

Machine-Assisted Election Auditing

Joseph A. Calandrino^{*}, J. Alex Halderman^{*}, and Edward W. Felten^{*,†}

^{*}Center for Information Technology Policy and Dept. of Computer Science, Princeton University

[†]Woodrow Wilson School of Public and International Affairs, Princeton University

{jcalandr, jhalderm, felten}@cs.princeton.edu

Abstract

Election audit procedures usually rely on precinct-based recounts, in which workers manually review all paper ballots from selected polling places, but these recounts can be expensive due to the labor required. This paper proposes an alternative audit strategy that allows machines to perform most of the work. Precincts are recounted using recounting machines, and their output is manually audited using efficient ballot sampling techniques. This strategy can achieve equal or greater confidence than precinct-based auditing at a significantly lower cost while protecting voter privacy better than previous ballot-based auditing methods. We show how to determine which ballots to audit against the recounting machines' records and compare this new approach to precinct-based audits in the context of Virginia's November 2006 election. Far fewer ballots need to be audited by hand using our approach. We also explore extensions to these techniques, such as varying individual ballots' audit probabilities based on the votes they contain, that promise further efficiency gains.

1 Introduction

Security analyses of computerized voting systems, including DREs and optical scan machines, have exposed numerous vulnerabilities that could compromise the integrity of elections performed using these devices (see [7, 5] and references therein). One proposed defense against such attacks is to produce voter-verified paper records and audit them to ensure that they support the totals claimed by the machines.

The most common auditing method is the precinct-based audit [3, 9, 10, 12, 13], in which workers count all paper ballots from selected precincts and compare the results to the reported precinct tallies.¹ Unfortunately, performing precinct-based audits can require considerable

time, labor, and expense. These costs are multiplied by the complexity of the ballots in many elections, which may include dozens of contests. In a trial recount of a DRE paper trail performed in Cobb County, Georgia, workers took an average of 5 minutes per ballot to audit 976 votes at a total cost of nearly \$3,000 [4]. Unless efficiency can be improved, performing a similar recount of 3% of precincts in New Jersey could cost more than \$200,000. Slow, expensive manual recounts limit the level of confidence that can be achieved within a fixed election budget, and they may delay the detection of errors until well after election results have been announced and losing candidates have conceded.

In this paper we propose an alternative audit strategy that substantially reduces these costs by using specialized machines to automate most of the work of recounting paper ballots followed by a manual audit of the machine results. The problem with machines, of course, is that the ones used for the recount are not necessarily more trustworthy than the ones used in the initial count. They may be useful for catching inadvertent errors (especially if they use a different technology and independently developed software), but a determined attacker could still target both sets of machines. What we desire is software independence—an assurance that any tampering with the machines will not cause undetected changes to the election outcome [11]. To achieve this, we pair recount machines with efficient statistical auditing techniques that allow humans to confirm that the election outcome is correct.

Statistical “ballot-based” audits are an alternative to manually recounting every ballot from selected precincts. Workers sample from all the paper ballots in all precincts and use the sample to assess the accuracy of the original count. Ballot-based audits tend to be more efficient than traditional precinct-based audits [8], since fewer ballots need to be recounted to achieve the same level of

¹H.R. 811, now under consideration in Congress, would mandate a complete manual recount of 3%, 5%, or 10% of precincts, depending

on the margin of victory. However, the bill permits alternative recount methods so long as they provide an equivalent level of confidence. [1]

confidence in the result. For example, in a state-wide race in New Jersey, fewer than one ballot per precinct (4,599 ballots total) would need to be sampled to achieve 99% confidence that the outcome had not been shifted by more than 0.2%. By contrast, over 150,000 ballots (6.9% of precincts) would need to be recounted using standard precinct-based audits (e.g., [13]) to achieve the same confidence.

Neff [8] and Johnson [6] were among the first to propose combining ballot-based audit techniques with electronic voting. Neff assumes that the voting machines link each paper ballot to its electronic counterpart using, for example, a unique identifier printed on the paper ballot and stored with the electronic ballot. When voting is complete, each precinct commits to its set of electronic ballots, then demonstrates that the paper ballots in a given random sample match the corresponding electronic ballots.

