State-wide Elections, Optical Scan Voting Systems, and the Pursuit of Integrity

Tigran Antonyan, Seda Davtyan, Sotirios Kentros, Aggelos Kiayias, Laurent Michel, Nicolas Nicolaou, Alexander Russell, Alexander A. Shvartsman

***Abstract*—In recent years, two distinct electronic voting technologies have been introduced and extensively utilized in election procedures: direct recording electronic (DRE) systems and optical scanner (OS) systems. The latter are typically deemed safer, as they inherently provide a voter verifiable paper trail that enables hand-counted audits and recounts that rely on direct voter input. For this reason, optical scanner machines have been widely deployed in the United States. Despite the growing popularity of these machines, they are known to suffer from various security vulnerabilities that, if left unchecked, can compromise the integrity of elections in which the machines are used. This article studies general auditing procedures designed to enhance the integrity of elections conducted with optical scan equipment and, additionally, describes the specific auditing procedures currently in place in the State of Connecticut. We present an abstract view of a typical OS voting technology and its relationship to the general election process. With this in place, we lay down a "temporal-resource" adversarial model, providing a simple language for describing the disruptive power of a potential adversary. Finally, we identify how audit procedures, injected at various critical stages before, during, and after an election, can frustrate such adversarial interference and so contribute to election integrity.**

We present the implementation of such auditing procedures for elections in the State of Connecticut utilizing the Premiere (Diebold) AccuVote optical scanner; these audits were conducted by the UConn VoTeR Center, at the University of Connecticut, on request of the Office of the Secretary of the State. We discuss the effectiveness of such procedures in every stage of the process and we present results and observations gathered from the analysis of past election data.

***Index Terms*—Electronic voting, election, optical scan, audit.**

I. Introduction

Optical Scan (OS) systems are the most widely used electronic voting equipment in present United States elections. Indeed, 57% of counties nationwide (corresponding to roughly 60 million voters), incorporated OS usage in the November 2008 Presidential Elections, with 41 out of 50 states using some type of OS machine in at least one of their counties [23]. Table I presents an overview of the types and the usage of OS systems in these elections. The OS systems rely on an optical ballot reader to scan voter ink markings on specially-designed paper ballots. An important benefit of optical scan technology is that it naturally yields a voter-verified paper audit trail (VVPAT)—the ballots marked by voters. This enables hand-counted audits and recounts that can directly assess voter intention.

The other major voting option is based on Direct Recording Electronic (DRE) technology, where voters record their votes using touch-screen electronic machines. The DRE devices can be equipped with printers that can produce a printed record. However, establishing a verifiable connection between the voters' choices and the DRE-printed records is a logistical and technological challenge that is beyond the ability of currently deployed DRE technology. This may be among the main reasons why DRE technology is not as widely adopted as OS technology [19], [4].

Following the widespread adoption of electronic voting equipment in order to comply with the Help America Vote Act (HAVA) [10], several research efforts identified security concerns with electronic voting technology (e.g., [20], [3], [18], [24], [5], [12], [25], [9]). Some of these concerns apply to OS technology [16], [17], [7], [11], [25], [2], revealing important security flaws and vulnerabilities and, in several cases, describing specific attacks that could interfere with election integrity.

TABLE I
Types and Usage of Optical Scan Voting Machines in the 2008 US Presidential Elections

Machine Name (Manufacturer)	# of Counties	% of Total OS Usage
Model 100 (Election Systems & Software)	670	36.16%
AccuVote-OS (Premier Election Solutions (Diebold))	398	21.48%
Model 650 (Election Systems & Software)	238	12.84%
Optech III-P Eagle (ES&S / Sequoia Voting Systems)	197	10.63%
eScan (Hart InterCivic)	105	5.67%

A general election process is an enormously complicated process involving elaborate distributed coordination of personnel, procedures, and equipment. The problem of ensuring integrity is one that must necessarily involve such disparate issues as equipment custody, voting day procedures, election official selection and training, voter training, tabulation procedures, and, finally, faithful behavior on the part of the actual physical apparatus. In this article we focus solely on the technological aspect of an election and, in particular, posit an *adversarial model* for elections that focuses on the "electronic" dynamics of the election. We proceed by identifying the characteristics that govern the family of optical scan systems and we incorporate an *election process schema*, embraced by any election that deploys optical scanners. Based on the derived election process we define adversarial strategies in terms of

Voting Technology Research Center and Computer Science and Engineering Department University of Connecticut,Storrs, CT 06269, USA.
Emails: {tigran,seda,skentros,aggelos,ldm,nicolas,acr,aas}@engr.uconn.edu

Research funded by the Secretary of the State of Connecticut.

the chronological stage and the resource of the election they exploit. To tackle and limit the effectiveness of various adversaries we propose injection of auditing procedures in critical stages of the election process. We include, as a case study, our work with the Accu-Vote Optical Scan tabulators used in the State of Connecticut. We present the implementation of the proposed technological auditing procedure by the UConn VoTeR Center that was used in recent elections in Connecticut, complementing the hand-counted audits performed by the State and analyzed by the VoTeR Center. The overall process includes testing, comparison, and analysis of the data collected during the audits. We conclude by presenting the results and useful observations extracted from the application of the auditing process.

Background. We begin with a summary of some previous security evaluations of OS systems.

The AV-OS voting terminal has been the subject of the report of H. Hursti [11], pointing out that the AV-OS memory card lacks cryptographic integrity checks, thus potentially leading to serious security vulnerabilities that can be exploited with the help of specialized (third party) hardware. These findings lead many jurisdictions employing the AV-OS to insist that memory cards be sealed in the terminal with a tamper-evident seal for the elections and further require that terminals be delivered to and returned from polling locations with such seals in place.

The Connecticut Secretary of the State commissioned a follow up study to confirm Hursti's findings and identify other vulnerabilities. The study by the UConn VoTeR Center [16], [17] identified an additional attack that can be successfully launched against the AV-OS even if the memory card is sealed in. Here the attack exploits the flawed authentication on a communication port of the machine and results in transparently modifying the contents of the memory card. This was made possible because no cryptographically authenticated data are transmitted between the terminal and the election management system (GEMS). The same attack also exploits vulnerabilities of the machine's proprietary language, called `AccuBasic`, used for reporting election results.

Previously it was assumed that the firmware of voting terminals in general, and of the AV-OS terminal in particular, can be treated as a trusted component of the system. However the report [7] proved this assumption to be incorrect, showing that someone with physical access to a terminal may manipulate firmware in a way that will be undetectable during election day testing.

In a study of the ES&S M100-OS optical scanner commissioned by the State of Ohio [2], the authors identify and report problems that affect critical components of the system. Deficiencies discovered in those systems illustrate ineffective protection of firmware and software, and insufficient cryptography and data authentication. The vulnerable components include the removable devices (PCMCIA memory cards) that contain sensitive election data, and the firmware code responsible for the functionality of the OS terminal. Based on the attacks, the unauthenticated election data can be altered using commonly available systems equipped with a PCMCIA reader/writer. Moreover, the unauthenticated firmware image on that same memory card can also be maliciously modified; such firmware images are loaded by the OS terminal without hardware and/or password authentication. Other threats concern the centralized election management system, called *Unity*, used for the programming and electronic tabulation of the election results. The authors show that the software suffers from undetectable buffer overflow attacks; these enable an attacker to gain access control on the database that stores the sensitive election data.

In [15] a series of measures are proposed for auditing elections, including a) comparison of poll center turn-outs with the number of ballots cast, b) comparison of the number of ballots cast with total of votes cast, c) hierarchic comparison of results during tabulation, d) auditing of the chain of custody, e) recounts and f) parallel testing. The first four suggestions are procedural measures that should and can be applied in all elections. The later two are aimed to address weaknesses introduced in elections by the adoption of new technologies. Thus the first four measures, although necessary, fall out of the scope of this work and will not be further discussed. The later two measures are discussed in Section IV.

These and many other findings underscore the importance of a methodical approach to deploying voting technology in ways that ensure election integrity.

Contributions. Our goal is to derive a theoretical framework that describes the general family of optical scan voting technologies and their deployment in elections. Based on that framework we aim to identify security vulnerabilities of such election systems and to propose effective solutions that prevent or limit the possibility that any of those vulnerabilities can be exploited. Though the principal focus of this paper is OS election systems, some of the procedures presented may naturally find applications in DRE voting technologies. In more detail, we present the following in this report.

