, and new costs. A DRE terminal may cost thousands of dollars, and many must be purchased to allow busy precincts to limit voter waiting times to avoid the problems observed, for example, in Franklin County, Ohio.

Modern DREs are, at their core, general-purpose programmable computers. Some even run Microsoft's Windows CE operating system. This gives DREs the flexibility to support a variety of attractive features including large text, speech synthesizers, and multiple languages, all of which help making voting accessible to a wider demographic of voters. This same flexibility, unfortunately, significantly increases the ease with which someone might tamper with the software. Such tampering could occur where the machine was manufactured or anywhere else from the moment the machine leaves its manufacturer to the day of the election. Anyone who has uninterrupted physical access to a DRE voting system for any length of time could potentially tamper with its software. Consider software updates. As with normal consumer software vendors, DRE vendors are constantly improving and modifying their software to satisfy the needs of their customers. They then submit this software for "certification" by an Independent Testing Authority. There are three U.S. companies currently serving as Independent Testing Authorities. However, in cases where outside computer security firms or academics have had the opportunity to independently examine DRE software, they have found significant and wide-ranging flaws. As such, it appears that the ITAs do not have the skills to properly audit voting system software. We also observe that ITAs make no warrant that voting systems are actually suitable for use in an election. Rather, much more weakly, they claim that voting systems "satisfy FEC standards", which unfortunately require almost nothing with regard to software quality or security, or even about usability or accuracy. More elaborate standards are in development, but are nowhere near adoption.

A fundamental attribute of all modern DRE systems is their elimination of the paper trail we have with optical scan systems. While these systems will allow voting totals, or even individual votes in some cases, to be printed at the end of the election, this does not

[2] Logic and accuracy testing for an optical scanner generally involves running a "test deck" through the machine. After scanning the deck, the tally is read from the machine. The scanner's tally can be compared to the known totals. Unfortunately, a well-designed Trojan horse can tell when it's being tested, either by identifying that, in fact, it's seeing the same test deck it always sees, or even by observing that the test ballots are arriving much faster than "normal" voters might cast their ballots.

provide a hedge against software failures in the DRE. It’s entirely possible that a DRE voter could vote for one candidate, which would be displayed on screen, while an entirely different candidate could be recorded internally as having received that vote. If such an error occurred, neither the voter nor any election official would be able to undo the damage after the fact. If such an error occurred systematically, it could swing the outcome of an election. And, if the faulty software was deliberately placed in the machine, it could even be programmed to modify itself to eliminate any traces of its having been present. *If such fraud were occurring, it would not be visible to poll workers or election observers.*

As with any other voting system, DRE votes must ultimately be centrally tabulated. This information may be communicated over a modem or carried by hand in a computer memory card. As with traditional ballot boxes, such data may be subject to tampering while in transit. However, while ballot boxes are large objects that can be easily observed and tracked, computer memory cards are small and sleight-of-hand can allow for quick substitutions. Likewise, telephone lines are not terribly secure against attackers who can climb telephone poles. While appropriate cryptographic techniques can mitigate against all of these risks, many DRE vendors either use no cryptography at all or do it improperly, leaving the data effectively unprotected while in transit. Once the data arrives at the central tabulation facility, it is typically stored in off-the-shelf personal computers running a Microsoft operating system and some form of database. These computers, themselves, may be subject to attack by election insiders. Anyone with physical access to these computers and the appropriate tools could execute a database script to directly modify the database records, overwriting any original data without leaving any evidence of such tampering. Furthermore, in the case that these machines are *ever* connected to the Internet, perhaps to deliver results to an election web server or to the press, these machines could be attacked over the Internet. Even if all the latest security patches have been applied, attackers may well keep other security attacks in reserve, specifically to attack such election computers.

Internet voting systems

The Department of Defense commissioned a voting system to allow overseas soldiers to cast their votes on the Internet, using web browsers and other off-the-shelf components available, even in remote locations. A report, written by several experts asked to study this system, concluded that both the end-user computers and the central tabulation machines were fundamentally at risk of security attack. Present software technology is not good enough that we can make any guarantees about such systems’ robustness against attack. And, if such a system were deployed, adversaries ranging from disaffected local voters to foreign intelligence services would have incentives and opportunities to go after the system. The Department of Defense scrapped the project.

While many other attempts to introduce non-traditional voting schemes may increase voter turnout by making it easier to vote, they introduce significant risks along these lines. Any opportunity for an attacker to electronically communicate with either a voter or the tabulation facility makes it easier than ever before to perform election fraud.

Likewise, such systems have all of the same bribery and coercion issues present in vote-by-mail systems.

Voter-verifiable paper trails and other DRE improvements

A number of proposals have emerged from the computer science community to improve the security and robustness of DRE-like voting systems. The simplest proposal is to attach some form of printer to a DRE system. Voters would use the same computerized user interface as before. However, when voters indicate that they are done, a printer would generate a printed representation of their ballot. Voters could read this ballot and, if they agree, it would become the official ballot, the primary record of their voting intent. The DRE system could keep its own internal tally, but as with precinct-based optical scanners, the paper records would take precedence in a recount. There are many variants on voter-verifiable schemes. One variant, the "Mercuri method," holds the ballot under glass. Voters can read it but cannot touch it. This defeats a vote-buying scheme called chain voting[3], and also prevents voters from accidentally removing ballots from the polling place. Another variant simply uses a computer to mark a traditional optical scan paper ballot which is then deposited into a standard ballot box (see Figure 3). An intriguing benefit of such systems is that only one per precinct needs to be purchased to satisfy HAVA requirements. Voters who need the accessibility features of DRE systems can use them, and voters who do not can use standard pens. With limited budgets, this becomes an attractive option for many counties, particularly those already using optical scan voting systems.