The primary weakness of this method is that it establishes the link between electronic and paper ballots at the time that votes are cast. This raises problematic voter privacy issues. For example, if the ballots are linked using sequentially increasing serial numbers, observers could correlate votes with the order in which they were cast, which can reveal the identity of voters. While a cryptographic link might protect privacy, opaque, random-looking identifiers printed on ballots may provide covert channels for leaking voter identities. Even if used securely, they might aid malicious parties who seek to intimidate voters by undermining their confidence in the secrecy of the ballot. Our audit strategy postpones linking paper and electronic records until the recount phase, which allows it to achieve equivalent confidence without jeopardizing privacy or resorting to cryptography.

Johnson alternatively proposes delaying both vote tallying and serial number printing until after all ballots are submitted, allowing voting machines to be simple, memory-less ballot printers [6]. Voters submit their ballots, which, once polls close, are randomized and scanned/tallied. The tallying machine is therefore able to print serial numbers while scanning without privacy risk. Unlike Johnson, we assume that the voting machines maintain an electronic tally, which helps deter traditional attacks against paper-based voting, such as ballot-box stuffing, and, as we will show, provides opportunities for improving the efficiency of the audit.

Our main contributions are:

- We propose a novel audit approach wherein ballots are recounted using recounting machines, and their output is manually audited by humans using ballot-based auditing techniques. (Sections 2 and 3)
- We evaluate the efficiency of our method using data from Virginia’s November 2006 elections, and we

find that it enjoys significant gains compared to the traditional precinct-based approach. (Section 4)

- We suggest several extensions to address practical considerations and to further improve efficiency, including a means of using knowledge of ballot contents to reduce the sample size. (Section 5)

We present additional details in the full version of this paper, available at <http://itpolicy.princeton.edu/voting>.

2 Machine-Assisted Auditing

We propose replacing manual precinct-based audits with machine-assisted audits. Poll workers, rather than recounting ballots manually, feed them through a specialized recount machine that functions like a combined optical scanner and printer. As it scans the contents of each ballot, it prints a unique serial number that is stored along with the ballot contents. At the end of the scanning process, the machine outputs a list of votes on each ballot together with the ballot’s serial number. If the recount tallies differ from the initially reported electronic count, discrepancies clearly exist and a wider investigation should be conducted.² If both tallies match, the workers perform a secondary audit to check the accuracy of the machine’s recount. They first quickly flip through the pile of numbered ballots to ensure that it increases sequentially from one to the reported ballot total without repeats.³ They then take a random sample of the electronic ballot records, retrieve the corresponding paper ballots, and verify that they match.

Since the ballots are serialized and fed out of the machine in order, retrieving a particular ballot for verification requires very little effort. The most significant labor required may be to check for repeats, which given sequential ordering, is a rapid single-pass process.

In practice, separate devices may be used to perform the printing and scanning functions of our proposed recount machine. When voting is complete, a printer device could place serial numbers on the ballots, and then a separate scanner could read the numbers along with the votes. In precincts utilizing optical scan machines, properly designed machines could perform both the initial count and the recount: this option decreases costs but reduces redundancy. If the same machine performs counts, recounts, and printing, officials must have some

²Depending on circumstances, an appropriate response might be to inspect the corresponding machines, other machines of the same model, other ballots in that precinct, etc. A comprehensive set of precise responses to various circumstances is beyond the scope of this paper, but we note the importance of establishing procedures to avoid partisan bickering if discrepancies arise.

³This check helps protect against collusion between voting and recount machines, as described shortly. A more efficient means of ensuring that the number of paper and electronic ballots match may exist.

means of mechanically disabling the printer while polls are open, such as removal of the printer head. Printers also must be physically unable to alter the record of the vote on the ballot. They could be designed so that they cannot reach outside of a predefined empty margin on ballots, or they could utilize a kind of ink that would be immediately apparent when ballots were inspected.

Security The redundancy of combining electronic and paper-based systems increases the security of the overall system. With high probability, the manual audit process detects any discrepancies between the sets of electronic and paper ballots that are substantial enough to impact the election's outcome. Because the process checks the correspondence between the sets of ballots, measures improving the integrity of either set increase the overall integrity of the election result. Since both the electronic and paper ballot sets must remain similar for a discrepancy to avoid detection, combined systems are more likely to detect malfunctions, and they increase the sophistication necessary to commit fraud.