1) We examine the general architecture of a group of OS election systems, identifying a) Election Management System Software, b) Optical Scan terminal, and c) Central Tabulator.
2) Based on the proposed OS-based election model we define and illustrate the process that any election that deploys OS terminals should follow. This process is independent of any state-specific processes and we recommend that it is embraced by any audience that uses such systems within any electoral procedure. The process identifies the flow of information (i.e., election information, counter information, executable code, etc.), as well as the interaction of external entities with the electronic equipment during the process.
3) Given the proposed election process we identify the attack-prone components and we divide the election process into three chronological periods: a) Before Election, b) During Election, and c) After Election. Based on this division we describe and define the characterization of an adversary in terms of the time during which an election can be affected and the resource(s) that the adversary attempts to exploit. Some known attacks on OS systems are presented and expressed in terms of our adversarial model.

4) Once we identify potential problems in the election process we present means of preventing or limiting the possibility of election corruption. We suggest the injection of auditing procedures in critical stages of the process to cover most of the spectrum of possible technological exploits. For each suggested procedure, we present the potential adversaries it foils by analyzing the time periods and the resources affected by an adversary.
5) We present real world application of a subset of the proposed audit procedures, as implemented by the UConn VoTeR Center on request of the Connecticut Secretary of State. Our team has participated in examining and auditing a number of elections for the State of Connecticut that deployed the AccuVote-Optical Scan (AV-OS) system manufactured by Premiere Election Solutions (formerly Diebold). As a case study, we present the development performed and the steps followed by the team to ensure accurate and timely analysis of the critical components of the AV-OS with the aim of preserving integrity of the elections.
6) Finally, we present and discuss the results of the audits in which we participated for the November 2008 elections. In particular, the audits validate the previous anecdotal evidence that a non-trivial percentage of memory cards used with the AV-OS terminal contained corrupted and unreadable data. Furthermore the analysis reveals procedural misconceptions and a certain lack of adherence to the established electoral procedures.

II. Computational Model and Adversity

An electoral process that deploys electronic OS election systems should provide security guarantees that are analogous to an electoral process utilizing sealed envelops and a ballot box. There are obvious advantages of using OS election systems, including fast generation of tally reports and the auditability of the election process. However, the use of OS election systems also introduces new adversarial possibilities: ones that exploit the new components of the electoral process. In this section we introduce a general model for an OS electoral process and define the adversaries that could interfere with such a process.

Before proceeding into the details of our proposed adversarial model we present a set of security and integrity properties that should characterize a general election process. We categorize them into *procedural* and *technological* characteristics. The first category refers to properties that will be enforced due to the procedures carried out by the participating entities, while the latter deals with properties that are supposed to be provided by the equipment used during the elections.

Privacy (Procedural+Technological). The voting system should ensure the privacy of the ballots in the sense that it should be impossible to extract any information about a voter's ballot beyond what can be inferred from the published tally. One can see that a combination of procedures at the poll center and careful design and use of the technology are needed in order to ensure this property.

Ballot Verifiability (Technological). The voting system should assure the voters that their correct voting preferences are reflected in the cast ballots, i.e., that each ballot was cast as intended. In the case of an OS voting system, cleartext paper ballots are always used and barring any other issue in terms of interface design they capture the true intent of the voter. Still, incorrectly printed ballots (e.g., circumstances where ballot layout is inconsistent with the ballot processing equipment) can lead to effective loss of the voter's intent.

Voter Verifiability (Procedural+Technological). The voting system should enable the voter to challenge the procedure in the post-election stage and verify that his/her ballot was included in the tally. This property is sometimes hard to achieve (though not impossible [6]), due to the fact that it interferes with the *Receipt-freeness* and *Coercion Resistance* properties presented later. OS systems are generally not designed to provide Voter Verifiability.

Universal Verifiability (Procedural+Technological). The voting system should enable anyone, including an outsider, to be convinced that all valid cast votes have been correctly counted in the final tally. The existence of an auditable paper trail in OS systems gives a natural way to verify that cast votes have been properly included in the final tally. Indeed, the major thrust of this article is to describe how this property can be achieved assuming trustworthy auditors with appropriate election access. We note that the trust placed on auditors has a two-fold benefit: on the one hand, it relaxes security issues of privacy and coercion that arise when verifiability is open to the public (that in part may act adversarially). On the other hand, it is—in principle—consistent with current election safety practices that rely election monitoring by trusted organizations (e.g., the Organization of American States, the Organization for Security and Co-operation in Europe).

Voter Eligibility (Procedural). The voting system only permits eligible voters as listed in the electoral roll to cast a ballot. At the same time, the system should ensure that no eligible voters are disenfranchised. These characteristics are enforced by the official electoral procedures, and OS voting systems are not concerned with it. Or, to put it differently, once the voter is standing in front of the machine he/she is assumed to be eligible from the machine point of view.

One-Ballot-One-Vote (Procedural+Technological). The voting system should not permit voters to vote twice. Guaranteeing that one voter casts one ballot is a procedural issue, on the other hand guarantying that each ballot is counted only once is a technological issue in OS systems.

Fault tolerance (or Robustness) (Procedural+Technological). The voting system should be resilient to failures within the formally specified tolerances for each item of equipment and its components or parts.

Fairness (Procedural+Technological). The voting system should ensure that no partial results become known prior to the official end of the election procedure.

Receipt-freeness (Procedural+Technological). The voting system should not facilitate any way for voters to prove how they voted. OS electoral systems are not generally designed to enforce this property. While in an OS procedure no receipt is furnished to the voter, optical scanners read only specific areas of a ballot, leaving many options for someone who wants to produce a mark or identify their ballot. This weakness of the

OS systems can be alleviated by procedures. For example, policies can be put in place that prevent public access to the paper trail of the election.

Coercion Resistance (Procedural+Technological). The voting system should not enable anyone to coerce voters to vote in a certain way. This can be provided procedurally, through careful supervision of the polling places, and in conjunction with the receipt-freeness characteristic as it is a necessary property for coercion resistance.

A. General Model of an OS Election System

We now establish a general computational model for the election systems that use optical scan voting machines. We aim to identify and list all the components that provide an exact characterization of an OS system. In general, an OS voting system consists of the following major components: (a) Election Management System Software, (b) Optical Scan Terminal, and (c) Central Tabulator. A schematic representation of the OS election system model and the interaction between its components appears in Figure 1. Below we explain in greater detail the aforementioned components.

1) Election Management System Software: The election management system software (EMSS) is responsible for: a) maintenance of the election data, b) design and production of the ballot sheets and c) delivery of election and execution data to the optical scan terminal.

Election Data: Election data describe the details of a particular election including candidates, races and precinct details. The EMSS is responsible to store such data, usually using a database, and provide the flexibility to the election officials to update the data accordingly.

Ballot Sheets: Every OS machine (independent of the manufacturer) should have a corresponding software that allows the design of the paper ballot sheets for a particular election. This is also one of the responsibilities of the EMSS. The system is responsible for the mapping between the ballot layout and the election information, and designing a paper ballot readable by the optical scan terminal it is designed for. Note, that each paper ballot may require different markings depending on the optical scan terminal for which the ballot is designed (e.g., filling/blackening a circle, completing a broken arrow, drawing a line through a rectangle), however, the idea remains the same.

EMSS ***and Optical Scan terminal communication:*** Finally, the EMSS maintains means of communicating with the optical scan terminal for information exchange. Data exchanged between the EMSS and OS includes election information, ballot layout, and executable code. Each system provides its own communication medium, for example, a serial communication port. The communication can be also facilitated through the removable media that is used by the terminal.

2) Optical Scan Terminal: The optical scan terminal consists of: a) hardware components, including input/output devices, b) executable code, and c) removable programmable media. The OS terminal itself may be thought of as the most technologically vulnerable component of an OS election system since it is movable to and from the polling place, it spends substantial periods of time in potentially unattended storage, it is exposed to the voters and other personnel during the election periods, and it is responsible for the collecting and locally storing the election votes.

Hardware and Input/Output Devices: A typical OS terminal is comprised of an on-board processor, fixed memory/storage, optical scanner, electro-mechanical ballot handling devices, printer, and other peripheral and input/output devices, all in a single enclosure. Users of the OS terminal may input or retrieve information through peripherals, attached on the OS terminal.

Input devices are mainly used to activate specific functions on the terminal, for communicating with external sources and for scanning voting ballots. Naturally the ballot reader falls into this category. The reader can be characterized based on: a) the type of ballots it recognizes, and b) the volume of ballots it can read per time unit.

Output devices are used for informative, reporting and troubleshooting purposes. For example, an LCD display would provide the status of the machine and present conditional queries to the users. A printer would be used to print election totals, zero counter reports, vote receipts or even audit log details.