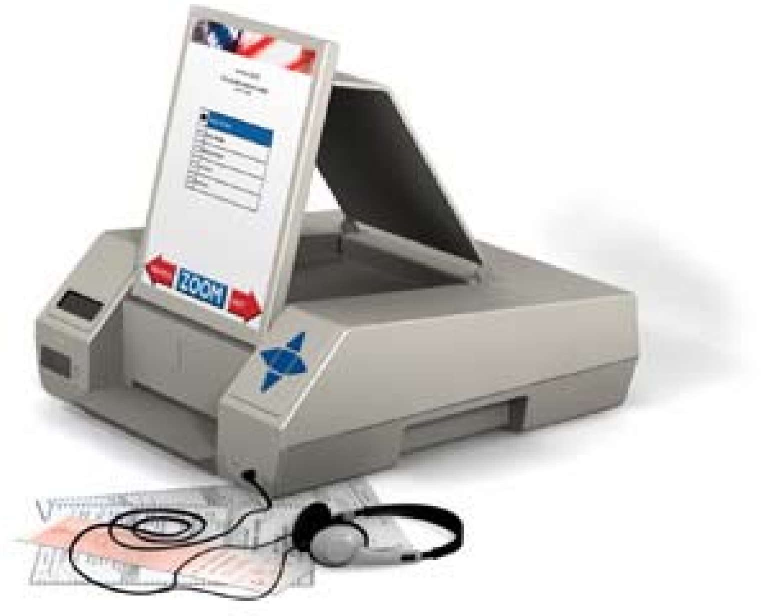

Figure 3: ES&S / Vogue Automark (computer-assisted optical scan ballot marking device)

[3] A typical chain voting attack on a paper ballot system has the attacker standing outside the polling place, offering to buy votes. A voter who wishes to sell a vote is given a ballot, already marked by the attacker and is told to pocket this ballot, go get a fresh one, and swap them. The previously-marked ballot is deposited in the box, and the fresh, unmarked ballot is returned to the attacker for the payment.

Computer scientists and cryptographers have also developed a variety of intriguing cryptographic schemes using advanced mathematical techniques to allow voters to go home with just enough numerical evidence that they can verify their vote is part of the final tally without being able to prove to a third party what their vote actually was. Such schemes generally allow independent third parties to perform their own tallies of the election, based again on cryptographic evidence. To date, such schemes have not been used in any elections and questions remain about both whether the cryptographic schemes can be broken and whether these systems would be usable by the broad voting population.

Recommendations

- Precinct-based optical scan systems are the most "accurate" voting systems available today. They are also reasonably priced and can satisfy HAVA requirements in a cost-effective manner with devices such as the ES&S AutoMark (see Figure 3).
- Current DRE systems are not engineered to meet the needs of elections. They are extremely expensive to procure and maintain. They are not sufficiently robust against fraud. They are less usable to the broad population of voters than earlier, simpler technologies.
- Existing standards and practices for the certification of voting systems are insufficient to the security requirements of DRE systems. Significant effort will be needed to create the next generation of standards.
- Few quantitative studies have been performed on the usability of different voting technologies. Vendor claims of improved usability should not be considered meaningful until they perform significant user studies under controlled conditions. Existing anecdotal evidence, including event reports, are at best mixed in their opinions of different voting systems' usability. Election official should perform controlled, scientific studies of their own populations using their own voting machines to truly understand where they might be experiencing usability problems.
- Most voting system vendors consider their software to be proprietary trade secrets and generally resist any attempts to disclose and discuss their designs in public. *Private, vendor trade secrets have no place in public elections*. Vendors are welcome to protect their intellectual property with copyrights and patents, but their full designs must be subject to public scrutiny. As elections become increasingly electronic, such scrutiny is critical to maintaining transparency and public confidence in elections.
- Computer software, at every stage in the process, might be buggy and could well be malicious. Different strategies are necessary to mitigate against this threat, depending on what voting system is used.
    - Paperless DRE voting systems generally print precinct-level tallies at the end of the election. These printouts are generally signed by the election officials working in the precinct. Those signed printouts should be treated

as important evidence as to the result of the election and should be preserved for recounts and post-election auditing.

  - Precinct-level optical scanners might incorrectly tally votes as well. The original marked ballots should be independently counted, or at least randomly sampled and compared to the electronic results, before an election result is certified.
  - Paperless DRE systems should be upgraded to voter-verified paper trail systems. The printouts should be treated in exactly the same fashion as optical scan ballots: they should be carefully preserved as evidence of voter intent and should be randomly sampled and compared to the electronic results.
  - “Parallel testing,” where some DRE voting systems are pulled out of general use and are tested, on election day but under controlled conditions, is an pragmatic and valuable test that should be performed whenever such voting machines are being used.
  - The computers used to tabulate election results are a tempting target for election fraud, and as such, require more significant controls, including well-chosen passwords and physical access restrictions. They should never, in their entire lifetime, be connected to the Internet or to any modem or communication device. Instead, an “air gap” style of security should be used. Data can be released to the public through simple measures such as burning a CD with election results and hand-carrying such a CD to a separate, network-enabled computer.
- Election officials need to hire “penetration testing” (also called “tiger team”) consultants to examine the security of their election systems. Where such teams have been hired in the past, significant vulnerabilities have been discovered. Such teams should be hired on a recurring basis to audit voting machines as well as the entire voting process, from registration through tabulation.
- The timely publication of detailed precinct-level election statistics is critical to the public confidence in an election result, and such data is often not available in its entirety for every county. Such statistics can be easily derived from local voting tabulation systems and should be quickly and electronically reported in a standardized fashion.