For an error to go undetected, the voting machine must report an incorrect electronic tally, the recount machine must support the incorrect tally, and the manual audit process must not detect a discrepancy between the paper and electronic ballot records.

A malfunctioning or dishonest voting machine may add, subtract, or switch votes to introduce errors in its electronic tallies. If election officials maintain an accurate sign-in list for the precinct, any significant discrepancy in the total number of reported votes will be detected. Therefore, the voting machine is limited to switching votes from one candidate to another.

For a recount machine to support an incorrect electronic tally, either the set of paper ballots must match the incorrect tally, or the machine must fail to detect a discrepancy. The set of paper ballots can only match the tally if either the voting machine printed an incorrect set of paper ballots or another party modified that set. If voters generally verify their paper ballots, the ballot box will likely contain an accurate paper ballot for most voters when polls close. While the voting machine may print additional, incorrect ballots, this would cause the number of paper ballots to exceed the electronic vote total, which mirrors the number of voters, so an accurate recount machine would detect this discrepancy. The simple, sequential nature of machine-assisted auditing also reduces opportunities for adversaries to modify paper ballots during the audit. Assuming that no adversary can modify the set of paper ballots after polls close, only recount machine malfunction, whether accidental or malicious, would allow the discrepancy to go undetected.

A malfunctioning recount machine may report incorrect electronic ballots that agree with any incorrect elec-

tronic tally regardless of the true paper ballots. The machine may even collude with other parties by omitting or printing incorrect serial numbers on paper ballots to hide errors. For example, a voting machine may print additional paper ballots with desirable votes, and a recounting machine may reuse serial numbers on certain undesirable voter-verified paper ballots to effectively replace them with the additional ballots. The manual check of serial numbers detects duplicate or omitted serial numbers and ensures that the number of paper ballots matches the total reported number of electronic ballots.

If no errors are detected before the sampling phase, we know that we have a set of electronic ballots from the recount machine that support the initial electronic tally and an equal-sized set of paper ballots with corresponding serial numbers. We designed the sampling process specifically to detect discrepancies between these sets significant enough to affect the election's outcome. Unless an error or adversary modified both the initial electronic tally and the paper ballots, the manual audit should catch any remaining errors with a high level of confidence.

Privacy Our technique avoids many of the privacy issues inherent in some earlier ballot-based audit methods that involve placing identifiers on ballots during the voting process. In our technique, the ballots do not receive serial numbers until the recount phase, so they are likely to become at least partially reordered before being numbered. Well-designed ballot boxes and cut-and-drop paper trail systems assure that the papers are somewhat shuffled as they are inserted. Since voters widely trust these methods to frustrate correlation with voter check-in times, this provides significant practical privacy benefits. Should alternative ballot shuffling methods offer greater protection, officials may substitute such methods without modifying the audit process. In any case, the recount machine has no more information about the order of votes than would workers performing a manual recount.

Another benefit of this technique is that a voting machine need only maintain tallies rather than electronic copies of individual ballots. Thus, voting machine designers do not need to worry about properly shuffling electronic ballots to protect voter privacy or about maintaining storage for those ballots. However, if the same machines perform counts and recounts, they must have some means of attaching extra memory during the recount for storing the ballot scan results.

3 What to Audit

Due to the popularity of plurality voting systems in the U.S. we exclusively consider those systems, though machine-assisted audits may be useful in many other vot-

ing systems. With plurality voting, voters may choose a number of candidates equal to the number of seats available.⁴ If k seats are available, voters may select up to k candidates, and candidates receiving the top k vote totals are the victors. This definition is an extension of the familiar single-seat contest.

An audit process need only sample enough ballots to confidently detect the minimum amount of fraud that would have affected the election’s outcome. To modify the fewest ballots while changing the outcome, an adversary would swap the positions of the losing candidate with the most votes and the victor with fewest votes. Switching votes directly between these candidates requires the fewest ballot changes, as each switch alters the relative difference by two. To do so, the adversary would take ballots with votes for candidate A but not B and change them to contain votes for B but not A. Therefore, we need only audit enough ballots to discover fraud that alters a number of ballots equal to half the difference (rounded up) in vote totals between the “just losing” and “just winning” candidates.