Executable code: Perhaps the most critical component (along with the removable programmable media) is the executable code of the OS system. The executable code is responsible for any behavior and/or computation performed by the machine. It controls the output and the input devices and presents or collects sensitive information, during the voting process. Included in the executable code is the operating system, which for some machines is embedded in the hardware (e.g., AV-OS), while in others it is stored in removable media (e.g., M100-OS, Optech-OS). Code not embedded in hardware is usually dynamic and election dependent. Thus such code may be generated and transferred to the system (usually by the EMSS) at the beginning of each election process, and remains unchanged throughout the election it was intended for.

Removable Programmable Media: Every OS system contains a programmable memory storage device that provides the flexibility of reprogramming the machine with multiple and different election data. Examples of such a programmable media includes the EPSON memory card used in AV-OS and the PCMCIA cards used in M100-OS and Optech-OS. The contents of the programmable memory can be divided into four major logical parts:

(a) Vote Totals Memory (VTM): This is the part of the memory where the election totals are kept. In some cases, such as Optech-OS, this can be a separate memory card, while in other instances, such as in AV-OS and M100-OS, it is combined with election information into one memory card.

(b) Election Information Block (EIB): In some cases (Optech-OS) this block is on a separate memory, while in other cases (AV-OS, M100-OS) it is combined with vote totals into one physical memory card. All the information regarding an election, including precinct, races, parties, candidates and ballot layout is kept on this memory block.

(c) Event Log (EL): A space in the programmable memory is reserved to record all the actions involving the machine during

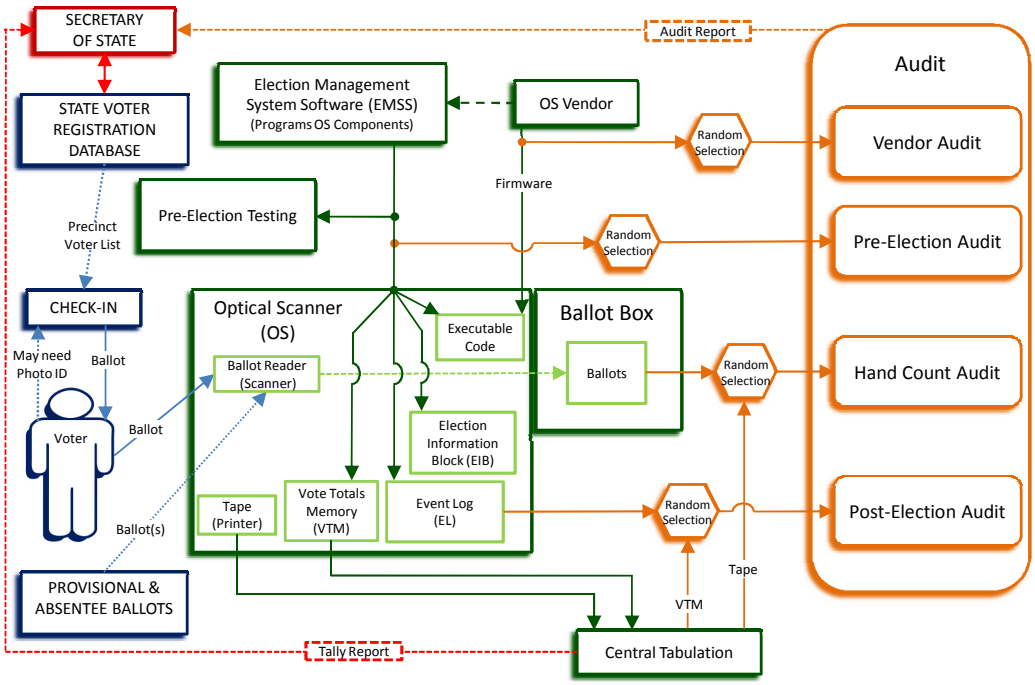

Fig. 1. ELECTION PROCESS DIAGRAM FOR OPTICAL SCAN ELECTION SYSTEMS

the election procedure. (The events that a machine may log may or may not be adequately implemented, depending on the specific voting system.)

(d) Executable Code: Removable media may be used to store executable code. The code might play a modest role, such as regulating the printing process, where in other cases it may serve as the critical application of vote tabulation, or be used to update the firmware of the machine. In general every OS system involves some customizable code whose purpose is to comply with the election parameters for each electoral district; this is managed by EMSS and transferred to the system in advance of the election.

3) Central Tabulation: Each OS election system includes a central tabulation process or mechanism that can be devised as either manual or electronic process; in case of the latter, it can be implemented in hardware and/or software. The purpose of the tabulation is to collect and tally the election results that were accumulated and/or counted by the individual OS terminals.

Electronic Tabulation: A software central tabulator provides the capability of tallying the results uploaded from multiple OS terminals. Sometimes this function is provided by the EMSS system. The election data can be conveyed from the OS terminals to the tabulation system by various communication means, for example, using a communication port, via a telephone connection, or by means of removable media. Some tabulators employ high speed scanner voting terminals and are used to count batched voting results, as in the case of the absentee votes. This class of tabulators can be included in hardware tabulation systems, where their executable code is embedded in the hardware of the voting terminal (e.g, M650-OS).

Manual Tabulation: Manual tabulation avoids the extra communication between the terminal and an external tabulation system, and instead relies on the printed results extracted from the output devices on the OS terminal. Along these lines, the results may be collected for each individual terminal and then the election officials proceed to parse and tally the results manually.

B. Modeling the Election Process

Figure 1 in Section II-A also presents the election process flow when using an OS election system. We next describe the election flow in more detail.

Before The Election Day. Election preparations begin at least 30 days prior to election day. The programmable components of OS machines are programmed for each precinct. The machines also undergo routine maintenance and testing to detect failures within the design parameters of the test function. Once the programmable components, i.e., the Election Information Block (EIB), Vote Totals Memory (VTM), Event Log (EL), and optionally Firmware are ready, they are securely transported to the polling locations and installed into the OS machines. Elections officials then conduct the specific pre-election tests on all the machines.

On Election Day. The following activities take place on the election day.

Before The Polls Open: On the morning of the election day, before the polls open, the poll workers and/or registrars of voters need to verify any seals present on each OS machine, ensure they are not tampered, set the machine(s) to "election mode" and verify that the machines are properly initialized that includes making sure all candidate counters are set to zero.

While The Polls Are Open: Each eligible voter is entitled to a single ballot that they get once they are verified against

the voter registration database. Once the voter fills the ballot he/she proceeds to feed the ballot to the machine's optical scanner.

After The Polls Close: If electronic tabulation is possible, the election officials remove the media with the VTM data. If manual tabulation is the chosen method, the election officials print the totals report directly from the OS machine. In some jurisdictions both methods are used: After printing the totals report from the OS machine, election officials remove the media with the VTM data. The printed tape and/or the VTM media is delivered to the central tabulation process where the totals are computed and reported to the authority, e.g., the Secretary of the State Office, for certification. Usually the central tabulation is done on municipal or county levels.

C. Modeling Adversity

We characterize our adversarial model in terms of the chronological election periods and the election resources they exploit. As mentioned before, we concentrate on technological attacks that affect the "forensic" data trail of an election. Non-technical and social engineering issues are outside the scope of this study. Following the election process presented in Section II-B we first identify the time frames and resources that may be affected by an adversary that intents to interfere with the proper conduct of an election. An electoral process can be divided into three time periods:

Pre: Pre-Election, up to the point the polls open.

In: In-Election, from the time the polls open and till the election results are certified.

$Post$: Post-Election, after the election results are certified. Adversaries that perform their attack during elections typically have restricted computational power, operate within small windows of opportunity to perform their attacks, and control a small subset of the resources. On the other hand, pre-election or post-election adversaries, can have unlimited computational power and can control a wide variety of resources. For example the pre-election adversary may be able to replace some or all of the ballots in a precinct, replace one or all the memory cards (removable media) of a precinct, or even compromise the programming of the EMSS system.

Each adversary may control one or more of the following resources: (i) $EMSS$ – the software and/or the communication of the EMSS system, (ii) $Ballot$ – the paper ballots used for voting, (iii) $Media$ – the removable media that contains the election information, totals counters, executable code, and Event Log (EL), (iv) $Machine$ – the OS machine, and (v) $Tabulator$ – the Central Tabulator/Tabulation.

An adversary is defined by the time period he launches an attack and by the resource it controls. For example we denote by $\mathsf{A}^{\mathsf{Pre}}_{\mathsf{Media}}$ the adversary that launches a pre-election attack on the removable media. We define an adversarial strategy A as a collection of adversaries. For example $A = \{\mathsf{A}^{\mathsf{Pre}}_{\mathsf{Media}}, \mathsf{A}^{\mathsf{In}}_{\mathsf{Media}}, \mathsf{A}^{\mathsf{Post}}_{\mathsf{Tabulator}}\}$, is an adversarial strategy that tries to corrupt the removable media before and during the election and tries to affect the tabulation system after the election. An adversarial strategy signifies that an adversary attack can occur at different moments and leverage one or more resources. The objective of such an adversarial strategy may be to compromise one or more of the properties discussed in the beginning of Section II. We focus exclusively on attacks that are enabled by the introduction of optical scan technology (procedural attacks being outside of our discourse).

