

Voting System Risk Assessment via Computational Complexity Analysis
Dan S. Wallach
Department of Computer Science
Rice University

August 19, 2008

This is a pre-print. The final article will appear in a special issue of the William & Mary Bill of Rights Journal, Volume 19 (forthcoming Dec., 2008).

Abstract

Any voting system must be designed to resist a variety of failures, ranging from inadvertent misconfiguration to intentional tampering. The problem with conducting analyses of these issues, particular across widely divergent technologies, is that it's very difficult to make apples-to-apples comparisons. This paper considers the use of a standard technique used in the analysis of algorithms, namely complexity analysis with its "big-O" notation, which can provide a high-level abstraction that allows for direct comparisons across voting systems. We avoid the need for making unreliable estimates of the probability a system might be hacked or of the cost of bribing key players in the election process to assist in an attack. Instead, we will consider attacks from the perspective of how they scale with the size of an election. We distinguish attacks by whether they require effort proportional to the number of voters, effort proportional to the number of poll workers, or a constant amount of effort in order to influence every vote in a county. Attacks requiring proportionately less effort are correspondingly more powerful and thus require more attention to countermeasures and mitigation strategies. We perform this analysis on a variety of voting systems in their full procedural context, including optical scanned paper ballots, electronic voting systems, both with and without paper trails, internet-based voting schemes, and future cryptographic techniques.

1. Introduction

The U.S. Elections Assistance Commission recently solicited submissions for how it might assess the risks of voting systems [EAC08]. According to their solicitation:

The first phase will create reference models to be used in the assessment. This includes developing election process models to describe the operational context in which voting systems are used. It also entails developing voting systems models by generic technology type. This is needed because the types of threats encountered and their potential impacts vary by technology.

The EAC asked the public to suggest how they might develop these models, with submissions due in April 2008. While these submissions have not yet been made available to the public, we will discuss some prior work on this topic and then propose our own solution to this problem.

Clearly, we need an objective, quantifiable method for comparing voting systems. Election officials who might purchase one system over another need to be able to concisely understand the relative insecurities of one product versus another. Security analysts, testing authorities, and regulators need common ground for both setting a lower bound on acceptable security and for explaining how much better a system is than whatever the minimum standard requires.

Qualitative analyses are, for better or worse, the standard method used to make arguments. To pick a well-known example, the U.S. military was investigating the possibility of allowing its soldiers to vote, from overseas locations, via the internet. They convened a panel of experts to conduct a security review. Several of them wrote a “minority report” expressing their concerns with the project [JRSW04], leading to its cancellation and replacement with a fax-based system. Alvarez and Hall [AH08] criticize this outcome (“In the end, a small but vocal segment of the scientific community opposed the use of scientific experimentation in voting systems and technologies.” – p. 77). While the SERVE report’s authors were concerned about the fundamental unsuitability of standard consumer platforms (e.g., Microsoft Windows XP plus Internet Explorer), with the risks of viruses, worms, or other forms of malware that could easily be engineered to compromise an election, Alvarez and Hall felt that “the central argument in this critique was overly general, ignored the reality of UOCAVA voting¹, and ignored what would be a broad array of project, procedural, and architectural details of the SERVE registration and voting system, which in all likelihood would have minimized or mitigated their concerns had the system been used in the planned trial.” [AH08, p. 85]

The SERVE report’s authors are absolutely correct in that the weaknesses of internet voting are a fundamental properties of the how internet clients and servers operate. However, Alvarez and Hall have a valid point when they later note that other overseas votes are cast by fax, with its own attendant risks. With disputes like these, we need a useful way of getting to some kind of ground truth.

2. Overview of Risk Assessment Frameworks

a. Security Metrics

There are many different ways to quantify how secure or insecure a voting system (or any system) might be, but most of them require us to make approximations and guesses, among other problems. Burstein et al. [BHMW08], in their submission to the EAC request, make the point that there may well be no one-size-fits-all solution to this problem. Nonetheless, we can consider a variety of possibilities.

Threat taxonomies. Jones [J05] discusses the possibility of cataloging known threats against voting systems and creating a taxonomy of those threats. The main benefit of an activity like this is that it would help security analysts identify areas where additional analysis might be desirable. The presence or absence of any particular flaw does not, in and of itself, make it any easier or harder to attack a voting system. An attacker only needs to find one serious flaw in order to cause significant damage.

Voter performance metrics. Voters are regularly surveyed both before and after elections as to whom they support for a political office. Voters are also surveyed, perhaps more rarely, about how well they enjoyed the voting experience, how easily they could find their local precinct, and so forth. Ultimately, any survey or scientific study of voters can measure three useful properties of a voting experience: accuracy, efficiency, and satisfaction. Accuracy refers to how well a voting system captures the voter’s intent. While “residual votes” (e.g., spoiled ballots) offer a proxy measurement for accuracy, these measurements can also be taken in controlled laboratory conditions where experimental subjects have no anonymity (nor need for it), and thus their actual votes can be examined directly. Efficiency is simply the amount of time it takes to complete the voting procedure. Accuracy and efficiency can be objectively measured, while satisfaction relies on the voter’s subjective opinions, stated in response to standardized questions. All of

¹ The Uniformed and Overseas Citizens Absentee Voting Act (UOCAVA) was enacted by Congress in 1986. The UOCAVA requires that the states and territories allow certain groups of citizens to register and vote absentee in elections for Federal offices. In addition, most states and territories have their own laws allowing citizens covered by the UOCAVA to register and vote absentee in state and local elections as well. Details can be found at the Department of Justice’s web site: http://www.usdoj.gov/crt/voting/misc/activ_uoc.htm

these dimensions of voter performance are quantifiable, and provide valuable metrics for examining how well a voting system performs.

Unfortunately, none of these performance metrics have a strong relationship to security. Two electronic voting systems could appear absolutely identical, from the voter's perspective, while having completely different underpinnings with radically different security vulnerabilities. A malicious voting machine could change a vote without detection, and the voter may well love the experience. Voter performance metrics do matter in one critical fashion: voting system security procedures often depend on voters or poll workers following procedures correctly. For example, voters should verify the output of VVPAT printers to ensure that it matches their intent. If voters cannot do this, then the security mechanism's effectiveness is blunted. We will discuss usability issues in more detail when we discuss particular voting systems.

Probabilities of failure. We can measure so many things about elections, but we cannot measure the odds that a security attack will occur. Election security in the past can never be considered to be a good predictor of election security in the future. For example, new vulnerabilities may be discovered which change the security landscape overnight. Also, if insecure systems were used and successfully compromised, we would not necessarily have any evidence that the compromise had occurred. The lack of evidence for prior security attacks against electronic voting machines simply has no bearing on the degree to which they are vulnerable or might be compromised in the future.