We describe two techniques for selecting which ballots and precincts need to be audited. The first technique has the benefit of a constant sample size given the number of ballots, the margin of victory, and the desired level of confidence. Sample size may vary with the second approach, but that approach is more amenable to extensions that we propose later.

3.1 Constant Sample Size Method

The hypergeometric distribution describes the number of bad ballots an auditor can expect to find when sampling without replacement. Assume that auditors desire a confidence level c that no fraud significant enough to change the election’s outcome occurred. By [9], given N total ballots and a minimum of B incorrect ballots, the probability mass function of the hypergeometric distribution dictates a minimum sample size, n , of:

$$n = \min \left\{ u \mid 1 - \prod_{k=0}^{u-1} \frac{N - B - k}{N - k} \geq c \right\} \quad (1)$$

A machine recount of a precinct is only necessary if a ballot will be selected for manual verification in that precinct. For this method, auditors could use the initial electronic tallies to perform a mock ballot selection before the machine recount. Any precinct which would have contained a chosen ballot given the mock selection will undergo a machine recount. Following the machine recount, officials may randomly select a single ballot from each recounted precinct and randomly draw the

⁴This is a mild misuse of the term plurality system: other forms of plurality voting for multiple candidates exist [2].

remaining required ballots from the full pool in all recounted precincts.

3.2 Varying Sample Size Method

Rivest [10] proposes an efficient precinct-based auditing technique in which, rather than drawing a given-size sample from the population of precincts, auditors instead randomly select each precinct with a given probability. The same idea is also useful in the context of ballot-based auditing. Assume that, to change the results of an election, the set of ballots must contain a minimum of B bad ballots. To achieve a confidence level of c that at least one bad ballot will be sampled, auditors may select each ballot with probability p chosen such that $(1 - p)^B \leq 1 - c$, or $p \geq 1 - (1 - c)^{1/B}$.

To determine which precincts need to be audited, we may calculate the probability that one or more of the v_i ballots in precinct i will be sampled as $1 - (1 - p)^{v_i}$. Auditors may select each precinct based on the probability that it contains a sampled ballot. If so, officials perform a machine recount in that precinct. Given that at least one ballot is sampled in a precinct, the probability of sampling k ballots in that precinct is:

$$\frac{\binom{v_i}{k} p^k (1 - p)^{v_i - k}}{1 - (1 - p)^{v_i}} \quad (2)$$

Following the machine recount, officials randomly select the precinct’s sample size based on this distribution.

3.3 Comparison to the Method of Rivest

Assume use of the audit method in Section 3.2, and let $p = 1 - (1 - c)^{1/B}$. The probability that precinct i requires a machine recount is therefore $1 - (1 - c)^{v_i/B}$. If an adversary can steal any number of votes in a precinct, Rivest [10] proposes a logistic precinct-based approach that yields the same precinct audit probability. For machine-assisted auditing, however, auditors need only manually review a subset of the recounted ballots.

Rivest presents his logistic approach as a non-optimal heuristic [10], so the usefulness of this link seems limited. Furthermore, the percentage of votes in a precinct that one may steal without generating suspicion is more likely 10–20% than the 100% assumed here. In light of this, a performance comparison between Rivest’s optimal precinct-based techniques and our methods under realistic circumstances would be informative.

4 Evaluation

To evaluate the efficiency of machine-assisted auditing (and ballot-based auditing in general) versus precinct-

based auditing,⁵ we consider both techniques in the context of available data from Virginia’s November 2006 elections, both for local and statewide races.⁶ In all cases, we seek a 99% confidence level.

Virginia contains 2,599 precincts and approximately 4.6 million registered voters, nearly 53% of whom cast ballots during the November 2006 election. The general election decided nineteen issues: four statewide issues, including a U.S. Senate race and several statewide initiatives, and fifteen smaller races, such as U.S. House races. In addition, voters considered numerous local ballot issues, ranging from city council elections to school constructions projects [14]. Because auditing is typically both more important and more labor-intensive in closer races, we focus on such races, excluding consideration of races for which modification of 10% or more of the ballots would have been necessary to change the outcome. This choice rules out many of the races but leaves a set of 49 remaining. Seven of those remaining were general election issues and forty-two were local issues.