III. Security Vulnerabilities in Elections using OS Election Systems

This section presents security vulnerabilities that are introduced by the use of optical scan systems. It demonstrates that along with the adoption of a new technology, new procedures should be added in the electoral process to compensate for the technological vulnerabilities. We do not intend to be comprehensive for the election process as a whole. Here we focus on the attacks targeting the technological aspects of OS voting systems without considering procedural attacks. For instance, attacks that erase media or destroy ballots by breaking chain of custody can seriously affect the auditability of the election but are beyond the scope of this paper.

Media vulnerability. The first vulnerability exploits the existence of removable media in an OS machine. A removable media provides the needed flexibility to customize the equipment from election to election. As explained in Section II, this media holds election information, such as counters, ballot layout, and sometimes executable code responsible for the presentation of the results. It may even include the operating system for the OS machines or its subcomponents. The current implementations were shown (cf. [2], [11], [17], [16]) to lack cryptographic integrity and authenticity, rendering the media vulnerable to attacks. Such attacks were demonstrated for both ES&S M100 OS in [2] and for the AV-OS in [17], [11], [16]. Attacks can occur prior to the election and target the media or the EMSS system and correspond to the adversarial strategies $\left\{\mathsf{A}^{\mathsf{Pre}}_{\mathsf{EMSS}}, \mathsf{A}^{\mathsf{Pre}}_{\mathsf{Media}}\right\}$. A strict custody policy for the OS machine and its media can curtail $\mathsf{A}^{\mathsf{Pre}}_{\mathsf{Media}}$. Given its prominence in the process, the EMSS system itself should be physically secured and only handled by trustworthy parties. Note that, even if cryptographic protocols are in use, a successful attack against the EMSS system could always compromise memory cards. Pre-election testing in the poll centers, along with pre-election, post-election and hand counted audits, can limit the capabilities of such an adversary.

Attacks based on media vulnerability target the ballot and universal verifiability. More specifically they attempt to alter the final election result, by affecting how a vote from a ballot is counted, or how the totals counters are interpreted. The existence of the paper trail (physical ballots) is the best defense against those attacks. Note also that such attacks are not always intentional. For example, [1] presents is a partial list of incidents involving vote swapping due to mistakes during the programming of the machines.

Ballot vulnerability. In the case of an adversary of type $\mathsf{A}^{\mathsf{Pre}}_{\mathsf{Ballot}}$, ballots with swapped positions could be injected among blank ballots. Such an attack could be prevented with strict procedural and custody policies, and pre-election testing. Note that in some cases, paper ballots with rotated candidates' positions are used in order to reduce the chance of voter fraud.

Programming errors as the one detected in the Pottawatamie County incident in Iowa during the June 6, 2006 primary elections [1], could have major impact on the election results. Procedures should ensure that pre-election testing is able to confirm that machines are properly programmed for all types of ballots used during the election, in case more than one version of the ballot is used. As mentioned before, strict procedural and custody policies should monitor the printing, storage and shipping of the ballots, to prevent the generation and inclusion of maliciously altered ballots in the polling centers. Such policies should include at least a) random sampling and auditing of the ballot batches that leave the printing facilities, b) sealing of ballots upon arrival with tamper evident seals, c) strict chain of custody for the ballots during transfer and d) random hand-counts after the election.

Firmware vulnerability. As it is pointed out in [7], it is possible to launch an attack by infusing the OS machines with malicious firmware. This would be the case of adversary $\mathsf{A}^{\mathsf{Pre}}_{\mathsf{Machine}}$. For some implementations, the firmware can be flashed directly from the memory card while in other implementations a physical EPROM must be changed (like AV-OS). Audits similar to the ones performed to prevent removable media attacks can also help detecting malicious firmware flashed on the memory card. Such audits may also include pre and post election examination of the firmware with the goal of detecting attempts of trace or event log hiding by the malicious code. The EPROM modification attacks can be foiled by a direct firmware audit that obtains the contents of the EPROM from the audited machine and compares it against a verified system code. Such pre-election and post-election firmware audits are based on white box testing, since they do not rely on the execution of the firmware code, but rather on the direct examination of the firmware content.

Central Tabulation vulnerability. The central tabulation process offers another avenue for attackers. Clearly, any adversary that compromises the central tabulation system itself, e.g., using an adversarial strategy $\mathsf{A}^{\mathsf{Post}}_{\mathsf{Tabulator}}$, can invalidate the integrity of tallying. Similarly, an adversary that gains access to the partial tallies (as reported on the printed tape or the electronic VTM) while they are being transferred to the central tabulation system would achieve the same result. In general, depending on the way central tabulation is performed it could be possible to introduce unauthenticated results to the tabulation process or selectively suppress the incorporation of some of the actual results. Attacks of this type can only be defeated through procedural means; in the case of electronic central tabulation it should be ensured that only valid election results are incorporated into the tallies by authenticating the VTM data as well as ensuring that no real VTM data are dropped.

DRE systems. Although the vulnerabilities introduced here are presented and analyzed for OS systems, some of them can directly apply to DRE election systems, specifically, media, firmware and central tabulation vulnerabilities. One limitation in some DRE systems is the lack of a paper trail as it removes the option of a hand count as a counter measure. In this setting, the ballot vulnerability can be associated with the calibration of the screen (in case of touchscreen DREs) or swapping the labels with candidate names (in case of a machine with buttons). As before, strict procedural and custody policies, and pre-election testing can be applied in order to prevent such attacks.

IV. Auditing Scheme for Integrity

The introduction of optical scan technologies into the electoral process creates new opportunities for potential adversaries who wish to interfere with its integrity. It is possible to detect (and therefore deter) such malicious activities and accidental errors associated with the technology by introducing the following seven audit procedures: *Vendor Audit, Procedural Audit, Pre-Election Audit, Parallel Testing, Post-Election Audit, Hand Count Audit, and Meta Audit.*

Audits are valuable in deterring potential adversaries who now face the risk of being detected and possibly lose the ability to conduct their attack successfully. To maximize audit reliability one should conduct the audit on the complete set of resources utilized in an election. Given the large scale elections we are considering, performing such an exhaustive audit may be prohibitively long or expensive. This suggests a random sampling approach for auditing in elections where a complete audit is unfeasible or impossible, cf. [21].

Vendor Audit. The vendor audit is meant to ensure the validity of the executable that the vendor installs in each voting machine. Whenever a new version of the vendor executable code (e.g., the firmware) is released it should be thoroughly examined to detect any malicious code. A sophisticated adversary $\mathsf{A}^{\mathsf{Pre}}_{\mathsf{Machine}}$ with ample know-how and access to voting terminal equipment can design a malicious firmware that has total control of the terminal and can thus corrupt any election characteristics presented in Section II. The purpose of this audit is to make certain that the vendor code complies with its expected behavior, i.e., that it correctly tallies the ballots scanned by the machine.

Procedural Audit. It is important that election officials and any other party involved with the election process strictly follow the safety procedures established prior to the conduct of any election. These procedures may involve chain of custody, pre-election testing of the removable media (at the programming facility as well as in every precinct), pre-election zero-count reports, and post-election tallies with proper machine-generated time stamps (the tallies must be produced after the closing of the polls). While an audit cannot enforce these processes, it can be helpful in exposing protocol violations. In particular, it is helpful to catch in-election adversarial strategies classified as $\mathsf{A}^{\mathsf{In}}_{\mathsf{Media}}$ and $\mathsf{A}^{\mathsf{In}}_{\mathsf{Machine}}$. The first adversary attempts to alter the outcome of the election by, for instance, "stuffing" the counters. The second adversary interferes with the fairness of the election by producing intermediate tallies during the election. Apart from common audit procedures [15], an analysis of the EL can provide useful information regarding the actions executed on and by the terminal, if the media card or the firmware are not compromised.

