Of Paper Trails and Voter Receipts

Alec Yasinsac* and Matt Bishop+

Abstract

The Internet pervades virtually every aspect of our daily lives, and it seems there is no area that is immune from computing solutions. Computers can do things faster, with greater precision, more reliably, etc., etc., etc. Ironically, one area that most needs the mechanical rigor offered by computing solutions seems destined to abandon electronic solutions and return to paper as the operating medium of choice. As electronic voting falls from favor across America, we are concerned to hear talk of paper receipts provided to voters[1]. Though the department store receipt model is appealing in its simplicity, we posit that when this model is applied to voting systems, it introduces a complex combination of dangerously conflicting properties. We describe these properties and offer an alternate framework to address paper receipt concerns. We then extend this notion into a discussion of paper records and their contribution to forensics for election systems.

1. Introduction

At the end of the day, elections are about counting votes. Since computers have always been particularly good at counting, it seems logical that computers offer great promise in improving vote count accuracy. The now-infamous "hanging chads" of the 2000 presidential election and the 2002 "Help America Vote Act [1]" triggered a mass exodus of elections officials transitioning from paper ballot systems to computer-centric and computer-aided digital vote capture and count models.

Fueled by public reports [2, 3, 4, 5, et al.] that electronic voting machines are prone to malicious manipulation, public discomfort levels are rising. This discomfort is founded in the difficulty of gathering confirming digital evidence available in Direct Recording Electronic (DRE) voting systems and in electronic voting systems overall.

Even as the field of digital forensics expands and new capabilities emerge at a breakneck pace, digital examination is bounded by fundamental computing limitations. The recent and rapid expansion of electronic voting leaves many questions regarding the magnitude of these limitations. For these reasons, many voting integrity advocates encourage a return to the familiarity of paper records to reestablish public trust in the electoral process. Even in the absence of rigorous study of the security properties of paper records, momentum builds to capture every vote on paper in some form.

Proposed federal and state initiatives could mandate paper trails in all elections covered by their jurisdiction. In the U.S. Congress, HR 811, sponsored by New Jersey Democratic Congressman Rush Holt and introduced in the third consecutive session, recently moved from committee to the Congressional floor and could energize a companion bill in the Senate.

One stumbling block to widespread paper trail acceptance is disagreement regarding the type of paper trail that should be required. Subtle properties such as durability, reliability, lifetime, print clarity, simplicity, privacy properties, and voter-friendliness have caused some paper trail advocates to oppose the otherwise popular Holt bill. These discussions are healthy and will ultimately result in defining important properties for voting system paper records that will effect voting system record keeping and audit policies and forensics opportunities that the paper records enable.

Another group of voting paper record properties face unfortunate misconception in this debate. Many voters, and even election advocates, mistakenly utilize the term "receipts" when referring to voting system paper records. While some have proposed systems that may provide voter receipts [6, 7, 8], such systems are largely academic exercises and are not considered for wide spread use.

In this paper, we address misconceptions about voter receipts and show that existing voting paper record systems do not carry with them properties that are integral to receipts. We further show how digital evidence and paper records can provide complementary parts in the voting systems forensics process. We further show that they lack the properties essential for digital forensics information to reconstruct the events that occur on an electronic voting system.

2. Defining "Paper Receipt"

It certainly seems like a simple concept: a receipt is a printed record of a transaction, traditionally a transaction

* Florida State University. This work is funded in part by Department of Defense grant H98230-06-1-0232 and Army Research Office grant DAAD19-02-1-0235.

+ University of California, Davis

[1] e.g. http://www.wired.com/politics/security/news/2003/11/61298

where something is received (usually a payment) by the party that provides the receipt (usually a vendor). When we go to the grocery or department store and make a purchase, we are given a paper record of the financial transaction...it's a receipt; what's hard about that? If we can get a receipt when we buy a pair of shoes, why not when we cast a ballot, particularly since the vast majority of ballots are cast on, or into, computers that could easily print ballot receipts?