The remaining general election issues include a U.S. Senate race with a margin of victory of 0.39%, four U.S. House races, a race for the Virginia House of Delegates, and a state constitutional amendment. For those races, machine-assisted auditing would require a manual review of approximately 437 ballots on average—0.06% of the 796,469 average total ballots. Only the smaller House of Delegates race would require review of greater than 1% of the ballots (1.05%), and five of seven races require audit rates under 0.1%. Precinct-based auditing would review approximately 177,849 ballots on average—22.33% of the average total ballots. In each case, precinct-based auditing requires an expected hand count of more than 42 times as many ballots. The closely contested U.S. Senate race would require review of 2,337 of 2,370,445 ballots with machine-assisted auditing and 1,141,900 ballots on average with precinct-based auditing.

While less overwhelming, the results for local ballot issues are highly favorable as well. In this case, machine-assisted audits would review approximately 224 ballots on average—2.28% of the 9,842 average total ballots. Precinct-based audits would require manual review of

⁵For machine-assisted auditing, we use the techniques in Section 3.1. For precinct-based auditing, we use the methods and assumptions in [13]: auditors choose precincts uniformly at random, an adversary may switch no more than a set percentage of the votes in a precinct without arousing suspicion (we use 10%), and the adversary may switch votes in the largest possible precincts.

⁶We consider all races from the available Virginia data [14]. Some local races are absent, so we ignore those. Due to minor absences in the data set, we assume that no voter submitting a ballot abstains from voting on an issue and that voters for multi-seat races submit multiple ballots rather than a single ballot with multiple selections. While these assumptions slightly affect the realism of the tests, they likely had only a minor impact on the overwhelming results.

approximately 3,928 ballots on average—39.91% of the average total ballots. Only five of the forty-two races would require a manual review of more than 50% of the ballots with machine-assisted audits. In contrast, only six of the forty-two races would require a review of *less* than 50% of the ballots on average with precinct-based audits. Precinct-based audits would require a complete recount in more than half of the cases.

The races that are particularly difficult for machine-assisted auditing are town council, city council, and school board races with 7/492, 5/849, 12/769, 7/246, and 3/2409 margins of victory—requiring manual review of 68.3%, 78.4%, 53.4%, 68.3%, and 90.0% of ballots respectively. In each of these cases, precinct-based auditing would require a full recount.

If comparing machine-assisted audits and precinct-based audits purely on the number of manual ballot reviews, these results indicate a conclusive advantage for machine-assisted audits.

5 Extensions

One way to further reduce the number of ballots that need to be verified by hand is to take into account the ballot contents when sampling. Consider a two-candidate mayoral race in which electronic results indicate that Alice beat Bob 10,001 to 10,000. Traditional audit techniques would require that officials consider ballots containing votes for either candidate even though the primary objective is to discover whether any votes for Alice should have been for Bob. Examining only ballots reported to contain votes for Alice could cut auditor work nearly in half, as auditors seek to discover an equivalent amount of fraud in a far smaller pool of ballots. In any two-candidate single-seat race, auditors may restrict their consideration to the subset of ballots claimed to contain votes for the the winning candidate. By considering the contents of ballots, officials may reduce the number of manual verifications required in nearly any race. In the full version of this paper, we generalize these ideas to races with arbitrary numbers of candidates competing for any number of seats.⁷

The full version of this paper also presents a number of additional methods for increasing the efficiency, practicality, and utility of machine-assisted audits:

- We suggest existing work and tweaks to our techniques to address the possibility that auditors might misread a paper ballot.

⁷Similar tricks may also be useful given only reported initial electronic vote tallies: for example, a precinct containing only votes for Bob could not have contributed to discrepancies affecting the election’s outcome, so both machine-assisted and precinct-based auditing could ignore that precinct entirely.

- As some precincts may encounter delays in reporting results, we propose a means of using educated or worst-case estimates to begin the audits in a timely manner for other precincts. (We compensate appropriately if the estimates turn out to be wrong.)
- We argue that the possibility of occasional machine malfunction need not prevent application of machine-assisted auditing.
- We discuss complementary work that allows candidates to either gain additional assurance or, given the small possibility that the audit process failed to detect significant fraud, uncover that fraud themselves.
- Auditors might wish for ballots from certain precincts to be selected with higher probability than ballots from others, and we explain how to do so.
- Finally, we remark on a continuum that exists between pure precinct-based auditing and pure machine-assisted ballot-based auditing.