Dollar or time-costs. A tempting metric is to attempt to quantify the dollar cost or man-hour cost of a security attack. This becomes even more attractive if the cost can be expressed in terms of dollars per vote flipped or man-hours per vote flipped because it can then be compared against other forms of manipulating voters' behavior (e.g., television advertisements, direct mail, or phone calls). We could even talk about the dollars or man-hours that must be invested "up front," perhaps as part of a reverse-engineering process conducted by an attacker, followed by a "per vote" cost afterward. These metrics are somewhat satisfying, in that they can be easily compared across different voting technologies, but there will be huge errors in any estimates. Will it take one month or three months to reverse-engineer a voting machine and find an exploitable security bug? Imagine that the bug in question allows an attacker to violate voters' anonymity but not flip their votes. What, then, is the cost per-voter for bribing them to vote in a particular fashion? A related metric, raised by the Brennan Center Security Report [N06], is the number of people who must be involved in a conspiracy to compromise a voting system. If more conspirators are required, it would be that much harder for the attack to go off without detection. As with dollar or man-hour costs, it's difficult to estimate these numbers accurately, except when only a single person is necessary to accomplish an attack. In many scenarios, this is unrealistic.

Complexity analysis. If we reject all those metrics, then we need an alternative. For this, we turn to a technique used widely in the analysis of algorithms: complexity analysis. Consider the case where we are given a pile of insurance forms, each coded with the client's customer ID number. We are assigned to take this unsorted pile and sort it such that the customer IDs are strictly increasing through the pile². We might start a new pile, on the side, in sorted order. Then, for each page from the unsorted pile, we flip through the sorted pile to find where it belongs. This first technique, called an *insertion sort*, will take time proportional to the number of pages (N), squared, as we have to insert N pages, and each might require us to flip through the whole stack, which can have up to N pages in it. We might instead start by sorting by the *last* digit in the ID (the one's digit) into ten separate buckets, putting everything back together again, then sorting again by the second digit (the ten's digit), and so forth. This second technique, called a *radix sort*, takes time proportional to the product of the number of pages and the number of digits. Another technique, cleverly called a *quick sort*, has us randomly pick a "pivot," separating the pages into two piles (those less than and those greater than the pivot). Then for each smaller pile, we pick another pivot and repeat. That process turns out to take an expected time proportional to the number of pages multiplied by the logarithm of the number of pages.

² A mind-numbing summer job performed by the author while he was in junior high school.

The benefit of this style of analysis, called complexity analysis, is that we can take very different procedures for sorting a pile of paper and focus on a very high level consideration of the effort we might expend. We don't have to worry about whether it's slightly faster to shuffle the pages in a radix sort or in a quick sort. We don't even worry about the variations that might occur from one run of the procedure to the next. Instead, we just roughly count up the number of operations that we need to perform in proportion to the size of the problem.

This style of analysis is typically expressed with “big-O” notation. Formally,
 $O(g(n))$ = a function $f(n)$, such that there exist positive constants c and n_0
such that $0 \leq f(n) \leq cg(n)$ for all $n \geq n_0$.

In other words, if we say an algorithm is $O(N^2)$ (pronounced “order N -squared”), we're saying that the cost of the algorithm approaches some constant times N^2 . If, however, N is small, then we make no guarantees, but we don't really care because the running time of the algorithm won't be big enough to matter. In this framework, we can say that insertion sort is $O(N^2)$, radix sort with k digit-numbers is $O(kN)$, and quick sort is $O(N \log N)$. Among other clever properties, you can add these O-expressions together, and the more expensive operation will dominate the cheaper one. For example, if you have to first pick up the pile of pages from a mess on the floor, an $O(N)$ process, then quick-sort them, taking $O(N \log N)$ time, the sorting time will dominate the pick-up time, so the total cost $O(N) + O(N \log N)$ is equivalent to $O(N \log N)$. This sort of complexity analysis is fundamental to the theory of computer science, and detailed treatments of it can be found in most any CS theory textbook (see, for example, Cormen et al. [CLRS01], chapters 2-3).

Similarly, big-O notation is a convenient way to evaluate election security. We can refer to attack costs in terms of the number of precincts or poll workers³ (P), or the number of voters (N), among other possibilities⁴. This style of analysis tends to discount the up-front cost of discovering a bug (it's a constant, even if it's a big constant) while focusing on the scalability of the attack. Big-O analyses provide a more precise characterization for what are commonly referred to as “retail” versus “wholesale” attacks. If an attack requires you to bribe each voter individually, we would say it's an $O(N)$ attack, regardless of the number of attackers required to accomplish the attack.

Our model must indicate how we deal with scalability. If we have P poll workers attacking N ballots, then each poll worker might well perform $O(N/P)$ work. Ultimately, we're interested in the incremental cost of corrupting N ballots. If each poll worker already handles N/P ballots, and the incremental cost of their corrupt behavior is negligible, then we would prefer to say that there is a total $O(P)$ cost in corrupting N ballots. For example, you may need to get all P poll workers involved if you want to tamper with N ballots while they are still in the precinct. In this case, the dominating cost of the election fraud is getting the poll workers on board, while the incremental work per poll worker is negligible. In this way, our complexity analysis can subsume the Brennan Center Security Report's notion of counting the people required [N06], when that's the limiting factor.

The strongest possible attacks would then be $O(1)$, “constant” cost attacks. If one or a handful of attackers can change an arbitrary number of ballots, with minimal effort, that's as bad as an attack could

³ Since there are a presumably constant number of poll workers per precinct, the O-expressions will be exactly the same in our analysis whether P refers to the number of poll workers or precincts. All that would change is the constant c , which disappears when we do this style of analysis.

⁴ While one could argue that the number of poll workers is typically a constant fraction of the population (e.g., one poll worker for every hundred voters), and thus $O(P) = O(N)$, the two order-of-magnitude difference between the sizes of these populations has a real impact on the viability of attacks conducted by poll workers versus attacks against individual voters. It's helpful if our analysis can distinguish corruption requiring P complicit poll workers versus N complicit voters.

possibly be. When we identify $O(1)$ attacks on a voting system, that doesn't mean the voting system is necessarily unsuitable for use, but it means that extreme procedural measures must be taken to thwart such attacks. If such measures are infeasible, then the voting system really should not be used.

b. Attacker Goals

Before we conduct complexity analyses on specific voting systems, we first need to consider what an attacker's goals might be. The standard goals of any attacker in any security-relevant context are typically attacks against integrity, confidentiality, and availability. Applying these to elections:

Integrity attacks aim to change the election totals. A successful integrity attack is one that survives tallying, recounting, and any subsequent legal challenges. A simple integrity attack is stuffing ballots into a paper ballot box, combined with forged signatures in a poll book, to keep everything consistent. An integrity attack, in this definition, must go undetected, lest the results be overturned or the election redone from scratch.