2.1. Paper Receipts and Voter Privacy

At face value, it seems reasonable to many that we should simply print a copy of each voter's paper ballot, and let them take it with them as receipt for their votes. The canonical reason that a voting receipt cannot be given is that the receipt may allow a voter to prove how they voted to some third party. Preventing such proofs protects against two related voting irregularities: vote selling and voter coercion. In the former, the theory is that if a voter cannot prove how they voted, there is no viable model for wide-scale (or *wholesale*) vote selling. Conversely, if each voter received a receipt complete with their name and their ballot selections, an unscrupulous operative may simply offer to pay for receipts that reflect a pre-designated voting selection pattern[2].

Similarly, voter receipts can also facilitate vote coercion. If official receipts exist, a corrupt government official, employer, or other miscreant may demand to see the receipts under threat of harm, job loss, or other coercive method.

Traditionally, voting system developers have gone to great lengths to prevent any mechanism that allows voters to prove how they voted, though the rapid expansion of vote-by-mail systems challenges this fundamental voting principle. Still, several scientists continue to propose receipt mechanisms, largely based on cryptography, that allow voters to verify that their votes were properly cast, while not facilitating voter coercion or vote selling [6, 7, 8].

2.2. Foundation for a Valid Receipt: Connecting a Person to a Transaction

Receipts are ubiquitous in society today. With a rich history in documenting cash financial transactions, paper receipts are now used to record document payments of all types, including electronic credit and debit payments. While many institutions advocate a transition to electronic receipts, smart card entries, and other electronic acknowledgements, paper receipts are still the dominate mechanism for documenting financial transactions.

From a standpoint of societal acceptance, paper receipt ubiquity is a self-perpetuating situation. Paper receipts are inherently simple and people are comfortable with them and consider them as near-perfect security items. Their comfort is reinforced with the pervasive receipt environment, where essentially every transaction is accompanied by a receipt. This confidence is a fundamental element provided by receipts.

In addition to providing transaction confidence, receipts also provide evidence that can be used to correct errors in the transaction that may be detected after the fact. It is not uncommon to find an overcharge among items listed on a long receipt and to use the receipt to return to the store for a refund (or to find an undercharge and be faced with a common moral dilemma: to pay it back or not to pay it back).

Perhaps the most common use of a paper receipt is to allow a valid purchase transaction to be reversed, possibly due to change of heart (or maybe when the buyer recognizes that the color just doesn't work for her after all). Few vendors will provide a refund without the paper receipt.

In each of these cases, the paper receipt provides evidence of a transaction involving a buyer, a seller, the specific merchandise that changes hands, and the transaction amount. Each of these items is essential to the transaction. Presentation of the corresponding paper receipt by the buyer to the seller along with the subject merchandise constitutes verification of the transaction at the source. The foundational notion is that the parties to the transaction are able to fully validate both the occurrence and the precise nature and terms of the transaction. Thus, the adjustments are enabled by the precise transaction record and reconstruction of its primary elements.

If we draw a parallel between the voting transaction and a purchase, the natural correlation is that the voter serves the buyer's role, the supervisor of elections is the vendor, the currency is the voter selections, and receipt is the ballot, paper trail, or other perpetual voter selection record.

We point out that these two models diverge substantially here. In the purchaser/vendor model, the vendor collects the information (and/or cash) that they need at the point of sale in order to ensure payment, while the purchaser retains the paper receipt. In the present voting model, the elections official captures the voter selections at the point of sale (in the voting booth) and then the elections official also retains the paper receipt.

2.3. Paper Trails as Constrained Data Items

In any integrity-critical system (hereafter "critical system"), there are sensitive documents and mechanisms that embody the vital protected-system aspects. In information integrity theory, these sensitive items are termed Constrained Data Items, or CDIs.

There are many different types of sensitive information in elections systems. With DRE systems where voter

[2] While the voting pattern itself could be used as a signature that unscrupulous elections workers could identify among collected ballots, this attack requires insider cooperation and, comparatively, limits the attack magnitude.

responses are electronically captured and recorded such as with touch screen systems, the touch screen device itself is a CDI, as are any removable media that were connected to any election-related component during the election period. Paper ballots and other paper trails may also be CDIs.

All CDIs must be rigorously protected, as they themselves are security vulnerability points. Accordingly, an important security goal of critical systems is to reduce the number of CDIs. This reduces both the system security cost and the system security vulnerability.

