Final Report for the Project

Pre-Election and Post-Election Audit Solutions

for Optical Scan Voting Equipment

Connecticut Secretary of the State Office

and

Center for Voting Technology Research at the University of Connecticut

February 2, 2014

1. Project Landscape

The State of Connecticut, Office of the Secretary of the State is developing and documenting processes and best practices for pre-election testing and post-election audits. Through our on-going partnership with the University of Connecticut VoTeR Center we have identified a need to conduct thorough pre-election logic and accuracy testing as well as post-election audit processes.

The need for pre-election testing and post-election auditing has evolved in recent years because elections are counted using Accu-Vote optical scan electronic voting systems. In addition to ensuring that a problem does not exist with the hardware or software of the voting system that could ultimately produce result-changing errors, it has become increasingly important in Connecticut that efficient and reliable pre- and post-election procedures be developed so that the general public can be assured of an accurate and reliable result.

We have improved upon existing processes used in both the pre-election logic and accuracy testing as well as post-election auditing processes. We developed new tools that supplement the existing processes, and in particular, we developed a new process for analyzing election event logs. With new and improved pre-election logic and accuracy testing we hope to reduce the number of memory cards used in our optical scan voting machines that become unusable because of memory loss. With new and improved post-election auditing procedures we aim to increase the accuracy and reduce the amount of time and cost that our municipalities currently incur when they perform our current post-election audit of voting machines.

Another major component of this project is the development of a prototype for an *Audit Station (AS)*. The prototype will include a well-defined methodology and recommendations on its use and features. The development of AS is based on results of several years of research and direct participation in the Election Process in the State of Connecticut.

The vision for the Audit Station is a combination of hardware, specialized software, methodology and auditing procedures for automated hand count activities. The system is being built as a tool for auditors, addressing or identifying how to address requirements for tamper resistance of various degrees.

2. Progress Overview and Document Structure

This report covers the period from October 17, 2012 to January 31, 2014. For the detailed description of the project and program design we refer the reader to the full project proposal.

The purpose of this project is twofold. First, it is to research and develop logic and accuracy testing tools and methodologies to streamline and augment both the pre- and post-election audit processes. Second, it is to specify, develop and validate an *audit station* that enables *independent* counting and tallying of the election results. An *independent* audit station will assuage concerns about the validity of machine counting and significantly reduce the cost and burden that arise with error-prone human counting.

The project is progressed as planned. In particular, we have improved our memory card auditing and, as planned, we successfully collected data for the pre-election audit of memory cards for the November 2011 elections, April 2012 primaries, August 2012 primaries, November 2012 elections, and August 2013 primaries. We are currently preparing the reports documenting the audits for the November 2013 elections. We have also developed the first release of the Audit Station, and we successful used it in four pilots in 2013. In the next two sections we presents the progress in the major technical areas of the project:

Section 3, Logic and Accuracy Testing, deals with the research and advanced development aimed at logic & accuracy testing and the pre-election audits.

Section 4, Post-Election Auditing, deals with the research and advanced development aimed at post-election audits. Here we present our accomplishments in developing the Audit Station.

Section 5, Memory Card Reliability, deals with our evaluation of the new non-volatile memory cards that dramatically improves the reliability of the electoral processes as compared to the older, battery-powered volatile memory cards.

Section 6, Automating Event Log Analysis, discusses the new automated tool we developed for pre- and post- election analysis of election event logs collected from the memory cards of voting tabulators.

3. Logic and Accuracy Testing

The most critical part of the logic and accuracy (L&A) testing of the optical scan (OS) voting equipment is to ensure that the memory card is (1) programmed appropriately and (2) does not fail. The UConn VoTeR Center has worked closely with our local election officials to determine the most common failure points for OS equipment.

The auditing tools developed in this project are used in both pre-election and post-election auditing. The tools are designed to extract and analyze the content of the removable memory cards. In the pre-election testing, the tools insure that (a) the cards are correctly programmed, and (b) the OS terminals with the memory cards were properly prepared for elections. In the post-election testing, our tools check that (a) the cards are (still) correctly programmed, (b) the logs show terminal usage consistent with the proper procedures, and that (c) the internal counters match the district election results reported at the close of the election.

Memory Card Analysis Tools

The memory card used in Connecticut with the AccuVote OS terminals is a 40-pin 128KB Epson card. During the current reporting period we continued to improve the analysis tools and we used our tools in the Pre-Election Audit and Post-Election Audit of memory cards for the November 2011 elections, and April 2012 and August 2012 primary elections in Connecticut. We have published the audit reports for November 2011 elections,[1] April 2012 primary elections,[2] August 2012 primary elections[3,4], and November 2012 elections.[5,6] We are currently finalizing the report for the November 2013 elections. Once the analysis is complete, we will publish detailed reports (that will be made available on the UConn VoTeR Center web site and from Office of the Secretary of the State).

We now outline the specific analyses that are conducted within the pre-election audit.

Analysis of the Vote Totals (VTM) and Election Information Block (EIB). VTM contains the results of the election in a given district. EIB contains the election data and ballot layout. Our analysis tools extract VTM and EIB data and verify that EIB is correct by comparing it to the corresponding data in a trusted election database. In pre-election audit, the tool also checks that VTM contains zeroed counters. (In post-election audit, the tool will be used to verify that the counters in VTM match the official record.)