Confidentiality or anonymity attacks are a different game. If we can violate the anonymity of voters through whatever means, then we can potentially bribe or coerce them to vote however we want. In this case, the election tallies may be entirely accurate with respect to capturing the voters' stated intent, but the election results may nonetheless not represent the true preferences of the voters.

Availability attacks (also called denial-of-service attacks or DoS) are attempts to disrupt the election and can be as simple as physically destroying a voting system, stealing a ballot box, or cutting power to a building. DoS attacks, selectively applied, may introduce biases that effect the election outcome, but they make no attempt at stealth (e.g., destroying ballot boxes from precincts with known biases). A DoS attack may have as its goal to simply cause a chaotic outcome, or no outcome at all. Or, a DoS attack may be acceptable to the attacker if the whole world knows the results are corrupt but regular election procedures will have no way to correct the errors and might then accept the faulty results. In this respect, an integrity attack that is discovered but cannot be undone would, by our definition, be a DoS attack. (Perhaps such attacks could also be classified as "integrity" attacks, but the victor in such a race would lack some of the legitimacy that can come from a "clean" victory.)

When we consider these classes of election attacks against various election systems, we must make a number of simplifying assumptions. We will assume that an attack is focused on a specific county, voting uniformly on a single technology. An attack against one county would not have an impact on surrounding counties⁵. Instead, each county would need to be separately attacked. (Clearly, the cost of attacking a county with two separate systems is equal to the sum of the costs for attacking the individual systems.) We will also assume that prospective attackers already know the design and implementation of particular voting technologies. Security must come either from the strengths of those designs or from the strengths of poll workers and election officials following procedures correctly.

For every attack we consider, we will consider three variants: attacks originating from "normal" voters; attacks originating from poll workers, postal employees, and other intermediaries; and attacks originating from insiders at the election headquarters. Corrupt insiders obviously pose a greater threat than corrupt voters, requiring correspondingly stronger mechanisms to mitigate against the possibility of such corruption.

⁵ In Georgia and Maryland, the whole state votes monolithically on the same electronic voting system, with centralized resources to maintain their machines. In a scenario like this, a simplification is to treat the state as if it were a single, large county. For constant-cost attacks, this is likely to be an accurate model.

For this analysis, we do not consider pattern voting attacks (also sometimes called the “Italian attack”), where a voter uses unimportant races to encode a unique ID number, then votes as he or she has been bribed to do on the important races. To the extent that this attack is a problem, it’s a problem on absolutely any voting system that can produce complete ballot records. In schemes where nobody ever sees the raw ballots (including VoteBox [SKW08], discussed in more detail below), this would never be a viable attack.

c. Example Cases

Under this set of attacker goals, there are some ambiguous cases that are worth considering. In Zimbabwe’s recent presidential election, Robert Mugabe initially lost his reelection bid to Morgan Tsvangirai [B08] and then won a revote where his opponent withdrew amid widespread state-sponsored violence [DB08]. In this case, voters were still asked to vote, but there’s no question that the will of the people was not reflected in the outcome. The violence was inflicted upon large swathes of the Zimbabwe population, so we would consider it an $O(N)$ attack. Certainly, there was no stealth or subterfuge as part of the attacks. Instead, the forced withdrawal of Tsvangirai is best described as a DoS attack. Now, Tsvangirai and his immediate deputies himself were also personally subject to harassment and violence, and were forbidden access to state-controlled news media. This could be classified as an $O(1)$ DoS attack. There were also allegations of voters’ anonymity being systematically violated at the polls: an $O(N)$ confidentiality attack. Unsurprisingly, the legitimacy of Mugabe’s victory is easy to call into question.

A similar instructive example is the election of “landslide” Lyndon Johnson for the U.S. Senate in 1948 (widely discussed in the literature; see, e.g., [BH94]). Texas, at the time, was largely controlled by the Democratic Party, so the Democratic Primary election was to be decisive for who would win the Senate seat. In that election, which went to a runoff between Johnson and former Texas Governor Coke Stevenson, Johnson defeated Stevenson by an “87-vote landslide.” Much attention has focused on ballot stuffing in Jim Wells County’s infamous “Box 13,” but ballot box stuffing, among other fraudulent behavior, was apparently the norm across the state. Counties were allowed to report “revisions” to their tallies in the week following the election, allowing local party bosses to continuously adjust their vote totals to assist their preferred candidate. (Baum and Hailey note that, had the original, unrevised tallies been used, Johnson would have still defeated Stevenson by 506 votes [BH94, p. 610]. They conclude that Johnson won largely on the strength of his turnout efforts, or rather the lackluster turnout efforts of his opponent.)

Classifying the attacks in this race is tricky. Individual voters’ poll taxes were often paid by the political campaigns, but their votes were still ostensibly cast privately. Still, a well-funded operation to pay for more voters would be classifiable as an $O(N)$ integrity attack, since each voter was individually brought in. (This assumes a candidate’s operation can accurately select its supporters, to pay their poll taxes, without paying for its competition.) Ballot stuffing and poll book manipulation is clearly $O(P)$, since there are P ballot boxes that would need to be compromised, but is it an integrity attack or a DoS attack? The “Box 13” story is all about the discovery of clearly fraudulent signatures in the poll books, and similar fraud likely occurred elsewhere in the state without undergoing the same level of scrutiny. Based on our scheme, the clear presence of the “Box 13” fraud, combined with the inability of anybody to correct the problem, would make it a DoS attack, while other ballot stuffing events that went undiscovered would be integrity attacks. (In other words, ballot stuffing can be a successful attack, whether or not it is discovered, so long as the election results aren’t thrown out and the election redone from scratch.)

3. Applying Complexity Analysis to Voting Systems

a. Absentee paper voting (vote-by-mail)

All of Oregon's voters cast their votes by mail, with large number of voters in others states increasingly preferring this voting mechanism for its comfort and convenience. Of course, by voting in an unsupervised location (i.e., in the voter's home), there is no enforcement of anonymity for voters. This means that vote selling, by voters, is an $O(N)$ attack, because each voter must individually be bribed or coerced. The postal mail is likewise far from a safe and secure delivery channel. With P postal workers (simplifying assumption: the number of postal mail routes is proportionate to the number of voting precincts), one could imagine a variety of $O(P)$ confidentiality attacks (example: steaming open envelopes to read their contents, with each of the P postal workers reading N/P ballots). Once postal votes reach the election headquarters, confidentiality attacks potentially drop to as low as an $O(1)$ cost, requiring at most a handful of corrupt election workers.