While each ballot is a CDI, elections officials may attempt to logically reduce the CDI count by accounting for batches of ballots, rather for individual ballots. For example, ballots may be left in their original ballot box until the ballot boxes reach a central collection point. Thus, if the ballot boxes are protected from tampering, the ballots inside are also safe.

Traditionally, elections officials go to great lengths to protect paper ballots, since in most voting systems, paper ballots are the official vote record. Unfortunately, paper record protection is an inherently manpower intensive operation and is correspondingly subject to human error and to individual or collaborative malice. Three typical threats to paper vote records, all of which may be accidental or malicious, include: (1) lost or destroyed legal ballots, (2) altered legal ballots, and (3) injected illegal ballots. Elections officials create checks and balances to reduce or mitigate these threats; still, voting history is rife with carefully documented[3,4] and anecdotal[5,6] evidence of mishandled ballots before, during, and after elections.

2.4. Too Much of a Good Thing: Misusing Audit Data to Overturn a Valid Election

In elections systems, auditing is a double-edged sword. When rigorously engineered and methodically executed, audits can detect anomalies during and after elections, and can add significant confidence to the electoral process. Conversely, audit mechanisms that are incorporated into the system that they verify are potentially subject to attack themselves.

Audit mechanisms are designed with the primary goal of detecting manipulations of the target system. Thus, if this primary audit goal is met, invalid results, or manipulations of valid results, will always be detected.

On the other hand, there is little attention paid in the literature to protecting election audit systems themselves from manipulation. If the audit system is not properly protected, a malicious attacker may attack the audit system in order to overturn a valid election, not by attacking the election result itself, but by attacking the election audit system. Worse, what little data is present can no longer be used to perform even rudimentary forensics, as its reliability and trustworthiness is now suspect.

To illustrate, consider the process of election verification. Verification is a strong result of an audit. If two independent processes produce identical audit results, the likelihood that both of these independent mechanisms are incorrect, and incorrect in precisely the same amount, is low. If a third mechanism, independent of both previous approaches, is added, the agreement gives greater confidence in the result. Thus, were cost not a barrier, we might seek to install many independent mechanisms so that confidence can approach perfection. Sounds pretty promising, were this the end of the story.

The complication occurs when the independent audit mechanisms do not agree on the result. In this situation, some reconciliation process must take place to determine which, if any, audit result is correct. Reconciling conflicting audit results is never easy, since a deviation in any result draws all results into question. In order to reduce such conflicts, we may choose to minimize the number of audit mechanisms, limiting ourselves to the strongest methods, i.e. to the mechanisms that are scientifically the most sound, or are the easiest to implement correctly.

This begs the question of how to identify the "strongest", "easiest", or "most accurate" mechanisms, where we recognize the added complexity that qualification brings to the table. Clearly, a strong audit mechanism naturally detects system anomalies a high percentage of the time. We may call this property "low false negatives", as it means the mechanism is highly unlikely to certify or verify an invalid or anomalous result. Audit system developers sometimes layer two or more filtering mechanisms to prevent false negatives so that if an anomaly erroneously passes through one or more filters, it may be detected in subsequent layers.

2.5. False Positives in Audit Systems

Certainly providing audit systems with low false negative rates is critical. We contend that the converse is also true; that audit mechanisms must also ensure a low false positive rate, meaning that the audit system does not (accidentally or maliciously) falsely indicate that an anomaly has occurred. For example, consider a paper ballot scanning system. If the voting standard specifies only that a bubble be filled in to indicate the selected candidate, it would be counterproductive to utilize a dual-audit system where one mechanism detects any opaque marking in the bubble, while the second mechanism detects only, say, a number two pencil mark filling. Such incompatible detection

[3] Washington/Politics Section, "Judge upholds Washington governor's election", USA Today, June 6, 2005

[4] http://www.sfgate.com/cgi-bin/article.cgi?file=/chronicle/archive/2002/02/11/MN209475.DTL

[5] http://archives.cnn.com/2000/ALLPOLITICS/stories/11/08/ballotbox.found/

[6] http://findarticles.com/p/articles/mi_qn4196/is_20001128/ai_n10656662

mechanisms would inject natural conflict, and its corresponding doubt, into the respective counts.