6 Conclusion

A well-designed audit process assures the public that an election's outcome is the product of voters' intentions, not fraud or voting machine flaws. By adding a novel machine-assisted recount procedure to ballot-based audits, we can enjoy the efficiency benefits of those audits while avoiding privacy concerns and retaining the security benefits of combined paper/electronic solutions. Our tests using data from Virginia's November 2006 elections confirm the efficiency advantages of machine-assisted audits, and the extended techniques that we propose promise to reduce even further the number of ballots that need to be inspected by humans.

Though future work is needed to better estimate the costs of machine-assisted audits and to assess other practical challenges that election officials face, we believe that the techniques in this paper offer a promising alternative to traditional precinct-based auditing and warrant further study.

Acknowledgments

We thank Andrew Appel, Ari Feldman, Shirley Gaw, David Wagner, Harlan Yu, and Bill Zeller for helpful discussions. We also thank Ronald Rivest and Joe Hall for pointing us towards useful papers and resources and the anonymous reviewers for helpful comments.

Calandrino performed this research while under appointment to the Department of Homeland Security (DHS) Scholarship and Fellowship Program under DOE contract number DE-AC05-06OR23100. All opinions expressed in this paper

are the authors' and do not necessarily reflect the policies and views of DHS or DOE.

This material is based upon work supported under a National Science Foundation Graduate Research Fellowship. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- [1] 110TH CONGRESS. H.R. 811: Voter confidence and increased accessibility act of 2007.
- [2] ACE ELECTORAL KNOWLEDGE NETWORK. Plurality/majority systems, 2006. <http://aceproject.org/ace-en/topics/es/esd/esd01/>.
- [3] APPEL, A. W. Effective audit policy for voter-verified paper ballots in New Jersey, February 2007. <http://www.cs.princeton.edu/~appel/papers/appel-nj-audits.pdf>.
- [4] DUNN, S. Voter verifiable paper audit trail pilot project, Cobb County, Georgia, November 2006. http://www.gaforverifiedvoting.org/docs/Cobb_county_pilot_report.pdf.
- [5] FELDMAN, A., HALDERMAN, J. A., AND FELTEN, E. Security analysis of the Diebold Accuvote-TS voting machine. In *Proc. 2007 USENIX/ACCURATE Electronic Voting Technology Workshop (EVT '07)*.
- [6] JOHNSON, K. C. Election certification by statistical audit of voter-verified paper ballots, October 2004. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=640943.
- [7] KOHNO, T., STUBBLEFIELD, A., RUBIN, A., AND WALLACH, D. Analysis of an electronic voting system. In *Proc. 2004 IEEE Symposium on Security and Privacy*, pp. 27–42.
- [8] NEFF, C. A. Election confidence: A comparison of methodologies and their relative effectiveness at achieving it, December 2003. <http://www.votehere.net/papers/ElectionConfidence.pdf>.
- [9] RIVEST, R. L. On estimating the size of a statistical audit, November 2006. <http://people.csail.mit.edu/rivest/Rivest-OnEstimatingTheSizeOfAStatisticalAudit.pdf>.
- [10] RIVEST, R. L. On auditing elections when precincts have different sizes, April 2007. <http://people.csail.mit.edu/rivest/Rivest-OnAuditingElectionsWhenPrecinctsHaveDifferentSizes.pdf>.
- [11] RIVEST, R. L., AND WACK, J. P. On the notion of “software independence” in voting systems, July 2006. <http://vote.nist.gov/SI-in-voting.pdf>.
- [12] SALTMAN, R. G. Effective use of computing technology in vote-tallying. Tech. Rep. NBSIR 75-687, National Bureau of Standards, March 1975.
- [13] STANISLEVIC, H. Random auditing of e-voting systems: How much is enough?, August 2006. <http://www.votetrustusa.org/pdfs/VTTF/EVEPAuditing.pdf>.
- [14] VIRGINIA STATE BOARD OF ELECTIONS. General election — November 7, 2006. http://www2.sbe.virginia.gov/web_docs/Election/results/2006/Nov/htm/index.htm.