Integrity attacks and DoS attacks will follow the same pattern. Individual voters could individually modify ballots from their neighbor, spouse, tenants, and so forth – $O(N)$ effort – requiring nothing more sophisticated than steam to open an envelope, or fire to destroy one. Similarly, P postal workers could open envelopes and either spoil undesirable ballots, substitute alternative ballots, or simply fail to deliver those ballots – an $O(P)$ attack. Election headquarters employees can potentially mount $O(1)$ integrity and DoS attacks, since they're already individually handling each ballot. The postal system has large central sorting facilities. Employees in these locations could potentially mount $O(1)$ DoS attacks (e.g., destroying ballots from particular zip codes), but their ability to mount integrity attacks would vary. In a large city, the speed with which mail is processed would preclude centrally attacking individual envelopes. In smaller towns and counties, $O(1)$ postal attacks against integrity might also be possible.

The existence of these $O(1)$ attacks clearly demands strong procedural controls to prevent them, such as "separation of privilege" procedures, where one worker examines the outer envelope for voter registration information, passing along an inner envelope, holding the ballot, to a separate worker who does not know the identity of the original voter. Of course, both workers could collude and defeat the system. Additional measures, such as having many more employees dividing the labor, installing round-the-clock video surveillance of areas where ballots are handled, or having ballots with hidden marks passing through the system as a check against tampering would certainly be appropriate. Even if the resulting system is still vulnerable to $O(1)$ attacks, additional constant factors of difficulty are far better than nothing. Similarly, the risks from the postal channel would seem to require sophisticated tamper-evident envelopes. If the effort necessary to steam open an envelope without detection was sufficiently high, it might become infeasible for regular postal workers to open and modify ballots without taking too much time, and thus being more easily detected.

b. Paperless, electronic voting systems

Direct recording electronic (DRE) voting system vendors have made a number of claims about their security. Following the California "Top to Bottom" Reports (see, e.g., [CA-Hart07] and [CA-Diebold07]) and the Ohio EVEREST Reports [Everest07], we now know many of these claims to be false. Based on that work, for every major commercial DRE voting system used in the U.S., we can summarize these systems' security as follows.

Anonymity attacks require $O(P)$ effort. P poll workers would need to each record the order in which voters appeared before each machine – a minimal additional effort beyond the work they already perform as part of the voter registration process. (In some states, voter sign-in logs are already recorded in order, as a matter of course, dropping the complexity of the attack to $O(1)$, since no poll workers must be

corrupted!) After that, due to poor engineering of the DRE systems, including easy poll worker access to the ballot records, the order in which votes were cast can be recovered from the electronic records kept by these systems. A variety of engineering tactics can and should be taken to improve this situation, but none are yet deployed. In the best possible case, however, we could still imagine $O(P)$ attacks that do not rely on poor engineering of the voting machines. As an example, small video cameras could be surreptitiously installed in the ceiling above a machine to record its screen as it's used. Consequently, improving anonymity attacks beyond $O(P)$ will be a difficult problem to tackle.

Due to poor engineering practices of the DRE vendors, integrity and DoS attacks require $O(1)$ effort for normal voters – the worst possible case. The California and Ohio analyses discovered “viral” attacks against these voting systems. A single attacker, tampering with a single voting system, can inject a custom-engineered virus⁶. This virus can then take advantage of the poor software engineering of other parts of the voting system and spread from one voting machine to another, based entirely on the usual procedures carried about by election officials. In the case of Diebold's DRE systems, the virus can spread through the smartcards used to authorize the machines to accept votes. In the case of Hart InterCivic's DRE system, one DRE can infect a back-end system used in the election administration's warehouse for inventory control, among other purposes. This system can then infect any subsequent DRE connected to it. An infection introduced in one election can then be present in *every* voting machine used by the county in its subsequent elections. This infection could then have a variety of negative behaviors, ranging from flipping votes through simply deleting them en masse. If the vendors were to repair all the bugs found in these studies, and no other bugs remained to be discovered, these attacks would become $O(P)$ – poll workers would need to tamper with each physical voting system to cause it to misbehave.

A number of procedural measures are commonly suggested for dealing with some of these vulnerabilities. “Logic and accuracy testing” typically involves casting a handful of votes for each candidate and then tabulating the votes. It's trivial for a (tampered) machine to determine when it's being tested like this and to behave correctly. Likewise, L&A testing, conducted by election officials prior to an election, cannot detect tampering conducted by poll workers or normal voters once machines are in the field. L&A testing, if not conducted carefully, can even result in test votes that appear in the final election tally [W06]. Similarly, “hash code testing,” a process for verifying that the software inside a voting machine matches the software that is *supposed* to be there, is ineffectual against sophisticated attackers, who can arrange for a corrupt machine to respond as a proper machine would to any of the tests. (As an example, Hart InterCivic's system has a very detailed mechanism intended to detect code tampering. The California review found it to be easily subverted [CA-Hart07].)

Parallel testing takes a set of voting machines, randomly pulled from general circulation, and instead tests them, on election day, in a controlled and videotaped environment. Any misbehavior of the voting machines would then be caught. Parallel testing procedures can discover the presence of many (but not necessarily all) integrity and DoS attacks, but will have no effect on anonymity attacks, or attacks that are installed by poll workers or attackers, on the site of specific precincts when the election begins. Likewise, while parallel testing might discover such attacks, it offers no way to prevent them or recover from them. At its best, parallel testing converts integrity attacks, previously undetected, into DoS attacks, without itself offering any way to recover the original intent of the voters.

⁶ A side question is whether a normal voter might have sufficient access to open the case of a voting machine or otherwise conduct an attack. In practice, voting machines are often left unguarded in relatively public places. Furthermore, some vendors support a “curb-side voting” procedure, where disabled voters who cannot physically enter the polling place can be accommodated by bringing voting machines out to them instead. Curb-side voting gives attackers the necessary time and privacy to conduct an attack. Many of the attacks could even be conducted by a voter while in the voting booth.

c. Electronic voting machines with VVPAT printers

In response to objections from computer scientists and activists, and to requirements by many states, DRE vendors now offer a variety of “voter-verified paper audit trail” (VVPAT) attachments for their machines. The major U.S. vendors all use a reel-to-reel thermal printer, behind glass. Voters may see the record of their vote but may not touch it. Because the paper is never cut, these devices record votes in the order cast. For the purposes of this analysis, we will assume that these printers are actually working properly (versus anecdotal evidence that they suffer mechanical failures) and that voters will actually read them (even though controlled human subject experiments suggest that two-thirds would fail to see any discrepancies [E07]).