As a short case study, consider the elegant voter receipt offered by David Chaum [7]. The security of his "Secret Ballot Receipts" is based on visual cryptography [9]. In short, Chaum's scheme creates two visual cryptography ballot shares on translucent paper layers. One of the shares is selected by the voter and is retained by the voter as their receipt. The selected share by itself (or along with the receipt's included administrative information) does not reveal any information about the voter's selections.

Chaum's scheme is complete with proofs that the system cannot produce false negatives. As the author states:

> "...if your receipt is correctly posted [to the elections bulletin board or web page], you can be sure (with acceptable probability) that your vote will be included correctly in the tally. A receipt that isn't properly posted is physical evidence of a failure of the election system, and a refusal by officials to post it is an irrefutable admission of a breakdown in the election process [7, p. 40]."

Digital signatures are also employed, again seemingly to allow the voter to prove that their receipt matches a ballot: "If the signature doesn't pass, the physical receipt is direct evidence of system failure [7, p. 44]."

We ask if it is possible to forge a receipt in this system. Counterfeit prevention is difficult and expensive, particularly in this age of desktop publishing, requiring sophisticated watermarking, intricate printing, expensive papers, etc. For example, if counterfeiting is possible, voters may print their own receipts that do not match any ballot, and may show up *en masse* in the days after the election to dispute the election results with these receipts that reflect "...irrefutable admission of a breakdown in the election process".

It turns out that Secret Ballot Receipts are counterfeit-resistant, made so by signatures computed over the share, serial number, etc. [7, see the "More Formally" inset on p. 44]. This feature is not discussed as counterfeit detection, nor false positive prevention in the paper. It is offered only as a step in the algorithm to generate the receipt. The fact that it can prevent false positives is treated as an aside, or possibly as an obvious result.

Though the scheme is resistant to externally generated counterfeit receipts, we suggest a more subtle and sinister electronic threat triggered by malware; malicious software that may infiltrate the system through a virus, exploited buffer overflow, or other software intrusion approach. Means of malware infiltrating real voting systems have been identified, several of which are detailed in [10]. Once an attacker is able to inject custom software into the system, they may make minor operational changes that cause the software to create some percentage of receipts that can be visually acceptable at the polling place, but that will not match thereafter when the voter attempts to verify their receipt. For example, the software adjustments may cause the code to offset the electronic share version that is posted on the election web pages or bulletin board, or to mix serial numbers on posted receipt layers.

As with many demonstrated electronic voting system attacks, this attack may target specific candidates, or candidates for a specific party in multiple races. Such an attack could cause widespread (false) vote-fraud allegations. More strongly, if voters present compelling (although false) vote-fraud proofs, they may be able to overturn an election even though there is no anomaly in the original vote count and that original vote count is completely valid.

This suggests an important observation here: "More information is *not* always better than less information". Officials must be careful when selecting audit mechanisms to ensure that the mechanisms themselves are accurate, in the sense that they have a low false positive rate, a low false negative rate, and that they are secure against malicious manipulation that could cause either false positives or false negatives.

2.6. The Transparency-Security Paradox

A canonical approach to enhancing voter confidence is to make the voting process as transparent as possible. Unfortunately transparency has properties that are inversely proportional to security; that is, it is sometimes necessary to sacrifice transparency to ensure security. For example, many CDIs disappear from public observation by policy to ensure that they are adequately protected from post-election, accidental or malicious manipulation. When a ballot is scanned into a precinct optical scanner, the ballot may completely disappear from sight, allegedly stored intact in the opaque container that holds the scanning device. After the polls close, the containers themselves are removed from public observation to be locked in a truck or van for transport, or taken to a secure room. The containers and the ballots within them may not be available for public viewing until they are opened by elections officials, possibly days or weeks after the election if there are no audits involving that election.

These normal election procedures are designed to protect the ballots and are routine and necessary to prevent post-election manipulation. Nonetheless, they ensure that the election process is fundamentally *not* transparent.