Analysis of Bytecode (Executable Code, EC). The memory cards contain executable code in the form of bytecode that is responsible for the reporting procedures associated with an election. This code needs to be analyzed to ensure the absence of undesirable behaviors. In particular, the code must accurately report the election results and be devoid of malicious or erroneous computation. Current audits involve manual analysis of the official version of this code, and then our tools check that each card has the correct version. In this project we will develop computer-aided tools to automate the analysis of the bytecode to substantially speed up its validation. We expect to receive the new code from the vendor in January 2012, at which time the computer-aided tools will be explored.

Analysis of Event Log (EL). The event log contains a record of the actions on the machine. The tools under development include a module that automates the EL analysis. We have developed a new tool that that allows one to defining an improved computational model embraced by the electoral process (as defined by the official elections documentation). The tool automates the analysis of EL, and performs a more comprehensive analysis and produces substantially better user notification than the tool we used before. We discuss this tool later in this report.

---

[1] Technological Audit of Memory Cards for the November 8, 2011 Connecticut Elections, VoTeR Center, April 5, 2012 (http://voter.engr.uconn.edu/voter/wp-content/uploads/VC-audit-main-11.pdf)

[2] Technological Audit of Memory Cards for the April 24, 2012 Connecticut Primary Elections, VoTeR Center, August 29, 2012 (http://voter.engr.uconn.edu/voter/wp-content/uploads/VC-audit-main-12-04.pdf)

[3] Pre-Election Audit of Memory Cards for the August 14, 2012 Connecticut Primary Elections, VoTeR Center, October 16, 2012, Version 1.0, (http://voter.engr.uconn.edu/voter/wp-content/uploads/VC-pre-2012-Aug.pdf)

[4] Post-Election Audit of Memory Cards for the August 14, 2012 Connecticut Primary Elections, VoTeR Center, October 16, 2012, Version 1.0 (http://voter.engr.uconn.edu/voter/wp-content/uploads/VC-post-2012-Aug.pdf)

[5] Pre-Election Audit of Memory Cards for the November 6, 2012 Connecticut Elections, VoTeR Center, January 18, 2013, Version 1.0 (http://voter.engr.uconn.edu/voter/wp-content/uploads/VC-pre-audit-Nov-2012.pdf)

[6] Post-Election Audit of Memory Cards for the November 6, 2012 Connecticut Elections, VoTeR Center, April 4, 2013, Preliminary Version 0.5 (http://voter.engr.uconn.edu/voter/wp-content/uploads/VC-audit-main.pdf)

Analysis of Vulnerabilities. We are performing on-going analysis of memory card vulnerabilities. During this period we identified a new vulnerability.[7] We identified a new attack that can be delivered without opening the system enclosure, and without changing a single bit of the system's firmware. The attack is launched by inserting a maliciously programmed AV-OS memory card into the terminal. The card contains binary code that exploits careless runtime memory management in the system's firmware to transfer control to alternate routines stored in the memory card. Once the control is taken by the injected code, the voting system is forced to operate according to the wishes of the attacker. In particular, given that the attack results in the execution of the arbitrary code, an attacker can completely take over AV-OS operation and compromise the results of an election. It is also noteworthy that once a memory card is compromised it can be duplicated using the native function of the voting terminal. In some past elections it was observed that up to 6% of all memory cards were involved in card duplication. There exists a non-trivial possibility that the infection on one memory card can propagate virally to other cards in a given election. This development was performed without access to the source code of the AV-OS system and without access to any internal vendor documentation. We note that this work is performed solely with the purpose of security analysis of AV-OS.

4. Post-Election Auditing

The most critical part of any post-election auditing procedure is to ensure that the results reported by the voting equipment are accurate. To this end, the State of Connecticut enacted into law a comprehensive audit procedure for the new voting equipment. This procedure requires local election officials to hand count all ballots that were cast on audited OS machines and compare that count to the count provided by the machines on election night. If a discrepancy is found, additional audits and possibly a recount may be required. Working closely with our local election officials throughout the implementation of this new process we have found that the weakest point in this process is the hand counting itself. To date, in all cases when noteworthy discrepancies were observed between the machine and hand counts, follow up investigations identified numerous hand counting errors.

Post-election audits seek to confirm that the basic code of the memory card properly read and recorded the vote totals for any proper ballot cast; it also ensures that no erroneous or malicious code was installed on the memory card that would invalidate the counting. We covered the technological post-election audits in the previous section. The statistical analysis of the hand-counted audit returns is available in the published reports.[8,9] The analysis report for the November 2012 elections will be published as soon as it is available.

Automating Post-Election Audit of Cast Ballots

To address the challenge of hand counted audits, we have developing a prototype of an *Audit Station (AS)*, including a methodology and recommendations on its use. This work and the refinement of the auditing procedures address the weak point of the hand count–its lack of precision and high costs of

[7] R.J. Jancewicz, A. Kiayias, L. D. Michel, A. C. Russell, and A. A. Shvartsman. Malicious takeover of Voting Systems: Arbitrary Code Execution on Optical Scan Voting Terminals. In Proc. of 28th ACM Symposium On Applied Computing, SAC 2013, pp. 1816-1823, Coimbra, Portugal, March 18-22, 2013

[8] Statistical Analysis of the Post-Election Audit Data 2010 August Primary Election, VoTeR Center, October 27, 2010 (http://voter.engr.uconn.edu/voter/wp-content/uploads/2010-aug-hand-v11.pdf)

[9] Statistical Analysis of the Post-Election Audit Data 2011 November Election, VoTeR Center, June 7, 2012 (http://voter.engr.uconn.edu/voter/wp-content/uploads/Nov-8-2011-HCA-V11.pdf)

getting precise counts. The proposed AS will significantly cut the time required for the hand count, while providing a higher accuracy and efficiency.