Confidentiality attacks are still $O(P)$, based purely on the electronic results being easily deanonymized (as described above for paperless DRE systems, and would likewise drop to $O(1)$ if poll workers or electronic poll books recorded the order in which voters appeared as their standard operating procedure). Additionally, the recording of votes in the order cast on the VVPAT paper tapes similarly allows for $O(P)$ attacks. If the paper were cut after each voter and piled up in a random order in a bag of some kind, then the VVPAT attachment would no longer be the most obvious way to conduct a confidentiality attack.

Tampering and DoS attacks improve from $O(1)$, with paperless DREs, to $O(P)$, with VVPAT printers, by virtue of increasing the burden on the attacker. Now the attacker must also perform $O(P)$ effort to consistently tamper with the VVPAT printouts. This implies that VVPAT, even poorly implemented, yields a significant complexity improvement relative to existing DREs’ security performance. However, achieving this improvement requires that election officials reconcile the VVPAT results against the electronic results. Given Everett’s findings [E07], that many voters will fail to notice errors on VVPAT results, election officials must be particularly attuned to even a very low rate of complaints, creating a secondary vulnerability, namely that voters could falsely claim to have observed discrepancies in order to cause perfectly good equipment to be removed from service – an $O(P)$ DoS attack.

In response to the California and Ohio studies, Halderman et al. [HRSW08] considered mitigations that might improve the security of electronic voting systems. They describe detailed system-specific procedures that can mitigate many (but not all) $O(1)$ attacks against these systems. Their techniques expand on the concepts described above, and rely on having the VVPAT paper trails. In the end, they cannot prevent integrity attacks, but they can limit the degree to which they spread and can increase the confidence with which they may be discovered. On paperless DRE systems, variations on their techniques might improve some of the $O(1)$ integrity attacks to $O(P)$ attacks, as well. Many of the mitigations proposed by Halderman et al. require custom hardware “gadgets” for purposes such as clearing memory cards. No such gadgets exist, and it’s unclear that such gadgets could legally be used without being certified by the appropriate federal and state authorities.

d. Precinct-based optical scan systems

While some activists feel that the only safe way to vote is to count ballots by hand, optical scanning of ballots adds significant efficiencies to the process, particularly in modern American elections where voters may be asked to respond to 30 or 40 separate issues. There are two main styles of ballot scanners: precinct-based scanners and central scanners. The former is attached to the top of a ballot box while the latter is installed at the central election headquarters. Precinct scanners can issue warnings or errors in cases of under or over-voted ballots, potentially improving accuracy. From a security perspective, the hand-marked paper ballots constitute a voter-verifiable paper trail, providing an independent path to determine the voters’ intent, in cases where the electronic records may be called into question. For this analysis, we will assume that the paper records are, indeed, reconciled against the electronic tallies. Without mandatory reconciliation or auditing, the security analysis would be much more pessimistic.

Precinct scanners offer comparable anonymity to any other technology. A malicious scanner can record votes in the order in which they were cast – an $O(P)$ attack, since P different scanners must be corrupted. Vulnerabilities have been found in commercial optical scanners, as part of the California and Ohio studies, showing that even unmodified scanners are improperly engineered, allowing the order in which votes were cast to be recovered. If these known-flawed optical scanners were used in precincts that recorded the order in which voters sign in, as a matter of regular procedure, that combination would lower the complexity of an anonymity attack to $O(1)$.

The integrity guarantees of precinct scanners are quite good, assuming proper auditing of the paper ballots occurs after the election. If the electronic records were compromised, the paper records could be used to reconstruct them (either via a different scanner or by hand). If the paper records were compromised, the electronic records could be used as a backup. Successful integrity attacks are thus an $O(P)$ process. An attacker would need to modify both the electronic records and the paper ballots in a consistent fashion and would need to attack each precinct in turn.

Precinct scanners have reasonably good resistance to DoS attacks. Even in the worst case, if they are completely broken, poll workers can remove them from the ballot boxes, falling back on the central scanners. Still, a variety of $O(P)$ attacks exist, such as exist against any precinct (e.g., cutting the power), as well as attacks such as intercepting or spoiling the paper ballots while in transit.

Centrally scanned ballots, whether used as a fallback for failed precinct ballot scanners, or just used for mail-in votes, are clearly subject to $O(1)$ insider attacks. Auditing procedures would clearly be necessary to double-check that the electronic tallies correspond accurately to the paper ballots.

e. Internet voting (via home computers)

We can conduct all manner of high-dollar commerce on our home computers, so why not vote? Indeed, Estonia has done this, as have a handful of other elections. Of course, there are significant differences between the world of e-commerce and the world of e-voting. Commercial fraud happens all the time (identity theft, etc.), leading to a huge industry that deals with this fraud. Nonetheless, e-commerce succeeds because the total volume of legitimate commerce dwarfs the volume and expense of fraud.

Furthermore, commercial transactions are not anonymous. Customers, merchants, and their banks carefully record these transactions. If and when discrepancies arise, customers can challenge merchants, and everybody can settle on the proper outcome. In e-voting, however, votes must be anonymous. Without anonymity, bribery and coercion become valid $O(N)$ threats against the election.

If we want to use home computers to vote over the internet, that means dealing with Microsoft Windows and Internet Explorer, for which new security vulnerabilities are discovered on a depressingly regular basis (not to say that any other consumer computing platforms are necessarily any better in this regard). One of the key risks this creates is that attackers might focus their attacks specifically on people who are likely to vote using the internet. For example, if an attacker wanted to tamper with the Estonian online election, the attacker might first break into a popular Estonian web site (a new site, a bank site, etc.) and install some kind of malware that could then infect visitors to that site. Likewise, many voters' computers can be directly attacked over their network connections. Attackers can use a variety of "geo-location" techniques to focus their efforts exclusive on computers within the country being attacked. Any machine that has been compromised would then presumably act normally except when visiting the election authority's web site. At that point, it's trivial for the computer to show one thing to the voter while behaving in an arbitrarily different fashion on the network. There would be no way for the voters to know whether their votes had been flipped. In fact, a custom-engineered attack could erase itself subsequent to the election, leaving no trace of its presence.