3. The Paper Trail as Forensic Evidence

The goal of the voting phase of an election is for elections officials to collect and tabulate votes in a way that provides strong evidence that the reported results of every race are correct. This evidence must be conclusive to Secretaries of State who certify the elections and also be

convincing to the voting population. Moreover, the evidence must stand up against rigorous and contentious examination by (potentially armies of) lawyers in state and federal court systems in order to make voters confident in the reported results. It was natural that terms such as "E-voting Forensics" and "Election Forensics" emerged to capture the need for rigorous analysis of elections.

Forensic examination can address confidence at all levels, but due to its expense, it is most often triggered only for contentious races where the victory margin was small, and these occurrences are rare. This results in a classic paradox that is all too familiar to security professionals: the "Return-on-Investment (ROI) Challenge".

In general, forensic information is more detailed and voluminous than audit data. Thus it is expensive to gather and to retain. Since we cannot know *a priori* which elections will require forensic investigation, elections officials must capture forensic information for all races on the ballot. Most of this information will not be used. The ROI Challenge is to quantify the value of gathered forensic information that is not used. The challenge is even greater if forensic examinations using the gathered information do not always provide a definitive, compelling result.

3.1. Properties of Paper Evidence

Contrary to public perception, as security mechanisms go, paper receipts do not have strong security properties. Paper, printers, and ink are all widely available and desktop printing enables average citizens to produce store-quality receipts from the comfort of their homes. Since they are reasonably easy to forge, paper receipts do not have strong non-repudiation properties. These properties can be extended for more consequential transactions, simply because paper is not an inherently secure medium.

3.1.1. Forgery

The distinction between public perception that paper is a strong audit mechanism and the reality that paper does not have strong properties results from a failure to understand the contexts in which paper is used. The public sees that many important functions are founded on paper documents. For example, birth and marriage certificates are traditionally paper documents. Additionally, titles to our automobiles and deeds to our homes are routinely recorded on paper. In each of these examples, there is a physical person, place, or thing that the paper represents, which is fundamentally different from a ballot.

Currency, on the other hand, reflects two fundamental ballot properties that titles, deeds, and marriage and birth certificates do not have: (1) currency is abstract, not connected to any physical entity; and (2) currency is anonymous, not inherently tied to any individual.

The public accepts paper bills as currency without reservation. While counterfeiting schemes routinely show up in news stories once or twice per year, they involve such a small percentage of the currency inventory that few citizens have been directly impacted by counterfeit currency.

The difference in context between ballots and currency is that, unlike ballots, paper currency is reused, with a lifetime of several years. Thus, it is cost effective to invest significant resources into creating each bill. Currency is rigorously engineered and tediously manufactured to have strong, easy to verify anti-counterfeiting properties. The paper formula and production process are painstakingly designed and these designs are closely guarded. The inks used in currency face similar protective processes and mechanisms, and may provide forensic properties that can be illuminated with special lights or even when held up to the sunlight.

Ballots, on the other hand, are used once and must be produced in sufficient volume to ensure that all citizens that desire to vote have a ballot available. In precincts that utilize paper ballots, paper and printing costs dominate electoral budgets. Even a small percentage increase can devastate local budgets.

3.1.2. Loss

Paper records have an air of "fire and forget"; that is paper has no inherently traceable properties. For example, if you drop a piece of paper at some point on a walk and realize you have dropped it when you get home, there is nothing inherent to that paper that will help you trace it. If the paper was a one hundred dollar bill, we have no way to locate that particular bill in the currency pool. Even if we were to see the bill again, the only characteristic that could tell us that it was ours is the serial number.

Paper ballots and other paper trail artifacts share this fire and forget property, though with paper ballots, even serial numbers cannot be used. If ballots are lost, there is nothing inherent in the paper to lead investigators to them. Moreover, when they are found, there is nothing inherent in ballots to verify their authenticity.

We routinely hear stories of overseas ballots lost in the mail[7], misplaced[8], or never delivered to the voter[9]. Some of these ballots are never found or are found too late to be useful. Those lost ballots that are recovered may cast a shadow over the entire electoral process since, at a minimum the chain of custody was violated.

A perfect illustration of lost ballots occurred in 2004, where a case of "lost and found ballots" resulted in an election reversal in the Washington Governor's race in the 2004 election [11].