Pre-Election Audit. Adversaries can also attempt to interfere with the electoral process with pre-election strategies. Only some of the possible attacks on $EMSS$, $Ballot$ and $Media$

might be discovered in such audit. A pre-election audit occurs after the EMSS was used to program the memory card, but prior to the election, and its purpose is to validate the integrity of the data stored on the removable media. The audit procedure achieves this goal by first collecting and then comparing the content of a random sample of the removable media against a trusted database containing the expected media contents. Any adversary $\mathsf{A}^{\mathsf{Pre}}_{\mathsf{EMSS}}$ that controls the EMSS, would be thwarted since a malicious or corrupted piece of data that is loaded on the removable media would be detected. Further, necessary pre-election testing [20], [13] must include the verification of the ballot geometry with respect to the counters, and test the sensitivity of the ballot reader. Provided that pre-election testing is adequate, attacks against the ballots ($\mathsf{A}^{\mathsf{Pre}}_{\mathsf{Ballot}}$), will only work when accompanied by an appropriate modification of the removable media (i.e., an attack $\mathsf{A}^{\mathsf{Pre}}_{\mathsf{Media}}$) to accept the corrupted ballots. Consequently, such an attack will be detected by the combination of the proposed pre-election audit and testing.

An important class of attacks that occur prior to the election are "man in the middle" attacks interfering with the data transfer between the EMSS and the removable media. Once again, any such attack against the removable media ($\mathsf{A}^{\mathsf{Pre}}_{\mathsf{Media}}$) will be caught by the pre-election audit. Section V-A1 goes into more details regarding the audit of the removable media of AV-OS machine.

Parallel Testing. Sometimes pre-election adversaries may launch attacks that are activated during specific time periods in an election process, and remain inactive during testing or audit time periods to avoid detection. In their simplest form, such attacks could, for example, get activated at the time and date that the polls open and become inactive at the time and date that the polls close. Parallel testing [20], [13], [22] follows the black box testing approach and is a good way to detect such adversarial strategies. This testing is designed to simulate the real election and it should be performed with a randomly selected subset of the OS machines that were prepared to be used in the election. In particular, the selected machines follow the same procedures as the machines that are used at the day of the elections, but instead they are fed with specially marked ballots (known to the tester) that are otherwise identical to the ones supplied to the voters. Since any malicious software executed on the machine is not able to detect that it is being tested, it does not alter its behavior and hence it would be detected if it attempts to modify the election results.

Post-Election Audit. Once the polls are closed and the results are tabulated, post-election audit can catch various irregularities in the voting process. If for example an adversary $\mathsf{A}^{\mathsf{In}}_{\mathsf{Media}}$ substitutes the media card during an election this may be discovered by inspecting the event log (EL). Similarly, a different code planted in the media card to produce a biased output can be detected as well. Furthermore, if the central tabulator was corrupted by adversary $\mathsf{A}^{\mathsf{Post}}_{\mathsf{Tabulator}}$, then the examination of the counters on the removable media, in combination with each district's totals, may reveal inconsistencies in the counting procedure. The post-election audit occurs after the central tabulation has occurred. It consists of an analysis of the EL, the election information, and the executable code.

Hand Count Audit. Hand-counting the ballots after the election is useful to detect any discrepancies between the machine counts and the actual votes cast. The audit is helpful to ascertain the accuracy of the scanning device and the reliability of the counting process. Extended testings performed by [13], [14] present inconsistencies in the scanner sensitivity of some OS voting terminals, further motivating this class of audits.

The adversaries covered by this audit include the ones that modify either vote counts or the way they are reported, e.g., $\mathsf{A}^{\mathsf{In}}_{\mathsf{Media}}$ or $\mathsf{A}^{\mathsf{Pre}}_{\mathsf{Machine}}$. Note, however, that due to the fact that the ballot box is a part of the OS machine, an adversary $\mathsf{A}^{\mathsf{Post}}_{\mathsf{Machine}}$ could prevent or invalidate hand count audits by tampering with the ballots and an adversary $\{\mathsf{A}^{\mathsf{Post}}_{\mathsf{Media}}, \mathsf{A}^{\mathsf{Post}}_{\mathsf{Machine}}\}$ could go undetected. Manual Counting may reveal attacks such as counter or candidate swapping, error in totals, errors in the election data, and possibly errors in the ballot layout. The Achilles' heel of this audit lies with its human aspect, and time and financial costs.

Meta Audit. A basic assumption is that the auditor is trustworthy. One may assume, however, that the auditing process itself can be the subject of attacks. It may be desirable to conduct random audits of the auditing process itself to ensure the overall integrity. Note that combination of a variety of audit procedures may eliminate or weaken stronger adversaries and more sophisticated attacks.

V. Audit Scheme Implementation in Connecticut Elections Deploying AV-OS

We now present an implementation of the audit scheme described in the previous section. This implementation was used in several elections in the State of Connecticut, including the November 2008 elections. We survey the approach and highlights the effectiveness of the audit procedure. Additional details of our work in Connecticut can be found in [7], where we focus on the AV-OS terminal using the methodology that fits the general model of OS machine presented in Section II-A. Next we provide a brief description of the components of the AV-OS (explaining the role of each in the computational model), and then we provide details of the audit procedures comprised of removable memory audits (pre and post election) and hand count audits (post election).

A. The AV-OS Election System

GEMS. Global Election Management System, GEMS, is a vendor-supplied system that contains the functions of EMSS and central tabulator. GEMS can be used for ballot design and central tabulation. It is installed and operated from a conventional PC. GEMS uses several databases that include the data, ballot layout, and bytecode corresponding to the precincts participating in the election. This information is transferred via the serial communication port to (and from) the AV-OS terminal containing a memory card. In some states, including the State of Connecticut, the central tabulation feature of GEMS is not used. For the State of Connecticut, an external contractor is responsible for programming the memory cards.

Firmware. The main software component of the AV-OS is the *firmware* executable code stored in an EPROM chip and responsible for all the functions provided by the machine. To extract and process the binary representation of the firmware, we used third party hardware and software components. Obtaining the binary image of the firmware served two purposes. First, it allows us to confirm that the firmware does not contain malicious code and fulfills its intended purpose. Second, it enabled us to determine that we could not rely on the firmware in faithfully obtaining the contents of the removable memory card. To streamline the audit of the cards and to obtain true copy of their contents we implemented custom audit firmware that was used with the AV-OS machines in the the audits. We refer the reader to [8], [7] for further technical details.

Memory Card. The AV-OS terminals uses a 40-pin 128KB Epson memory card. It is installed into the 40-pin card slot (J40 connector) found in the right front side of the terminal. Note that Epson discontinued the production of this memory card, and reader/writers for this memory card are not readily available. The data on the card includes status information, an event log, ballot description, and counters. This was established by analyzing the firmware binary code and the communication between GEMS and AV-OS. Note that our analysis was performed without any technical documentation or source code from the vendor. The structure of the memory card contents is shown in Figure 2.

1) Memory Card Audit: To audit the memory cards we collected three types of data:

(a) Baseline data: Before the elections we used a standard AV-OS, GEMS, and the databases from the external contractor that were used to program the memory cards for the elections. Using these resources we programmed our own memory cards. We then *imaged* the contents of these baseline cards using our data collection methodology.

(b) Pre-Election Data: Prior to the elections the districts were instructed to send a randomly selected subset of their memory cards for testing. We collected images of each of these memory cards using our own tools. This forms the *pre-election data*.

(c) Post-Election Data: Similar to the pre-election data, randomly selected districts were instructed to send us their cards after the completion of the elections. We refer to the data collected from those cards as *post-election data*.

A focal point of the audit was the validity of the data collected and the integrity and reliability of the memory cards as a storage medium. The latter can be partially tested during the data collection as our tools identify cards containing an apparently arbitrary sequence of data values (that we call "junk"), or no programmed data. Below we present the steps taken for testing the pre- and post- election cards. The results and detailed description of the testing appears in Section VI.

Pre-Election Audit. Pre-election audits attempt to identify invalid or maliciously altered memory cards before the election and additionally check that the towns followed the correct testing procedures (based on the events recorded in the logs, and the values of the counters and state variables).

The first concern was to collect a sufficient number of memory cards to obtain a representative sample for our findings. Each polling center received four programmed memory cards the external contractor. According to their instructions each district is supposed to perform pre-election tests of the four cards. Immediately after the testing is complete, they are required to *randomly* select one memory card per district and send it to the University of Connecticut VoTeR Center for pre-election audit testing. The procedure for random selection of the memory card(s) described above only applies to precinct based tabulators and does not include central absentee ballot tabulation. Given the above procedures, each memory card received for pre-election auditing should be in "election mode" with counters set to zero and should have evidence of a pre-election tests in its log.

After collecting the necessary cards from the districts we test the validity of the cards by performing a semantic comparison between the pre-election and the baseline data. The potential problems we are testing for include incorrect ballot data or bytecode, non-zero counters, and incorrect states. Such problems could arise from either malicious attacks, accidents, human error, or failure to follow procedures.