The Audit Station was developed as a turnkey solution that consists of hardware (computer, keyboard, and scanner), software, methodology and auditing procedures for automated hand count activities. As a tool for auditors it will be deployed broadly in Connecticut in the post-election audits. In 2013 we have performed successful pilots of the Audit Station in several towns. Next we present the details of this development.

Our Approach to Automating Audits

Before attempting any approach to automating post-election audits, it is important to consider the question of whether hand count audits can or should be automated. Given the challenges and issues with hand count audits, it is tempting to develop a completely automated approach. However, there are serious concerns associated with the use of automation in post-election audits if the human auditors are prevented or excluded from being meaningfully involved in the audit procedure: *Quis custodiet ipsos custodes?*[10] For example, some proposals to automate audits permit the use of the same equipment to tabulate the ballots. This is clearly problematic: using the same tabulator, or even a different tabulator of the same design will not reveal problems that cause similar errors in interpretation, or even complete misinterpretation of ballots. Using equipment or software from the same vendor, or using equipment from a different vendor to perform completely automated retabulation of ballots is also problematic for similar reasons. In general, any opaque, unobservable, or unobserved automated auditing presents problems due to the fact that the only primary document in the election, that is, the voter-generated paper ballot, is never inspected by the audit officials. Automated audit systems that analyze ballot images and that separation ballot images from ballots are likewise troublesome. This is because "a subverted retabulation system could display arbitrarily many ballot images and correct interpretations thereof, yet every vote count could be misreported."[11] In general, any completely automated audit system where human auditors delegate all responsibility for the audit to automation cannot be recommended as a valid approach to retabulation. To sum it up, "relying on unaudited retabulations is dangerous and unwarranted."[12]

Nevertheless, given the cost, time, and accuracy concerns plaguing purely manual audits, it is desirable to provide some automation in assisting post-election audits. What kind of automation can be sensibly deployed? We consider it reasonable to provide the semi-automated means for assisting audits that are not subject to the same flaws as those found in the completely automated or unaudited approaches. (We separately document requirements for an acceptable audit process.) The main goal of the project is to specify, develop and validate a novel audit station that will enable fully independent counting and tallying of the election results. An independent Audit Station will assuage concerns about the validity of machine counting and significantly reduce the issues associated with error-prone human counting.

Audit Station at a High Level

Auditing involves automatic scanning of ballots in batches, where each ballot, and its suggested interpretation, is projected onto a large screen for auditors and the interested public to observe. Using auditor-specified definitions, the system identifies unambiguous and questionable votes and presents this information by means of color-coded overlays on the ballots. The auditors can accept the

[10] (Latin) *Who will watch the watchmen?*

[11] M. Lindeman, R.L. Rivest, and P.B. Stark. Retabulations, Machine-Assisted Audits, and Election Verification. 20 March 2013 (http://www.stat.berkeley.edu/~stark/Preprints/retabulation13.htm).

[12] Ibid.

automatic interpretation, or they can override it. Each batch can be scanned multiple times to increase the auditors confidence as needed.

The Audit Station, as presented here, was used in four pilot audits in the State of Connecticut. The system consists of the following main hardware components: 1) optical scanner, 2) computer, 3) printer, 4) and projector. In the audit mode, the system projects an image of the ballot together with its interpretation.

The setup of the Audit Station is shown in Figure 1. Each hardware component is a relatively inexpensive COTS component. In this paper we do not describe these components in detail and we do not delve into technical decisions that led to the selection of the computing platform. We mention that the current system runs on a mid-range Apple Mac mini and it can include any optical scanner that supports standard interfaces.