[7] http://news.bbc.co.uk/2/hi/americas/3960679.stm

[8] http://seattletimes.nwsource.com/html/localnews/2002122945_webballot17.html

[9] http://www.worldnetdaily.com/news/article.asp?ARTICLE_ID=15601

3.1.3. Replacement

Another weakness of using paper as an audit or forensic medium is that paper is susceptible to replacement. In testimony before a House sub-committee hearing [12] Michael Shamos highlighted the ballot replacement threat [paraphrased]: "Once a voter deposits their ballot, they have no guarantee that their paper ballot will actually be counted, or that their paper ballot will be present later for any recount". As we said earlier, protecting paper records is an inherently human process that is subject to human failings and to human malice.

We illustrate with a hypothetical example. Assume that an election worker with access to legal ballots desires to rig an election. If this worker can arrange to be in charge of transporting ballot containers to a central location where initial counts are conducted (which seems to be a common elections practice), they may one-for-one replace legal ballots with ballots prepared for this purpose. Since ballots cannot have unique identifiers, there would be no way to show that the replacement ballots are not valid.

Alternatively, the attacker may even replace the entire ballot box with a different, pre-prepared box, if given suitable planning opportunity. Again, there is nothing inherent on the paper ballots themselves that can prevent or detect this type of ballot swap.

3.2. Retail Versus Wholesale Fraud

Many issues raised to this point identify small scale, retail fraud that occurs in the polling place. Electronic voting mechanism properties highlight the opportunity for wide-scale elections fraud that could occur if malicious parties could control or influence the software that operates voting machines. Because software is so flexible, theoretically attacks could be written to favor a selected party in every locale where a specific vendor and software version are used.

Voting machine attacks are not limited to wholesale fraud, as a myriad of precinct, and even single machine level attacks are documented in the literature. However, these attacks are universally accomplished before or during the voting period. To date, we have not seen wholesale or retail attacks that are effective against electronic voting systems that can occur after the results are reported.

Conversely, paper-based elections systems seem to be inherently resistant to wholesale fraud. While we can envision sophisticated attacks that leverage subtle, subliminal messages in reams of blank ballots or ballot printers that systematically print names of candidates from one party slightly darker than those from other parties, no such attacks are considered serious threats.

To summarize, electronic voting systems are subject to wholesale attacks before and during elections, but are resistant to both wholesale and retail fraud after the polls close. On the other hand, paper-based systems are wholesale fraud resistant before, during, and after election day, but are subject to retail fraud in all periods. Table 1 also captures these comparative properties.

This observation naturally leads to a discussion of prioritization between audit media, which is presently a contentious issue among voting activists. The prevailing logic in the Holt bill is that there should always be a paper representation of each vote and that the paper record should be the ballot of record. That is, if the electronic and paper counts differ, all other things being equal, the paper count dominates.

Examination of Table 1 suggests that a different precedent structure may be more appropriate. Electronic vote counting is accepted as the most efficient process for establishing the election night results. Since paper systems can mitigate wholesale fraud during the election, it follows that paper should be the foundation for parallel audits, conducted during the election, right up until the results are reported.

	Paper		Electronic	
	Election	Audit	Election	Audit
Retail				x
Wholesale	x	x		x

Table 1. Comparative inherent attack resistance

Once the results are reported, electronic ballots provide the strongest fraud protection. While some claim that auditing electronic results is meaningless, our experience is quite the opposite. Post election auditing is possible and meaningful for electronic systems, even though it is dramatically different from paper based audits. Processes can be verified, logic paths can be checked and tested, and executables can be compared to baseline versions. Its effectiveness requires that forensic data be gathered during the election. We have presented some attributes of this data above; in the next section, we suggest some information to record, and some problems that arise in doing so.

Moreover, data from parallel audits can provide invaluable forensic information. We emphasize that audits that are conducted during the election with an immediate purpose of detecting anomalies on the fly, can provide essential information regarding the electronic vote count during post-election forensics. In combination, electronic system review coupled with results from parallel paper-based audits can provide strong evidence of a valid result, or can pinpoint anomalies.