Post-Election Audit. The post-election audit employs a procedure similar to the pre-election audit. The main goal however is to check the validity of the cards after the elections are closed. This audit focuses on the cards used during the actual election. To test the validity of a card, we compare the post-election data of a district with the corresponding baseline data. The status of each card along with the value of the counters were extracted and examined. The integrity and the reliability of the hardware of the memory cards was tested in this audit phase as well. Detailed results appear in the next section.

2) Hand Count Audit: It is a legislated requirement in the State of Connecticut to perform post-election hand count audit of 10% of the districts that are randomly selected after an election by the Office of the Secretary of the State (SOTS). An official hand counted recount is also mandated when the difference between the candidates is 0.5% or less. (We refer the reader to the Statutes of the State of Connecticut for the formal definitions of such audits and recounts.) Note, however, that there is a significant difference between a hand count audit and a recount. The intent of a hand count audit is to determine whether the machine counted correctly according to its specification. The purpose of a recount is to determine the true intent of all voters. For instance, a ballot with bubbles that are circled rather than blackened may count as an under-vote in an audit, while it may be attributed to the circled candidate in a recount, given the unambiguous voter intention.

For a hand count audit of the 10% of the districts, the totals for each candidate are counted and the results of hand counts are reported to the SOTS Office. The returns are then conveyed to the UConn VoTeR Center. Each entry in a hand count audit report represents information about a given candidate. Specifically, each record contains the following: date, district, machine seal number, office, candidate, machine counted total, undisputed hand counted total, questionable hand counted total, overall hand counted total, that is, the sum of undisputed and questionable ballots. Thus for the AV-OS, every record corresponds to the totaled "values" of the specific bubble on the ballot sheet. Hence, our goal is to evaluate the accuracy of

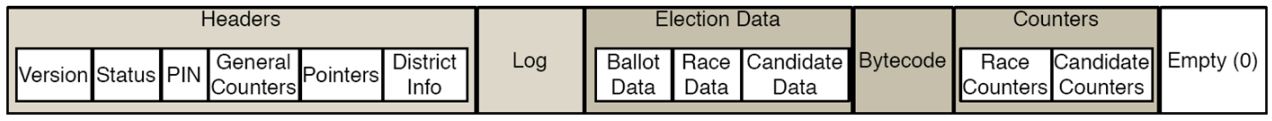

Fig. 2. FORMAT OF THE AV-OS MEMORY CARD

the AV-OS machine in obtaining the totals for each candidate running for a certain office.

In any given race, the difference between hand counted total and machine counted total is called *Discrepancy* (D). The discrepancy can be negative or positive. If the discrepancy is positive then we observe a machine undercount relative to the hand count H, i.e., the machine counted fewer ballots than the auditors. If the discrepancy is negative then we observe a machine overcount relative to the hand count H, i.e., the machine counted more ballots than the auditors.

Note that this assumes that the hand count does not contain (human) errors. This is not necessarily true in reality. In general it is not possible to ascertain whether the hand counted data is error free, and so we assume that the hand counted data is reported correctly. In Section VI-D we take a closer look at the returns received by the VoTeR Center.

VI. AUDIT RESULTS AND OBSERVATIONS

We now present the results of the various recent audits performed in Connecticut. We start by describing in detail the most important state values that are extracted from the memory cards and their meaning (Section VI-A). We then proceed to the presentation of the results for pre-election (Section VI-B) and post-election (Section VI-C) audits of memory cards. We conclude with the presentations of the statistical analysis of the hand counted audit returns (Section VI-D).

A. Audited States of the Memory Card

There are three aspects of interest regarding the states of the memory cards: *(a) Card Format*, *(b) Card Status*, and *(c) Counter Status*.

(a) Card Format: At a high level, the memory cards either contain properly formatted, recognizable data, or contain seemingly arbitrary, noise-like data that we term "junk". The "junk" cards are unrecognizable by the AV-OS. Such cards are readily detected by poll workers during pre-election testing.

On rare occasions it can also be observed that a card—while properly formatted and containing good and usable data—shows a few "specks", that is isolated bytes with unexpected values. These occurred outside the area that is used for election data that is usually filled with zeros. These "specks" are not detected by AV-OS and we have yet to discover an instance where they interfere with normal AV-OS operation.

To sum up, we distinguish the following three card formats: *Good Data (Clean)*, *Good Data (Specks)*, and *Junk Data*.

(b) Card Status: This refers to the current state of the memory card as indicated by a status flag in the header data. We identified the following states: *Not Programmed (Blank), Not Set for Election, Set for Election, Results Print Aborted, Election Closed, Results Sent/Uploaded, Audit Report Printed*.

(c) Counter Status: The Counter Status can be one of the following: *Zero Counters*, *Non Zero Counters*, *Non Zero and Set for Election*. Pre-election cards are expected to have zero counters. Note that a card with non-zero counters that is *not* set for election will be zeroed when the card is set for election. Post-election cards used in an election are expected to have non-zero counters and a status of "closed."

Pre-election cards should, at minimum, be in the Election Loaded state and, ideally, in the Set for Election state; they should never be Set for Election with non-zero counters. Post-election cards should, ideally, be in the Election Closed state and, furthermore, may indicate that the audit log has been printed. Post-election cards should never show Results Aborted or an election not yet closed states. When the election is closed, result printing is begun automatically and the aborted state is only indicated if the printing is canceled, or if the machine was not properly powered down. The Audit Report Printed indicates that the results and the audit log were both printed. For our Connecticut post-election auditing procedure we expected to observe an Election Closed state, since printing the native audit log is not part of the standard procedures.

In the next two sections we present the results of the pre-election and post-election audits of memory cards.

B. Pre-Election Memory Card Audit Results

Table II presents the results of pre-election audits of memory cards conducted in Connecticut for November 2007, August 2008, and November 2008 elections (pre-election audit was not performed for the February 2008 primary). The table shows the frequency of various states observed on the audited memory cards.

To enhance the readability of the data we annotate certain values to identify them as acceptable or unacceptable memory card states. In particular each line preceded by an asterisk "*" represents an expected state, while a state preceded by "x" is not acceptable. We also use double asterisk "**" to identify additional specific observations that are acceptable. The rest of the states are not expected, although they are acceptable because they are easily detectable and do not pose a threat.

(a) Card Format. Table II records the following statistics: *November 2007 Election:* 522 memory cards were audited. Almost 95% of all cards were properly formatted and contained good data. There were in total 18 cards which contained "junk" data, over 5%, or about one in 30. *August 2008 Election:* 185 memory cards were audited, out of which ten cards were identified as "junk," roughly one out of eighteen cards. *November 2008 Election:* 610 memory cards were audited, out of which 54 cards were identified as "junk," roughly one out of eleven.

TABLE II
PRE-ELECTION MEMORY CARD AUDIT ANALYSIS SUMMARY FOR NOVEMBER 2007, AUGUST 2008, AND NOVEMBER 2008 CONNECTICUT ELECTIONS: (A) CARD FORMAT, (B) CARD STATUS, (C) COUNTER STATUS

		Pre-Election Audit November 2007		Pre-Election Audit August 2008			Pre-Election Audit November 2008		
		Num	% Total	Num	% Total		Num	% Total	
	(a) Card Format								
*	Good Data, Clean Card	495	94.8%	175	94.6%		532	87.2%	
	Good Data, Some "Specks"	9	1.7%	0	0.0%		23	3.8%	
	Not Programmed	0	0.0%	0	0.0%		1	0.1%	
	Junk Data	18	3.4%	10	5.4%		54	8.9%	
	Totals:	522	100%	185	100%		610	100%	
	(b) Card Status								
	Not Programmed (Blank)	0	0.0%	0	0.0%		1	0.2%	
	Not Set for Election	218	43.3%	175	100%	**	551	99.1%	**
*	Set for Election	233	46.2%	0	0.0%		4	0.7%	
	Results Print Aborted	11	2.2%	0	0.0%		0	0.0%	
	Election Closed	42	8.3%	0	0.0%		0	0.0%	
	Results Sent/Uploaded	0	0.0%	0	0.0%		0	0.0%	
	Audit Report Printed	0	0.0%	0	0.0%		0	0.0%	
	Totals:	504	100%	175	100%		556	100%	
	(c) Counter Status								
	Not Set for Election, Non-Zero Counters	165	32.7%	175	100%	**	501	90.1%	**
	Not Set for Election, Zero Counters	53	10.5%	0	0.0%		50	8.8%	
*	Set for Election, Zero Counters	232	46.1%	0	0.0%		6	1.1%	
x	Set for Election, Non-Zero Counters	1	0.2%	0	0.0%		0	0.0%	
	Election Closed, Non-Zero Counters	42	8.3%	0	0.0%		0	0.0%	
	Election Closed, Zero Counters	0	0.0%	0	0.0%		0	0.0%	
	Results Print Aborted, Non-Zero Counters	11	2.2%	0	0.0%		0	0.0%	
	Totals:	504	100%	175	100%		555	100%	