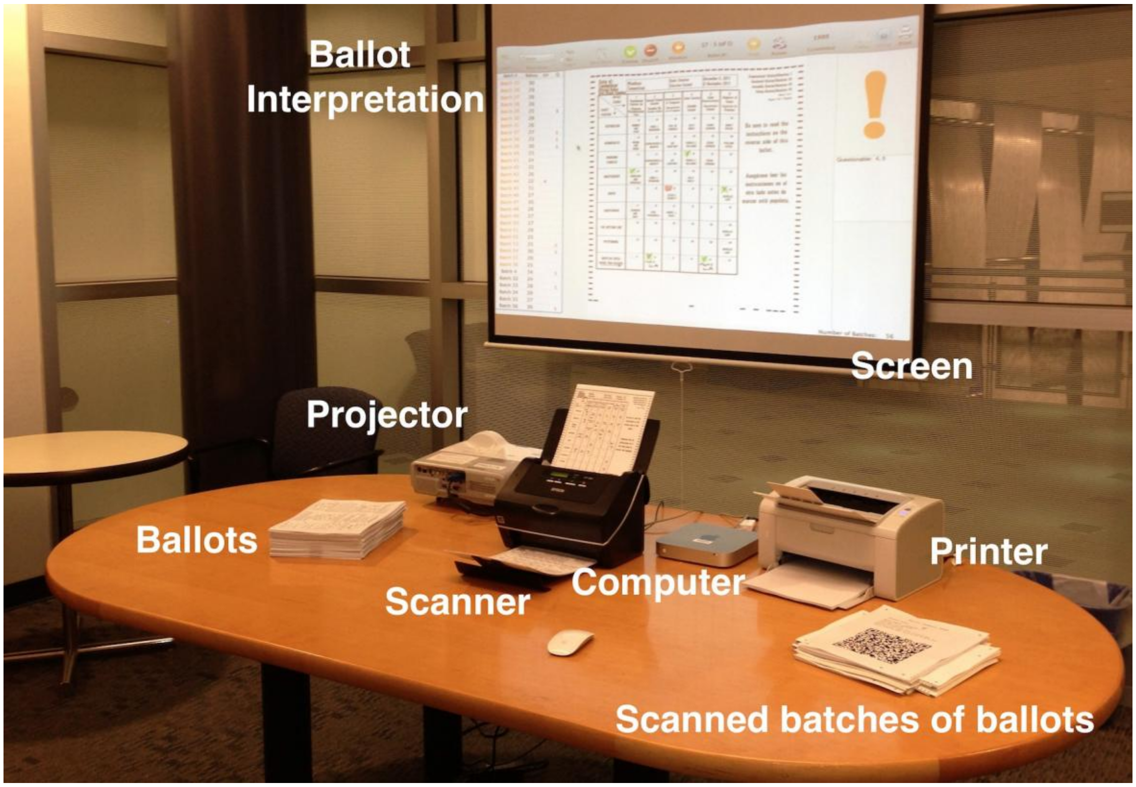

Figure 1. Audit Station Setup

The scanner in the figure is an inexpensive Epson GT-S80 model with which we achieve rates of up to 40 two-sided ballots per minute (this rate is currently only limited by the capability of the scanner). In supporting the batch-oriented audit process, it is most convenient to limit the size of the batch to the capacity of the automatic feeder in the scanner (the system with the shown scanner handles up to 40 ballots at a time).

Likewise, any standard computer projector can be used, and the screen is optional as the projector works quite nicely with any lightly colored wall. Note that no computer monitor is needed, since the projector also serves as the monitor. The setup shown in the figure is compact and it was easily transported for the audit pilots that we conducted in several towns.

The Audit Process

We now outline the audit process in greater detail. The computer-assisted audit process is designed to be used in conjunction with the post-election ballot audits in a polling district (precinct). The process assumes a batch-oriented approach, where all ballots in a district are divided into small batches, with each batch tabulated with the help of the audit station. The auditors can make decisions for any ballot on whether to accept the cast votes as analyzed by the Audit Station, or to revise the votes. The overall audit process is as follows.

1. Once a specific district is chosen for the post-election audit, the Audit Station is configured to audit the particular district using the official ballot definitions. (We do not present this function in detail. The system is designed to support an administrative interface that is used to prepare the system for the audit. This includes providing a ballot definition to the system and annotating it as necessary. Ultimately the information for the ballot definition will be obtained from four different sources: memory card of the optical tabulator, pdf file of the ballot, scanned ballot image, and the election management system database.)
2. On the day of the audit, the Audit Station is delivered to the district. The thresholds for determining what constitutes a “vote" and what constitutes a blank, unvoted bubble are initialized at the district.
3. The ballots are divided into batches. No pre-counting of the ballots is necessary---the auditors simply separate a deck of ballots based on its thickness to approximate the desired size.
4. Scanning a batch:
    a. A batch is scanned using the Audit Station using one of the two modes described above.
    b. If this is the first time the batch is processed, after the last ballot of the batch, the Audit Station generates a batch cover sheet that contains a unique batch sequence number. The cover sheet is human readable as well as encoded using a QR code.
    c. The batch cover sheet is placed on top of its batch. This is used to identify the batch if it is to be examined manually or rescanned using the Audit Station.
    d. If the Audit Station determines that some votes are ambiguous or cannot be processed, the auditors are informed. In any case, each scanned ballot can be compared with the corresponding paper ballot, and the auditors can revise the ballot interpretation and/or rescan the batch.
    e. Once the auditors accepted and commit the results for a batch, the Audit Station adds the totals from the batch to the election totals (replacing the previous interpretation of this batch, if this is not the first scan of the batch). The Audit Station displays the running election totals and the most recently scanned batch totals.
5. Any batch can be rescanned as many times as necessary. Each rescan of a batch produces a new result for the batch that overrides the results of any previous scan. For a rescan, ballots can be added to or removed from any batch that has not been committed by the auditors.
6. After all batches have been scanned, processed, and committed, the Audit Station produces the final tally based on the internally stored summaries. The Audit Station also exports the results for each ballot, each batch, and the overall totals for independent verification, depending on the official procedures (in particular, this enables the system itself to be audited in the style risk-limiting audits. Once this is done the results of the audit cannot be altered/revised using the Audit Station. (If changing the results is necessary, the entire district needs to be re-audited.)

We note that the decisions on the deployment logistics of the Audit Station and associated methodology have not been finalized, thus our presentation focuses on the selected capabilities of the system. We next provide additional details of the audit process.

Audit Station Details of the User Interaction

Ballot interpretation display. For each ballot in a batch the Audit Station displays the ballot interpretation as the ballots are being scanned. The auditors can page through the ballots in a batch to examine the interpretation, to compare the interpretation to the physical paper ballots in the batch, and to override the interpretation.

Figure 2 shows the automatic interpretation of the votes recorded on a ballot. The system uses a color-coded transparent overlay to show the automatically derived vote interpretations. The large exclamation sign indicates that the batch contains questionable votes on ballots 4 and 5, as shown below the sign. The currently displayed ballot is number 5 in the batch. The system identifies votes based on the thresholds set by the auditors; the interpretation is shown the color overlay. The marks that exceed the voted threshold are colored green by the system. The marks in the voting areas that are above the blank threshold but below the voted threshold are colored red.