4. The Way Ahead

One thing that became painfully clear from the 2000 presidential election is that voting systems were in a state of disarray. We have since iterated through three "silver bullet"

voting solution paradigms that each have reflected the common wisdom on voting systems for some period of time:

(1) Electronic Voting Machines

(2) Voter Verifiable Paper Trails

(3) Hand-Marked, Optical Scan-Counted Ballots

We are not presently close to reaching an effective, stable solution. Moreover, we are convinced that any effective, stable voting solution will support strong audit trails that allow auditors to forensically verify the entire voting process. For this, we offer some observations regarding voting systems and their audits.

4.1. Relaxing the "Vote Non-Provability" Principle

If allowed to occur, vote selling and voter coercion can devastate democracy, and we cannot count on laws to prevent these activities. However, we are seeing an increasing trend toward shifting protection from the voting system itself to human procedures in elections offices with the rapid expansion of absentee and vote-by-mail systems. If this type of protection is ultimately considered acceptable by the body politic, it offers opportunities for several procedural changes that can facilitate audits.

For example, one simple mechanism that can reduce post-election ballot-handling anomalies is to mandate use of numbered ballots. Traditionally, voting systems have spurned any attempt to number ballots, even if the number is difficult to visually understand, e.g. if the number is encapsulated in a bar code. However, similarly to the way we manage our personal checkbooks, serialized ballots would allow us to more easily detect and investigate missing ballots, and similarly, duplicate or invalid ballot numbers are easily detected and the ballot numbers can provide valuable investigative information.

Watermarking ballots can similarly help control post-election ballot manipulation. As with numbers, watermarks can prevent illegal ballots from being injected into an election. However, watermarking is not as promising as serializing ballots since it can be expensive and watermarking does not help detect missing ballots.

4.2. Other Marked Ballots Protection Mechanisms

Another issue that has come up in recent elections is whether ballots should include a "None of the Above" or "No Vote (NV)" option in each race. Candidates perennially oppose allowing voters to make this type of selection to indicate that having no one fill the position is better than having one of the available candidates fill the position.

Public perceptions aside, such a positive-intentioned record is vital to forensic investigation, particularly where undervotes are involved. The NV option would eliminate the unintentional undervote that is exacerbated by assigning an assumed intention to a non-action.

Improperly marked ballots may be reduced or eliminated by precision, machine marked ballots. Touch screen devices and other mechanical voter input devices have yet to pass the test of time. Nonetheless, they have the potential to provide both expanded accessibility and precise voter intent capture. When combined with machine marked ballots, these systems offer significant opportunity for high quality forensic information.

4.3. Independent Mechanisms

Corroboration occurs when independent mechanisms agree on a result. Such mechanisms are appropriate for both paper ballot, optical scan systems and direct recording devices that capture voter intent through touch screen, pointing device (mouse), audio feed-button response devices, sip-and-puff technology, or other electronic response systems.

4.3.1. Paper Based Independent Mechanisms

One approach to providing elections audit information is to combine a mathematically-based electronic count with a paper-based ballot count. This is a tenant of the voter-marked, optical scan count-recount paradigm.

As we mentioned earlier, a challenge to this paradigm is to ensure that no ballots are added to, or deleted from, the ballot box between the original count and subsequent recounts. We may create an independent mechanism for this purpose by adding a second device to precinct scanners that computes an independent value that uniquely identifies each ballot. Such a value may be based, for example, on a hash of a computation of ballot pixels, similar to the Chaum method [7]. Patterned or watermarked ballots could guarantee collision protection in such a scheme.

Bloom Filters [13] provide an efficient computation that balances collision resistance against storage requirements. Additionally, checking ballot inclusion against Bloom filters is a simple process that does not reveal any information about voter selections.

4.3.2. Direct Recording Electronic Independent Mechanisms

Some may claim that it is impossible to capture independent evidence regarding ballots cast on DREs. Of course, ingenuity can overcome many limitations. For instance, an approach offered in [10] is to have a program written by other than the voting machine vendor, digitally capture all voter actions (screen touches and button presses) taken during the voting period.

Using such a machine activity log and the corresponding ballot definition file, it is possible to mechanically reconstruct the votes cast from each machine. Moreover, given access to the machine from which the log was attained, an auditor has significant digital evidence regarding potential errors that may have occurred (e.g. mis-calibration, etc.).