TABLE III
POST-ELECTION MEMORY CARD AUDIT ANALYSIS SUMMARY FOR NOVEMBER 2007, FEBRUARY 2008, AUGUST 2008, NOVEMBER 2008 CONNECTICUT ELECTIONS: (A) CARD FORMAT FOR ALL CARDS, (B) CARD STATUS FOR WELL-FORMATTED CARDS, (C) COUNTER STATUS FOR USABLE CARDS

		Post-Election Audit November 2007		Post-Election Audit February 2008		Post-Election Audit August 2008		Post-Election Audit November 2008	
		Number of Cards	% Total Cards	Number of Cards	% Total Cards	Number of Cards	% Total Cards	Number of Cards	% Total Cards
	(a) Card Format (all cards)								
*	Good Data, Clean Card	92	92.0%	197	93.8%	231	82.5%	418	90.5%
	Good Data, Some "Specks"	0	0.0%	3	1.4%	6	2.1%	3	0.6%
	Unusable Cards, "Junk Data"	8	8.0%	10	4.8%	43	15.4%	41	8.9%
	Total:	100	100%	210	100%	280	100%	462	100%
	(b) Card Status (well-formatted cards)								
	Not Programmed (Blank)	1	1.1%	0	0.0%	0	0.0%	0	0.0%
	Not Set for Election	11	12.0%	19	9.5%	1	0.4%	52	12.4 %
*	Set for Election	44	47.8%	44	22.0%	83	35.0%	90	21.4%
	Results Print Aborted	4	4.3%	10	5.0%	9	3.8%	20	4.7%
*	Election Closed	32	34.8%	127	63.5%	144	60.8%	259	61.5%
	Results Sent/Uploaded	0	0.0%	0	0.0%	0	0.0%	0	0.0%
	Audit Report Printed	0	0.0%	0	0.0%	0	0.0%	0	0.0%
	Totals:	92	100%	200	100%	237	100%	421	100%
	(c) Counter Status (usable cards)								
	Not Set for Election, Non-Zero Counters	11	12.1%	18	9%	1	0.4%	41	9.8%
	Not Set for Election, Zero Counters	0	0.0%	1	0.5%	0	0.0%	11	2.6%
*	Set for Election, Zero Counters	43	47.3%	44	22.0%	83	35.0%	88	20.9%
x	Set for Election, Non-Zero Counters	1	1.1%	0	0.0%	0	0.0%	2	0.5%
*	Election Closed, Non-Zero Counters	32	35.2%	126	63.0%	141	59.5%	259	61.5%
	Election Closed, Zero Counters	0	0.0%	1	0.5%	3	1.3%	0	0.0%
	Results Print Aborted, Non-Zero Counters	4	4.4%	10	5.0%	9	3.8%	20	4.7%
	Totals:	91	100%	200	100%	237	100%	421	100%

We observe a clear trend of increasing incidence of "junk" cards from election to election. The very existence of "junk" cards suggest either poor testing procedures at the vendor site or post-programming card "decay," perhaps due to battery issues (the cards are battery powered).

(b,c) Card & Counter Status. All relevant memory card

states and counters are presented in Table II. The anticipated memory card state depends on the audit type (pre-election or post-election) and whether the card was actually used during the election (for post-election cards). In no case, however, do we expect to see a card in a "Not Programmed (Blank)" state, or in a "Results Print Aborted" state, especially if the card was used during the election.

In pre-election memory card audits we encountered a blank card only once. However, the existence of such a card implies that not all cards are tested by the vendor that programs the cards before they are shipped.

According to the instructions set up by the Office of Secretary of the State, after receiving programmed memory cards poll workers of each district must place the cards in the available machines and run a test election on each of them. Once tested the cards should be placed in "election mode." Putting the cards in "election mode" at this point resets the counters to zero.

The audit results for November 2007 Election identified that over 50% of the cards were not in the expected "set for election" with zero counters state. This observation indicates that proper pre-election testing procedures are either not uniform, or are not communicated effectively. We note that for the August 2008 and November 2008 elections, very few of the cards was "Set for Election." However in this case, this is due to the fact that the pre-election memory cards were received directly from the external vendor programming the cards, consequently these cards were not subject to pre-election testing by poll workers.

Finally, we note that one card was found to be in the "set for election" state with non-zero counters (specifically recording that 19 vote were cast). This was determined to be due to incorrect pre-election testing procedures. If such a card was to be used in the election, the required printing of the zero-counter reports would have revealed this situation, and the poll workers would have reset the card to zero the counters. Nonetheless, if poll workers are unaware of this requirement, then such a card could result in incorrect election results. (It is worth noting that for the district in question, the Secretary of the State subsequently received copies of the printout that contained the required zero-count report, indicating that correct procedures were in fact followed on the election day.)

C. Post-Election Memory Card Audit Results

Table III combines the results of post-election memory card audits conducted for November 2007 elections, February 2008 primary, and August 2008 primary, and November 2008 Connecticut Elections. We make the following observations.

(a) Card Format. November 2007 Election: 100 memory cards were audited. 92% of these cards were properly formatted and contained good data, while 8 cards, or roughly one out of twelve cards, contained "junk" data. *February 2008 Election:* 210 memory cards were audited out of which 10 cards were identified as "junk", roughly one out of twenty one cards. *August 2008 Election:* 280 memory cards were audited out of which 43 cards were identified as "junk," roughly one out of seven cards. *November 2008 Election:* 462 memory cards were audited out of which 41 cards were identified as "junk," roughly one out of eleven cards.

We note that the percentage of "junk" cards detected through post-election audit is high, ranging from almost 5% to over 15%, although we do not observe a clear pattern similar to the one observed in the pre-election audit. Nevertheless, the percentages are overall higher than observed in the pre-election audit. This lack of consistency is potentially explained by the fact that the cards examined in the post-election audit are not sampled randomly, instead they represent a mixture of the cards actually used in an election and "leftover" cards that were not used (each district receives four cards out of which one ends up being used in the election).

Additionally, our event log analysis reveals that up to 8% of the cards were duplicated. The electoral procedures explicitly do not allow card duplication. This exhibits another deviation from the intended procedures. It is possible that some cards were determined to be "junk" in the pre-election testing process and were "reprogrammed" using card duplication procedure of AV-OS. Although all duplicates contained proper data, it is nevertheless a source of concern and the Connecticut SOTS Office will amplify the no-duplication policy for future elections.

(b,c) Card & Counter Status. Table III also shows that during each election 3% to 5% of the memory cards were found in a "Results Print Aborted, Non-Zero Counters" state. This is not the expected state and it suggests a deviation from standard procedures and possibly an incorrect shut-down of the machines after the completion of the election.

The post-election audit also allows one to observe the pre-election status of cards that were submitted for the audit, but were not used in the election. Recall that the expected state of such cards is "Set for Election, Zero Counters". Table III indicates the following: *November 2007 Election:* 54 cards were not used in election and were properly formatted. Out of 54 cards 11 (20.37%) cards were in a "Not Set for Election" status. Hence, about 80% of the cards were in the expected ("Set for Election") state. *February 2008 Election:* 63 cards were not used in election and were properly formatted. Out of 63 cards 19 (30%) cards were in a "Not Set for Election" status. Hence, 70% of the cards were in the expected ("Set for Election") state. *August 2008 Election:* 84 cards were not used in election and were properly formatted. Out of 84 cards 1 (about 1%) card was in a "Not Set for Election" status. Hence, almost 99% of the cards were in the expected ("Set for Election") state. *November 2008 Election:* 140 cards were not used in election and were properly formatted. Out of 140 cards 52 (37%) cards were in a "Not Set for Election" status. Hence, 63% of the cards were in an expected ("Set for Election") state.

Finally we note that three cards were found to be in "Set for Election" state with non-zeroed counters. As mentioned in Section VI-B such cards, if proper procedures are not followed, can lead to incorrect election results. A follow up investigation by the SOTS Office determined that these cases were due to detected malfunctions; these cards were removed from the election process and the ballots were recounted using backup machines.