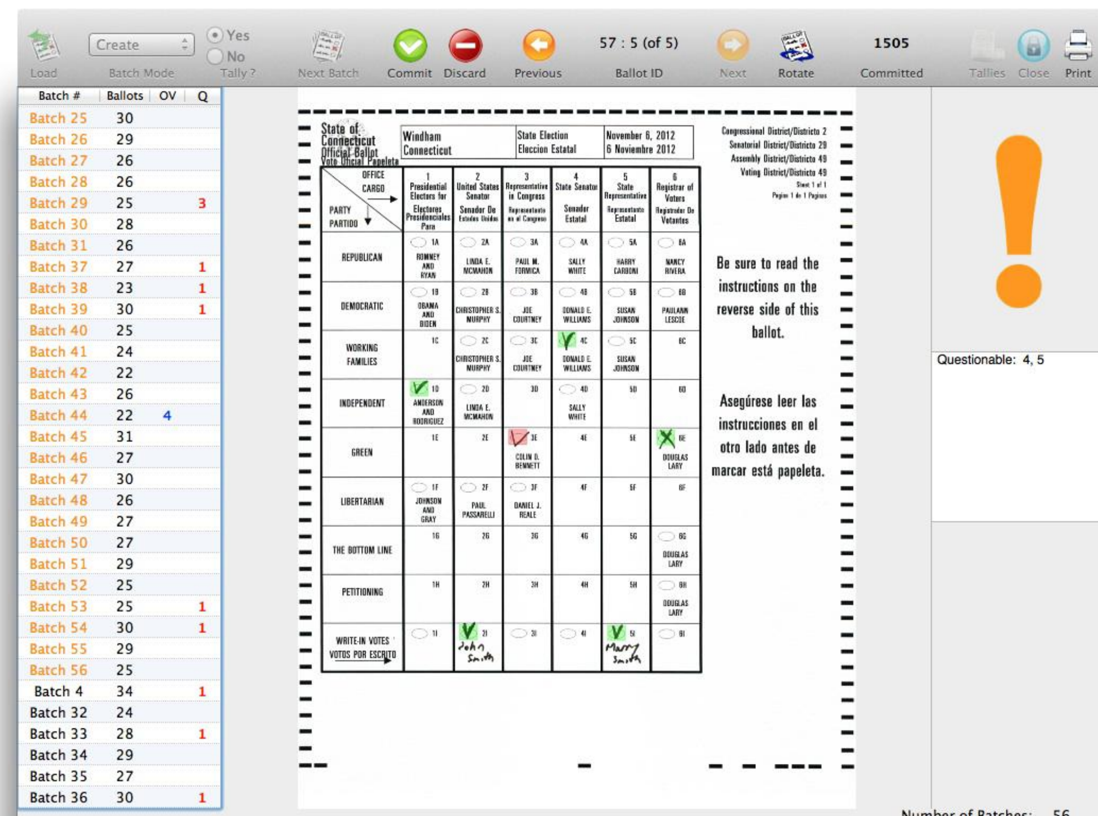

Figure 2. Audit Station ballot interpretation

Figure 3 shows the vote override interface. When the auditors examine a scanned ballot, they can override the interpretation of any bubble (vote area). If a vote interpretation is overridden, the system includes and displays an annotation in the ballot overlay.

The ballot number 5 in the current batch number 57 from Figure 2 contains a mark that the Audit Station considers ambiguous; it is highlighted red by the system. The auditor selected the corresponding bubble for inspection. The voting area is magnified and the system also displays the relevant thresholds and the score obtained by the mark. The auditors can now accept the automatic interpretation, or override it. The override can be done using any of the three choices: not voted, voted, or ambiguous.

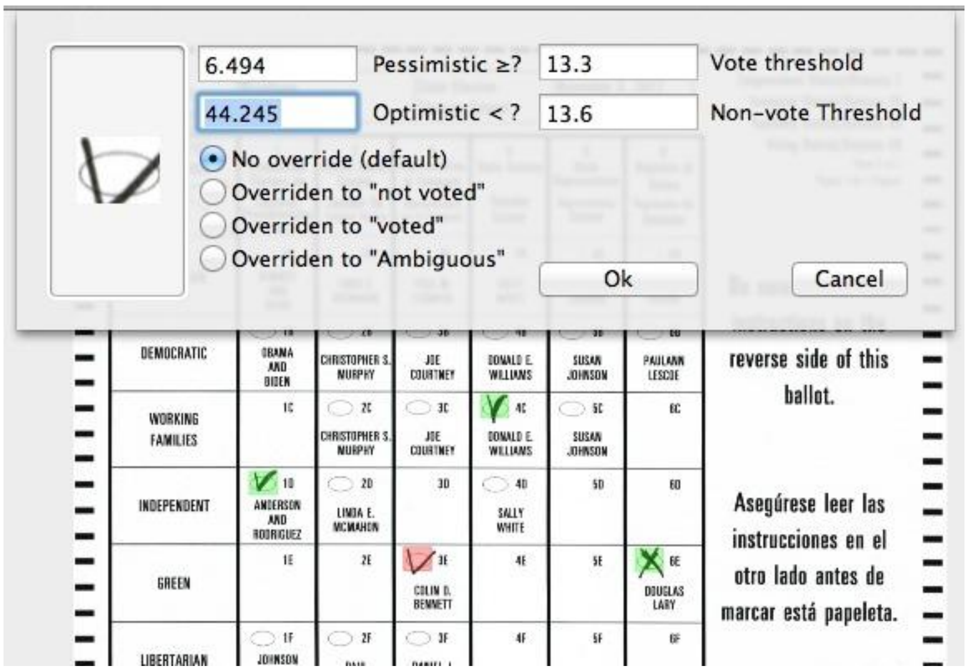

Figure 3: Examining and overriding automatic interpretation.