Also consider a less voluminous, more semantic digital log that records a per-voter, touch history complete with its semantic interpretation. Such a log may contain entries such as those in Table 2. From this log we can tell that:

(1) Voter 23 did not vote for any candidate in Race 2

(2) Voter 7 did not vote for any candidate in Races 1 or 2

(3) Voter 11 did not cast their ballot.

Voter 23	Voter 7	Voter 11
Select cand A, Race 1	Page forward	Select cand A, Race 1
Page Forward	Select cand E, Race 3	Select cand C, Race 2
Select cand D, Race 3	Select cand G, Race 4	Page forward
Page forward	Page forward	Select cand D, Race 3
Page forward	Select cand E, Race 7	Page backward
Cast ballot	Select cand E, Race 6	Deselect cand C, Race 2
	Cast ballot	Page forward

Table 2. Semantic voter action audit log

These techniques open a covert channel to identify voters, because the voter can use a specific sequence of screen touches to indicate identity. Relaxing the "vote non-provability" principle may ameliorate this threat, as might several classical defenses to limit the bandwidth of covert channels.

5. Conclusion

Audit trails are essential to voting systems. In this paper, we examine issues related to paper-based audit mechanisms to support voting systems and show that the phrase "paper receipt" is not appropriate to voting systems in wide spread use today.

We also identify a pitfall to audit systems that has received no interest in the literature: false positives, and illustrate why audit systems must consider the possible impact that malicious activity could have on these systems during the post election forensic process.

Finally, we suggest mechanisms and approaches that may help voting forensics in the future by combining the complementary properties of digital forensics and paper based mechanisms. While we anticipate many advances in voting technology and procedures, none will be more important than techniques that facilitate forensic investigation in close elections.

6. References

[1] Help America Vote Act of 2002, Public Law 107-252, 107th Congress, USA

[2] Tadayoshi Kohno, Adam Stubblefield, Aviel D. Rubin, and Dan S. Wallach, "Analysis of an Electronic Voting System", IEEE Symp. on Security and Privacy, May 9-12, 2004, pp. 27-40.

[3] Harry Hursti, "The Black Box Report, SECURITY ALERT: July 4, 2005, Critical Security Issues with Diebold Optical Scan Design"

[4] David Wagner, David Jefferson, Matt Bishop, "Security Analysis of the Diebold AccuBasic Interpreter", Voting Systems Technology Assessment Advisory Board

[5] Harry Hursti, Diebold TSx Evaluation, SECURITY ALERT: May 11, 2006, Critical Security Issues with Diebold TSx

[6] Ronald L. Rivest, "The ThreeBallot Voting System", Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, Oct. 1, 2006

[7] David Chaum, "Secret-Ballot Receipts: True Voter-Verifiable Elections", IEEE Security & Privacy, Jan/Feb 2004, pp. 38-47

[8] Andreu Riera, Josep Rifà, Joan Borrell, "Efficient construction of vote-tags to allow open objection to the tally in electronic elections," Information Processing Letters 75 (2000) 211–215

[9] M. Naor and A. Shamir, "Visual Cryptography," Proc. Advances in Cryptology (Eurocrypt 94), A. De Santis, ed., LNCS 950, Springer-Verlag, 1995, pp. 1–12.

[10] A. Yasinsac, D. Wagner, M. Bishop, T. Baker, B. de Medeiros, G. Tyson, M. Shamos, and M. Burmester, "Software Review and Security Analysis of the ES&S iVotronic 8.0.1.2 Voting Machine Firmware, Final Report", Security and Assurance in Information Technology (SAIT) Laboratory, Florida State University, February 23, 2007, http://election.dos.state.fl.us/pdf/FinalAudRepSAIT.pdf.

[11] Washington/Politics Section, "Judge upholds Washington governor's election", USA Today, June 6, 2005

[12] Michael Shamos, Testimony Before the Committee on House Administration, US House of Representatives, Sept. 28, 2006.

[13] Burton Bloom, "Space/time trade-offs in hash coding with allowable errors," Communications of ACM, 13(7):422-426, July 1970.