The Estonian election had a number of particular security measures that they put into place, notably requiring the use of the national ID card, a “smartcard,” in order to authenticate to the voting web site. Estonia also allows voters to cast multiple votes, with the last one counting. They additionally allowed voters to subsequently vote in person, canceling out any prior electronic vote. With these mechanisms in mind, as they might well be copied with other internet-based voting schemes, we can analyze the security of many other proposed internet voting systems.

Confidentiality attacks that rely on in-person observation of the voting process (i.e., somebody watching over your shoulder as you vote) are ostensibly defeated by Estonia’s multiple voting scheme, but in fact the attackers could well “borrow” the voters’ ID cards, preventing subsequent online votes, and watching to make sure the voters do not attempt to go vote in person. These $O(N)$ attacks would be very difficult to defeat. More worrisome is the fact that the back-end computers carefully track the binding between the voter and the more recent cast vote. This could allow malicious code, in the back-end, to observe these bindings – an $O(1)$ attack. Likewise, malicious client-side code, installed in a viral fashion, could observe voters’ actual behavior – another $O(1)$ attack.

Integrity attacks on the client side would rely on breaking into the voters’ home computers. It’s difficult to estimate what percentage of these computers might be vulnerable to modern attack tools, but the cost of breaking in is $O(1)$ – it’s no more effort to break into one system than a thousand systems. There may also be $O(1)$ integrity attacks against the server side, depending on the particular details of how the server has been implemented, how the cryptographic exchange works with the smartcard, and so forth. Given the difficulty inherent in building a secure web-facing service, external attacks may always be feasible. And, of course, insider attacks would also be $O(1)$. If and when such attacks were detected, one could imagine heavy-weight solutions, such as discarding online votes and requiring voters to appear in person. Clearly, an attacker would have to avoid the temptation to exert too much influence on the election outcome in order to fly below the radar.

DoS attacks are trivially available with $O(1)$ effort. The elections web server has finite resources that can be easily exhausted by standard internet attack tools that leverage vast numbers of compromised “zombie” computers to make requests as fast as possible of the server. Such “distributed denial of service” (DDoS) attacks require extensive server-side resources to counter. Google and Yahoo! engineers will privately admit that, despite the massive scale of their operations, they still worry about DDoS attacks. If the biggest web services in the world have to worry about DDoS attacks, an elections web server is clearly in trouble. DDoS attacks can also be mounted against an entire country. Estonia, in fact, was attacked by (alleged) Russian hackers in 2007, crippling large parts of the country’s internet infrastructure [T07]. To date, nobody has been charged with a crime. (Wikipedia [Wiki07] has a detailed summary of the incident. Nazario [N07] has more technical details.) If this were to occur during an election, the internet voting infrastructure would almost certainly fail. One could imagine a variety of draconian solutions, including disconnecting the internet connection for the entire country. This would dampen attacks from outside the country, but would do nothing if attackers had compromised systems, within the country, set to mount their attacks autonomously.

In short, internet voting in Estonia, via the failed SERVE project, or really any other way using the internet to connect personal computers’ web browsers to web servers for purposes of casting votes, is subject to a variety of $O(1)$ attacks. While Estonians can at least fall back to other voting technologies, the push to adopt internet voting, even in limited trials, is inappropriately ignoring these glaring risks.

f. Future “end to end” cryptographic voting systems

Cryptographers for the last two decades have been promoting a family of mathematical techniques that lead to a property called “end to end verification” of elections. While specific techniques vary, there are several features these systems all have in common. Voters still cast votes in a polling place, requiring and

ensuring their anonymity. Whether using paper or an electronic voting machine, voters may optionally leave the polls with some kind of “receipt” that will allow them to identify a record corresponding to their vote on a public “bulletin board.” The bulletin board does not contain human-readable votes, but rather an encrypted form of the vote. The tallying process requires the participation of one or more trusted parties, and ultimately yields both the election totals as well as a “proof” that can be verified by anybody. Anybody can read the bulletin board and verify the computation that was used to derive the vote totals, even though they could not have performed the whole computation themselves.

Other properties vary. One clever idea from Benaloh [B07] allows any voter to “challenge” any voting system to prove that it is properly encrypting the voter’s ballot. The idea is that a normal DRE system would go through its normal dialog with the voter until the voter agrees with the contents of the summary screen. At this point, the machine would generate the encrypted vote, sign it, and physically print it. The voter sees the printing taking place. At this point, the voter gets one new question: “would you like to cast this ballot or challenge it?” If the voter says “challenge”, then the machine prints the encryption keys, allowing anybody to decrypt the vote and make sure that it’s what the voter intended; this vote will *not* be counted. If the voter says “cast”, then the machine throws away the key and casts the vote. These challenges allow roving election auditors to enter polling places while the election is ongoing, set up video cameras as in parallel testing, and verify that voting machines are behaving properly. Unlike parallel testing, however, the machines being tested are the very live machines being used by real voters. Variations on Benaloh’s technique have already been implemented in two separate research voting systems (VoteBox [SKW08] and Helios⁷ [A08]).

In order to be concrete, we will analyze the security of VoteBox [SWK08], which feels to most voters like a DRE voting system, and Scantegrity II [C+08], which feels to most voters like an optical scan voting system.

VoteBox, like any DRE, could be maliciously modified to record votes in the order cast. Scantegrity, like any precinct optical scanner, has the same property. Both are thus subject to $O(P)$ anonymity attacks.

VoteBox and Scantegrity have very strong mechanisms to detect integrity attacks. VoteBox’s challenge mechanism can detect a misbehaving machine, as the election is ongoing. Scantegrity looks and feels much like a regular paper ballot, so it inherits the integrity benefits of paper ballots. Scantegrity also has mechanisms that allow third-parties to verify the tally and to challenge individual ballots, all without revealing how any specific voter cast a ballot. Impressively, neither VoteBox nor Scantegrity is vulnerable *at all* to integrity attacks, assuming these systems are used and audited properly. If only a handful of votes are flipped, then they may go undetected, but the odds of successful integrity attacks decline quickly as more votes are flipped. As such, most attempts to compromise ballot integrity will quickly degenerate to DoS attacks, and even those may be recoverable.