Batch processing details. We now describe the processing of a batch in more detail. A batch is any collection of ballots, where any ballot can belong to only one batch at a time. Setting a sensible maximum number of ballots in a batch should reflect the common opinion that “small-batch auditing" is desirable. We found it convenient to limit the maximum size of the batch not exceed the capacity of the scanner's automatic feeder (e.g., up to 40 ballots). Using larger size batches is of course possible, but this requires that a larger batch is fed through the scanner in smaller sub-batches suitable for the scanner.

The audit of each batch can be done in one of two modes: (1) a batch is scanned automatically, then the ballot interpretations are browsed by the auditors, so that each ballot interpretation is examined and revised by the auditors as needed, or (2) the ballots in the batch are scanned one at a time, with the audit station pausing after each ballot, to let the auditors observe the results of the scan for each ballot. The first process is faster, while the second process provides an easier way for comparing the results of the automatic interpretation to the physical ballots.

Once a batch is scanned and processed, the system displays the batch summary as shown in Figure 4. All races in the election are counted at the same time. Recall that any batch can be re-scanned if deemed necessary, and as many time as the auditors consider necessary. When a batch is processed for the first time, its identifier is given in a light (orange) font in the left pane of the display. When two consecutive scans of the same batch yield identical results, the color is changed to black. The auditors can always commit the batch results based on the most recent scan, overriding all prior interpretations.

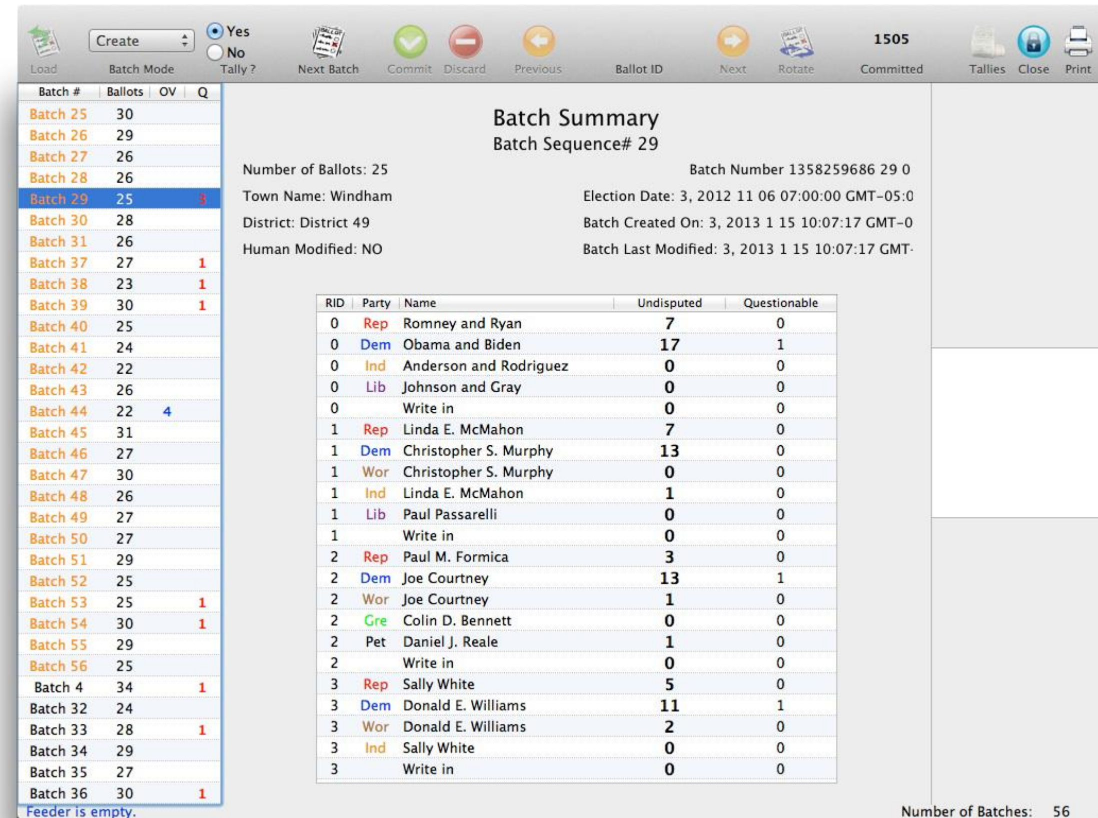

Figure 4: Batch summary display: the totals are given for the selected batch in the left column.

Finally we note that the system allows for multiple audit stations to be used in parallel in the same district level audit. The system automatically aggregates the results of the batch processing at different stations.

Audit summary. The Audit Station maintains an audit summary that provides a cumulative view of the batches scanned thus far. An example of the summary display is in Figure 5. Batch summaries are displayed in the left pane. The status of each batch is indicated by the font color (orange or black as described earlier) in the column labeled Batch #. The column labeled Ballots gives the number of ballots in each batch. The column labeled OV deals shows the number of overvoted ballots. The number of questionable bubbles is given per batch in the column labeled Q.

Audit summary display in Figure 5 shows in the left column the committed batches that are shown in the black font and the tentative batches are in a light font; the large "!" sign indicates that there are ambiguous votes in the election that may need to be resolved.