D. Analysis of the Hand Counted Audit Returns

Recall that after each election the State of Connecticut performs hand counted audits on a random sample consisting of 10% of the districts. The analysis of hand count audit returns involves the participation of the Connecticut Secretary of the State Office that performs follow-up investigation in all cases were non-trivial discrepancies are reported between the machine counted totals and hand counted totals. Noteworthy is that in no cases the discrepancies could be attributed to incorrect tabulation by AV-OS, and that in all cases where follow up investigations were performed, it was determined that the discrepancies were due to human error in either totaling the votes or (mis)attributing votes to candidates. Thus, in order to increase the value of the hand counted audits it is extremely important to improve the precision of the hand counting process. Here we present a summary of the statistical analysis performed on the audit returns.

Table IV combines the results of this analysis for the following elections: November 2007 Election, February 2008 Primary, August 2008 Primary, and November 2008 Election.

The results indicate that in the substantial majority of cases there is either no discrepancy or the discrepancy between the machine totals and hand count totals below three. The highest discrepancy is a single case of 10 (ten). Of course, such discrepancies do not immediately imply miscounts on the part of the machine: in these cases there typically not a small number of ambiguous ballots are involved. In fact, over all audits, it is reported that while the average discrepancy per race is about one vote, the number of ambiguous or questionable ballots is about five.

A much more detailed presentation of the audit results briefly summarized here is found on our web site at URL http://voter.engr.uconn.edu/voter/.

VII. Conclusion

In this article we described a family of auditing procedures designed to enhance the integrity of elections conducted using optical scan technology. We focus specifically on auditing the "electronic fingerprint" of an election and motivate our selection of auditing procedures by modeling both the relevant computational infrastructure and a wide class of adversarial behavior. With these models in hand, we explored how various auditing choices can frustrate both the adversarial and non-malicious disruptive interference with the conduct if an election, and to provide essential sanity checks, increasing confidence in the election outcomes. We augmented this general discussion with a detailed survey of auditing carried out in the State of Connecticut in recent years. In addition to helping ensure safe use of technology in elections, these audits also help monitor adherence to the established policies and procedures in each election. We believe that our approach is practical, and we are continuing to refine and enrich the auditing procedures that are now routinely used in Connecticut.

Acknowledgments. We thank the anonymous referees for a number of insightful comments and suggestions that helped us substantially improve the quality of the presentation.

References

[1] Vote-switching software provided by vendors a partial list reported in the news. http://www.votersunite.org/info/Vote-Switchinginthenews.pdf.

[2] Project EVEREST: Risk assesment study of ohio voting systems, December 14 2007.

[3] Bannet, J., Price, D. W., Rudys, A., Singer, J., and Wallach, D. S. Hack-a-vote: Security issues with electronic voting systems. *IEEE Security & Privacy 2*, 1 (2004), 32–37.

[4] Secretary of State Debra Bowen moves to strengthen voter confidence in election security following top-to-bottom review of voting systems. http://www.sos.ca.gov/elections/voting_systems/ttbr/db07_042_ttbr_system_decisions_release.pdf, 2007.

[5] Cell, R. I. S. Trusted agent report Diebold AccuVote-TS voting system, January 2004.

[6] Chaum, D., Carback, R., Clark, J., Essex, A., Popoveniuc, S., Rivest, R. L., Ryan, P. Y. A., Shen, E., and Sherman, A. T. Scantegrity ii: End-to-end verifiability for optical scan election systems using invisible ink confirmation codes., 2008.

[7] Davtyan, S., Kentros, S., Kiayias, A., Michel, L., Nicolaou, N. C., Russell, A., See, A., Shashidhar, N., and Shvartsman, A. A. Pre-election testing and post-election audit of optical scan voting terminal memory cards. In *Proceedings of the 2008 USENIX/ACCURATE Electronic Voting Workshop (EVT 08), July 28-29, 2008, San Jose, CA, USA* (2008).

[8] Davtyan, S., Kentros, S., Kiayias, A., Michel, L., Nicolaou, N. C., Russell, A., See, A., Shashidhar, N., and Shvartsman, A. A. Taking total control of voting systems: Firmware manipulations on an optical scan voting terminal. In *Proceedings of the 24th Annual ACM Symposium on Applied Computing (SAC 09)* (2009).

[9] Feldman, A. J., Halderman, J. A., and Felten, E. W. Security analysis of the Diebold AccuVote-TS voting machine. http://itpolicy.princeton.edu/voting, 13 September 2006.

[10] Help America Vote Act. http://www.fec.gov/hava/law_ext.txt.

[11] Hursti, H. Critical security issues with Diebold optical scan design. http://www.blackboxvoting.org/BBVreport.pdf, 4 July 2005.

[12] Hursti, H. Diebold TSx evaluation. Black Box Voting Project, http://www.blackboxvoting.org/BBVtsxstudy.pdf, 11 May 2006.

[13] Jones, D. W. Observations and recommendations on pre-election testing in miami-dade county. http://www.cs.uiowa.edu/~jones/voting/miamitest.pdf.

[14] Jones, D. W. Regarding the optical mark-sense vote tabulators in maricopa county. http://www.cs.uiowa.edu/~jones/voting/ArizonaDist20.pdf.

[15] Jones, D. W. Auditing elections. *Commun. ACM 47*, 10 (2004), 46–50.

[16] Kiayias, A., Michel, L., Russell, A., Shashidar, N., See, A., and Shvartsman, A. An authentication and ballot layout attack against an optical scan voting terminal. In *Proceedings of the USENIX/ACCURATE Electronic Voting Technology Workshop (EVT 07)* (August 2007).

[17] Kiayias, A., Michel, L., Russell, A., Shashidhar, N., See, A., Shvartsman, A. A., and Davtyan, S. Tampering with special purpose trusted computing devices: A case study in optical scan e-voting. In *Proceedings of the 23rd Annual Computer Security Applications Conference (ACSAC 2007), December 10-14, 2007, Miami Beach, Florida, USA* (2007), pp. 30–39.

[18] Kohno, T., Stubblefield, A., Rubin, A. D., and Wallach, D. S. Analysis of an electronic voting system. In *Proceedings of the IEEE Symposium on Security and Privacy* (2004), pp. 27–.

[19] Connecting Maryland's election debacle dots. http://blackboxvoting.com/s9/index.php?/archives/138-CONNECTING-MARYLANDS-ELECTION-DEBACLE-DOTS.html, 2006.

[20] on Voting System Security, B. C. T. F. The machinery of democracy: Protecting elections in an electronic world. Brennan Center for Justice, NYU School of Law, http://www.brennancenter.org, 2005.

[21] Pamela Smith, Verified Voting.org. Written testimony before the Committee on House Administration, Subcommittee on Elections U.S. House of Representatives. http://electionaudits.org/files/PamelaSmithTestimonyFinal_2007mar20.pdf, 20 March 2007.

[22] Shamos, M. I. Paper v. electronic records an assessment. http://euro.ecom.cmu.edu/people/faculty/mshamos/paper.htm, 2004.

[23] Verified Voting. http:/www.verifiedvoting.org.

[24] Vora, P. L., Adida, B., Bucholz, R., Chaum, D., Dill, D. L., Jefferson, D., Jones, D. W., Lattin, W., Rubin, A. D., Shamos, M. I., and Yung, M. Evaluation of voting systems. *Commun. ACM 47*, 11 (2004), 144.

TABLE IV
HAND COUNT AUDIT ANALYSIS SUMMARY FOR NOVEMBER 2007, FEBRUARY 2008, AUGUST 2008, NOVEMBER 2008 CONNECTICUT ELECTIONS: $|D|$ REPRESENTS THE ABSOLUTE VALUE OF DISCREPANCY

	Hand Count Audit November 2007		Hand Count Audit February 2008, RP		Hand Count Audit February 2008, DP		Hand Count Audit August 2008		Hand Count Audit November 2008	
	No. of Records	**% Total Records**	**No. of Records**	**% Total Records**	**No. of Records**	**% Total Records**	**No. of Records**	**% Total Records**	**No. of Records**	**% Total Records**
Discrepancy										
$\|D\|$ of 0	319	42.99%	610	97.76%	611	97.92%	171	91.94%	424	51.39%
$\|D\|$ of 1-3	337	45.42%	13	2.08%	12	1.92%	13	6.98%	303	36.73%
$\|D\|$ of 4-6	66	8.89%	1	0.16%	1	0.16%	0	0.0%	64	7.76%
$\|D\|$ of 7-9	20	2.7%	0	0.0%	0	0.0%	1	0.54%	34	4.12%
$\|D\|$ of 10	0	0.0%	0	0.0%	0	0.0%	1	0.54%	0	0.0%
Totals:	742	100%	624	100%	624	100%	186	100%	825	100%

[25] WAGNER, D., JEFFERSON, D., AND BISHOP, M. Security analysis of the Diebold AccuBasic interpreter. Voting Systems Technology Assessment Advisory Board, University of California, Berkeley, 14 February 2006.