DoS for Scantegrity is identical to precinct-based optical scan – all the same $O(P)$ attacks apply. DoS for VoteBox is likewise still $O(P)$, but VoteBox contains a number of data replication features that at least make it a constant factor harder. In particular, for precincts with many VoteBox systems, they would be networked together, with each machine holding copies of the votes from every other machine. An attacker would need to successfully attack *every* VoteBox machine in order to destroy the records of an election. Scantegrity, by virtue of the paper ballots, is more robust against electronic tampering. However, there are still only $O(P)$ ballot boxes, themselves subject to physical tampering. In the end, while the specific attacks would be quite different, VoteBox and Scantegrity both have $O(P)$ DoS vulnerabilities, allowing us to reasonably conclude that neither has a decisive advantage in this respect.

⁷ Interestingly, Helios is a voting system implemented using standard web servers and clients. However, it is explicitly intended only for “low coercion” elections. This makes Helios appropriate for electing the leader of an open-source software project or perhaps even for conducting Faculty Senate elections in a university setting. Helios was never intended for electing politicians to important offices.

The cryptography in VoteBox and many other end-to-end systems generally rely on a notion of “election trustees”, who hold “shares” of cryptographic key material and must collaborate together to tally and decrypt the votes. Several $O(1)$ confidentiality attacks exist if all the trustees collude together⁸. At that point, they could decrypt any individual ballot (violating confidentiality), learning how a specific voter cast his or her vote. They would be unable to delete ballots from the bulletin board, as voters would be able to detect the absence of their ballots. They could perhaps add additional ballots, but these would need to have corresponding entries in poll books. (Ultimately, preventing ballot stuffing attacks, in any voting technology, relies on having good records of whether a voter has or has not cast a vote, along with audits of these records for consistency with the number of cast ballots.) Scantegrity places all of this trust in a central election administrator. If corrupt, a Scantegrity election administrator could violate the confidentiality of every ballot – an $O(1)$ attack. (Generalizing Scantegrity to have a set of trustees rather than a single trusted election administrator would be a straightforward extension to the system.)

VoteBox has one additional $O(1)$ DoS vulnerability: the destruction or loss of trustees’ cryptographic key materials. While key shares could be copied or backed up, this process would require careful design. Scantegrity is not vulnerable to a comparable attack because its paper ballots are always human-readable and can be tallied using traditional scanners or by hand.

An important concern with VoteBox, Scantegrity, or any other e2e crypto voting system is what must happen if a fault is found. What should the proper procedure be if a voter can prove, beyond any doubt, that his or her ballot was not included in the final tally? How many voters must provide such proof before the election results are called into question? Ultimately, e2e crypto voting mechanisms provide a great way of proving that an election was done properly. Resolving errors, omissions, or fraud, once they are shown to exist, is likely to be a more labor-intensive process.

An intriguing issue for cryptography-intensive voting systems like VoteBox and Scantegrity is the issue of the strength of the cryptosystem itself. Modern cryptography rests on the computational difficulty of solving a number of “hard problems” such as factoring very large numbers. What happens if and when a future cryptographer invents an effective attack against the cryptography? A breakthrough like this would have ramifications well beyond the world of elections. Certainly, there would be a period wherein older broken cryptographic algorithms would need to be replaced with newer, replacement algorithms. The main concern, for elections, would be that votes cast prior to the breakthrough could potentially be decrypted. (A similar threat occurs if the “trustees” collude, as described above.) A cryptographic failure, assuming it was known and corrected prior to the election, would not be usable to compromise election integrity or availability.

4. Summary and Conclusions

The table below summarizes the complexity analysis of attacking a variety of different voting technologies.

⁸ Cryptographic key sharing mechanisms can be configured to require *every* share-holder to participate to perform a computation, or “threshold” schemes can be used that require some k out of n participants. If at least one trustee is honest (or, with threshold schemes, if a quorum of colluding trustees cannot be formed), then the system is secure. A malicious trustee could potentially conduct an $O(1)$ DoS attacks against the election by simply refusing to participate, and thus preventing the decryption of the election totals, but this would be impossible to do anonymously. Regular criminal penalties for this sort of behavior would hopefully provide the necessary deterrence.

	Confidentiality	Integrity	Denial of Service
Vote-by-mail	$O(N)$ – individual voters selling their votes $O(P)$ – postal workers $O(1)$ – election insiders	$O(N)$ – individual voters selling their votes $O(P)$, possibly $O(1)$ – postal workers $O(1)$ – election insiders	$O(N)$ – individual voters selling their votes $O(P)$, possibly $O(1)$ – postal workers $O(1)$ – election insiders
Paperless electronic	$O(P)$	$O(1)$ *	$O(1)$ *
Electronic + VVPAT	$O(P)$	$O(P)$	$O(P)$
Precinct optical scan	$O(P)$	$O(P)$	$O(P)$
Internet voting	$O(1)$	$O(1)$	$O(1)$
End-to-end crypto voting (Scantegrity)	$O(P)$ – precinct tampering $O(1)$ – corrupt election administrator	<i>impossible</i>	$O(P)$
End-to-end crypto voting (VoteBox)	$O(P)$ – precinct tampering $O(1)$ – trustee collusion	<i>impossible</i>	$O(P)$ $O(1)$ – trustee key loss

(* These $O(1)$ attacks may be partially mitigated using complex procedures described by Halderman et al. [HRSW08], which could strengthen the systems from $O(1)$ to $O(P)$.)

Based on a straightforward complexity analysis of different voting technologies, we can reach clear conclusions about their relative strengths and weaknesses. While no voting system is perfect, understanding the relative effort required to mount an attack is critical to allocating resources toward defeating these attacks. For example, if an election official is conducting a vote-by-mail election, then clearly they need to spend significant effort in their internal security procedures to detect and defeat the constant-cost threat of insider attacks. Likewise, the constant-effort attacks against internet voting, and the lack of adequate remedies, provide more than sufficient justification to use other modalities for conducting elections. (We do note that the internet is useful for connecting precincts back to election headquarters, so long as it's not relied upon for election correctness [SKW08].) Another consistent theme is that election registration procedures should never record the order in which voters appear, as this reduces the complexity of anonymity attacks across many different technologies. Of voting technologies presently on the market and certified for use in most states, precinct-based optical scanning systems appear to have the best resistance to security attacks, based on our complexity analysis. Future cryptographic end-to-end systems should be able to improve upon this, although it's an open question whether normal voters would find these systems to be usable.

Acknowledgements

The author wishes to thank Aaron Burstein, Ed Felten, Joe Hall, David Jefferson, Doug Jones, Eric Rescorla, David Wagner, and Nick Weaver for their comments on drafts of this article. This work was supported in part by NSF grants CNS-0524211 and CNS-0509297.