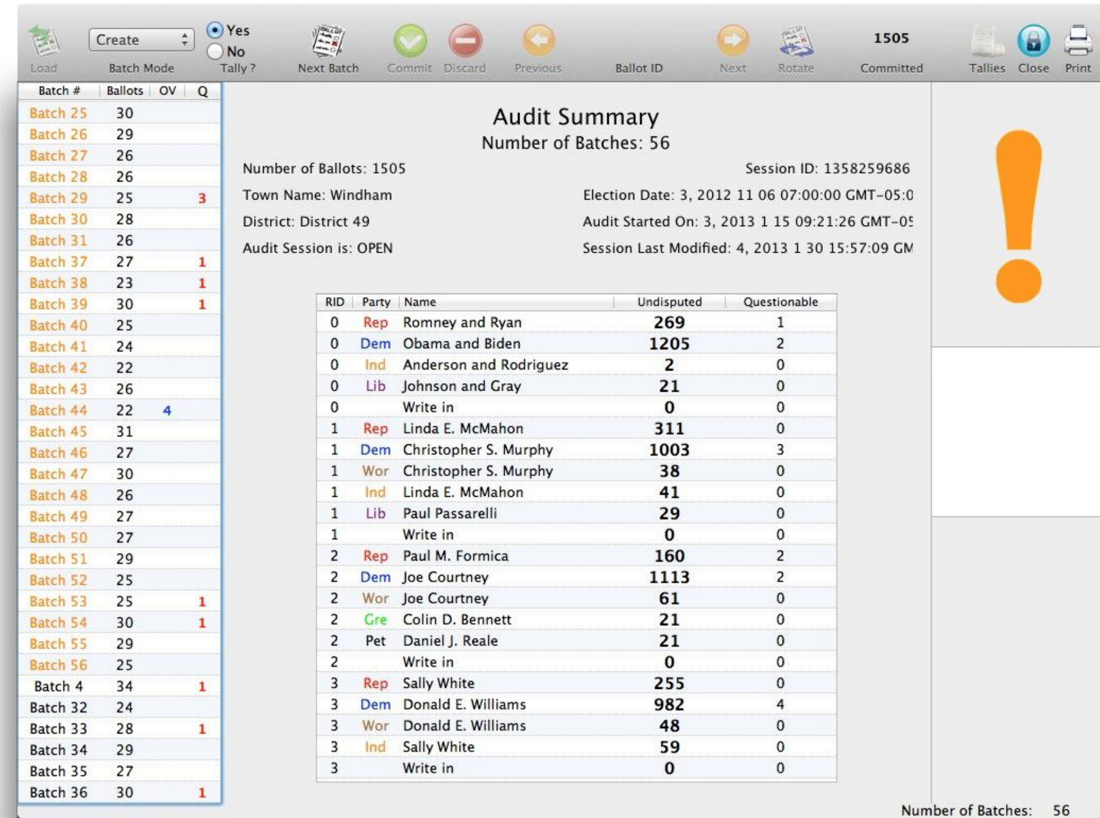

Figure 5: Audit summary display

Technical issues. Several technical issues were resolved during the development of the Audit Station. We have considered and evaluated three system platforms for this development (Windows, Unix, and Mac), and we have settled on the Mac platform as the most suitable. This determination was made on the basis of the platforms ability to support effective user interface development, its support for a variety of scanners, and its support of suitable image processing software.

Image Processing. We designed and evaluating algorithmic approaches to image correction and analysis. We have established that, depending on the type and make of scanners, the scanned images are distorted. This distortion is typically piecewise linear in the length of the scanned image (ballot). Our algorithms have been designed to correct the scanned image, including de-skewing, so that the digitized image is a faithful representation of the correctly printed ballot. We also developed algorithms that enable fast ballot analysis, with the goal of enabling scanning rates up to the ability of the specific scanner. Using the current scanner, we are able to push it to its limit of about 40 ballots per minute. We are planning to integrate commercial scanners that will enable up to 100 ballots per minute scanning rates, but the availability of such (affordable) scanners is an issue, and the vendor software support for such scanners is not sufficient for fast adoption.

Scanner Certification. Given that the Audit Station is designed as a turnkey solution, we intend to certify scanners for integration. In order to ensure the most faithful ballot image capture, we are developing techniques for evaluating scanners, and specifying scanner “signatures" to be included with our software. Given that different scanners have different scanning characteristics, only the scanners for which we develop signatures may be include in the Audit Station solution. Such a diligent approach is necessary to prevent the possible interference of scanner hardware variations from affecting the quality of captured ballot images.

Of independent interest, we note that our software accurizing of scanners allows for inexpensive commercial scanner to be used with high precision, comparable to that achievable in expensive scanners. Additionally, our efficient image processing algorithms enable fast processing of ballots using a modest overall off-the-shelf computer system.

A forthcoming report will cover the technical details of the Audit Station implementation.

Audit Station Pilots in Connecticut

The Audit Station was deployed in auditing pilots in four municipalities in Connecticut using the actual ballots from the November 2012 elections. In each case ballots from one district were audited, where from about 2,000 to 3,800 ballots were processed depending on the district. The total ballot counts matched the official counts, except for one case that was apparently due to a single misplaced ballot. The summary of the audits is in Table 1.

Town	**Number of Ballots**	**Hand Count Total Hours**	**Audit Station Total Hours**	**% Savings**
Tolland	3851	48	14	70%
Bloomfield	2272	40	7	80%
Windham	1963	n/a	5	n/a
Vernon	2544	79	7	90%

Table 1: Summary of the number of ballots audited and the number of hours spend on audits.

In the towns of Tolland, Bloomfield, and Vernon official hand counts were performed prior to the Audit Station pilot. In each case we recorded the total number of hours spent doing the pilot audits. Given that the total number of hours spent in the official audit is available for the three municipalities, we observe that at least 70% savings in time were achieved using the Audit Station.