Citations

[A08] Ben Adida. *Helios: Web-based open-audit voting*. *Proceedings of the 17th USENIX Security Symposium (USENIX Security '08)* (San Jose, California), July 2008. http://www.usenix.org/events/sec08/tech/full_papers/adida/adida.pdf

[SKW08] Daniel R. Sandler, Kyle Derr, and Dan S. Wallach. VoteBox: a tamper-evident, verifiable electronic voting system. *Proceedings of the 17th USENIX Security Symposium (USENIX Security '08)* (San Jose, California), July 2008.
<http://www.cs.rice.edu/~dsandler/pub/sandler08votebox.pdf>

[HRSW08] J. Alex Halderman, Eric Rescorla, Hovav Shacham, and David Wagner. You Go to Elections with the Voting System You Have: Mitigations for Deployed Voting Systems. *Proceedings of the 3rd USENIX/ACCURATE Electronic Voting Technology Workshop (EVT '08)* (San Jose, California), July 2008.
http://www.usenix.org/events/evt08/tech/full_papers/halderman/halderman.pdf

[C+08] David Chaum, Richard Carback Jeremy Clark, Aleksander Essex, Stefan Popoveniuc, Ronald L. Rivest, Peter Y.A. Ryan, Emily Shen, Alan T. Sherman. Scantegrity II: End-to-end verifiability for optical scan election systems using invisible ink confirmation codes. *Proceedings of the 3rd USENIX/ACCURATE Electronic Voting Technology Workshop (EVT '08)* (San Jose, California), July 2008.
http://www.usenix.org/events/evt08/tech/full_papers/chaum/chaum.pdf

[BMHW08] Aaron Burstein, Joseph Lorenzo Hall, Deirdre Mulligan and David Wagner. Comment on the U.S. Election Assistance Commission's Request for Information Regarding Voting Systems Risk Assessment Support; Submitted to the Election Assistance Commission by the Samuelson Law, Technology and Public Policy Clinic (2008).
<http://josephhall.org/papers/ucb-eacrificomment-20080428.pdf>

[DB08] Celia W. Dugger and Barry Bearak. Mugabe is Sworn in to Sixth Term After Victory in One-Candidate Runoff, *New York Times*, June 30, 2008.
<http://www.nytimes.com/2008/06/30/world/africa/30zimbabwe.html>

[EAC08] Election Assistance Commission, "Risk Assessment Summary", April 2008.
http://www.eac.gov/voting%20systems/voluntary-voting-guidelines/docs/risk-assessment-summary/attachment_download/file

[B08] Barry Bearak. Mugabe Foes Win Majority in Zimbabwe, *New York Times*, April 3, 2008.
<http://www.nytimes.com/2008/04/03/world/africa/03zimbabwe.html>

[AH08] R. Michael Alvarez and Thad E. Hall, *Electronic Elections: The Perils and Promises of Digital Democracy*, Princeton University Press (Princeton, New Jersey), 2008.

[Everest07] EVEREST Testing Reports, December 2007.
<http://www.sos.state.oh.us/elections/voterinformation/equipment/VotingSystemReviewFindings/EVERESTtestingReports.aspx>

[B07] Josh Benaloh. Ballot casting assurance via voter-initiated poll station auditing. *Proceedings of the 2nd USENIX/ACCURATE Electronic Voting Technology Workshop (EVT '07)* (Boston, MA), August 2007.
http://www.usenix.org/events/evt07/tech/full_papers/benaloh/benaloh.pdf

[CA-Hart07] Srinivas Inguva, Eric Rescorla, Hovav Shacham, and Dan S. Wallach. Source Code Review of the Hart InterCivic Voting System. *Report commissioned as part of the California Secretary of State's Top-To-Bottom Review of California voting systems* (Sacramento, California), July 2007.
http://www.sos.ca.gov/elections/voting_systems/ttbr/Hart-source-public.pdf

[CA-Diebold07] Joseph A. Calandrino, Ariel J. Feldman, J. Alex Halderman, David Wagner, Harlan Yu, William P. Zeller. Source Code Review of the Diebold Voting System. *Report commissioned as part of the*

California Secretary of State's Top-To-Bottom Review of California voting systems (Sacramento, California), July 2007.

http://www.sos.ca.gov/elections/voting_systems/ttbr/diebold-source-public-jul29.pdf

[Wiki07] 2007 cyberattacks on Estonia. *Wikipedia, The Free Encyclopedia*. Retrieved 15:46, August 15, 2008.

http://en.wikipedia.org/w/index.php?title=2007_cyberattacks_on_Estonia&oldid=231225092

[N07] Jose Nazario. Estonian DDoS attacks – a summary to date. *Arbor Networks Security Blog*, May 17, 2007.

<http://asert.arbornetworks.com/2007/05/estonian-ddos-attacks-a-summary-to-date/>

[T07] Ian Traynor. Russia accused of unleashing cyberwar to disable Estonia. *The Guardian*, May 17, 2007.

<http://www.guardian.co.uk/world/2007/may/17/topstories3.russia>

[E07] Everett, S. P. *The Usability of Electronic Voting Machines and How Votes Can Be Changed Without Detection*. Doctoral dissertation, Rice University (Houston, Texas), May 2007.

<http://chil.rice.edu/research/pdf/EverettDissertation.pdf>

[N06] Lawrence Norden et al., *The Machinery of Democracy: Protecting Elections in an Electronic World*. Brennan Center Task Force on Voting System Security, Brennan Center for Justice, NYU Law School, June 2006.

http://www.brennancenter.org/dynamic/subpages/download_file_39288.pdf

[W06] Dan S. Wallach, Security and Reliability of Webb County's ES&S Voting System and the March '06 Primary Election. Expert report in *Flores v. Lopez* (Laredo, Texas), May 2006.

<http://accurate-voting.org/wp-content/uploads/2006/09/webb-report2.pdf>

[J05] Douglas W. Jones. Threats to Voting Systems, *NIST Workshop on Developing an Analysis of Threats to Voting Systems* (Gaithersburg, Maryland), October 2005.

<http://www.cs.uiowa.edu/~jones/voting/nist2005.shtml>

[JRSW04] David Jefferson and Aviel D. Rubin and Barbara Simons and David A. Wagner. A Security Analysis of the Secure Electronic Registration and Voting Experiment (SERVE), January 2004.

<http://www.servesecurityreport.org/>

[CLRS01] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*, 2nd ed. MIT Press (Cambridge, Massachusetts), 2001.

[BH94] Dale Baum and James L. Hailey, Lyndon Johnson's Victory in the 1948 Texas Senate Race: A Reappraisal. *Political Science Quarterly*, Vol. 109, No. 4 (Autumn 1994), pp. 595-613.

<http://www.jstor.org/stable/2151840>