During each audit, we compared the official tally (and the official election-day hand count data where applicable) to the results obtained using the Audit Station. Overall, the system performed extremely well in terms of its precision. The semi-automated audit tallies were within one vote of the official count; this is summarized in a separate report.[13]

5. Memory Card Reliability

In recent years, technological audits in the State of Connecticut established that up to 15% of all memory cards fail by losing all data. These older memory cards rely on a battery to maintain data and we have established that battery depletion is the major cause of card failures.[14] The low battery

[13] T. Antonyan, Th. Bromley, L.D. Michel, A.C. Russell, A.A. Shvartsman, S. Stark. Computer Assisted Post Election Audits (Extended Abstract). State Certification of Voting Systems National Conference. June 20-21. 2013, Harrisburg, Pennsylvania, USA (URL: http://voter.engr.uconn.edu/voter/wp-content/uploads/AS-2013.pdf).

[14] T. Antonyan, N. Nicolaou, A. Shvartsman, and T. Smith. Determining the Causes of AccuVote Optical Scan Voting Terminal Memory Card Failures. USENIX/ACCURATE Electronic Voting Workshop, 15 pages, electronic ed. (2010)

warning from the OS may be too lax making it difficult to predict and anticipate card failures. This high rate of failure is naturally unacceptable for digital components involved in such critical applications as elections. In more detail, our earlier investigation determined that the primary reason for memory card failures is depleted batteries. Once the battery's store of energy is depleted, the cards lose their data. The electrical properties of the batteries are such that the battery voltage output can decrease precipitously as the battery reaches the end of its service life. Therefore one cannot expect to rely on the low battery warning system built into the AV-OS optical scan tabulator. Battery depletion may happen within days after a card was programmed and tested. Thus even if a card is successfully programmed, it can fail before it is tested prior to an election, or at any time after it is successfully tested.

Following our findings and reports from other States on memory card failures, the vendor (Dominion Voting Systems), developed a new memory card design that is based on nonvolatile (MRAM) memory. These cards do not require a battery to store the data.

The State of Connecticut received a small number of new cards for evaluation. We have performed preliminary evaluation of these cards and determined that they are functionally compatible with the optical scan tabulators used in Connecticut. We have also confirmed that these cards are not subject to the high failure rates associated with the older cards. These preliminary findings were sufficiently encouraging, and a pilot deployment of the new cards was done in the April 2012 primary elections in the Town of Vernon. The technical analysis prior and after the pilot deployment of the new cards showed that the new cards performed well, no failures were detected, and no such cards lost their data.

We performed a more extensive evaluation of the new non-volatile memory cards.[15] We performed tests on the nonvolatile cards using the old, volatile cards as the control group. The purpose of this testing was to evaluate the integrity and reliability of the nonvolatile cards over time. Our evaluation concluded that the failure rates for new memory cards are negligible (fraction of one percent) compared to the failure rates of the old cards (up to 15%). We conjecture that the very few failures of new cards can be attributed to other reasons (e.g., tabulators themselves). The very small numbers of failures observed do not permit us to determine the causes with high certainty. Regardless of the causes, the number of failures is reduced by at least two orders of magnitude – a dramatic improvement in reliability. We have also evaluated how well the new cards withstand frequent repeated reading and writing cycles. Depending on the test, we subjected the cards from 25M to 30B cycles. All new memory cards tested passed this test.

Based on the results of the testing and on the pilot deployment of new cards, the State of Connecticut is planning a broader deployment of the new cards in the near future.

6. Automating Audit Log Analysis

We developed an entirely new system to perform the analysis of audit logs produced by the AV-OS tabulators during the electoral process.[16] We designed and implemented a systematic approach to automating the analysis of event logs recorded by the electronic voting tabulators in the course of an election. An attribute context-free grammar is used to specify the language of the event logs, and to distinguish compliant event logs (those that adhere to the defined proper conduct of an election) and

---

[15] A.A. Shvartsman et al., Summary Results of the Tests of Nonvolatile Memory Cards, VoTeR Center, Version 0.2, July 30, 2012.

[16] L.D. Michel, A.A. Shvartsman, N. Volgushev. A Systematic Approach to Analyzing Voting Terminal Event Logs. Accepted for publication in the *Journal of Election Technology and Systems*, 2014.

non-compliant logs (those that deviate from the expected sequence of events). The attributes provide additional means for semantic analysis of the event logs by enforcing constraints on the timing of events and repetitions of events. The system is implemented with the help of commodity tools for lexical analysis and parsing of the logs. The system was rigorously tested against several thousand event logs collected in real elections in the State of Connecticut. The approach based on an attribute grammar proved to be superior to a previous approach that used state machine specifications. The new system is substantially easier to refine and maintain due to the very intuitive top-down specification. An unexpected benefit is the discovery of revealing and previously unknown deficiencies and defects in the event log recording systems of a widely used optical scan tabulator.

Election audits are a critical procedural component of the electoral process to guarantee the proper conduct of an election. Our work demonstrates yet again how audits can be valuable in the forensic analysis of data collected from voting terminals used during the election. Indeed, the audit process reveals several classes of problems ranging from voting terminal malfunctions and defects to deviations in the recommended behaviors for system operators. Our contributions encompass a new formalization of voting machine event logs to systematize a multi-layered compliance analysis that delivers detailed notifications characterizing election traces. The event log analysis uses attributed context-free grammars, making the system highly extensible and maintainable, and readily available for refinements that reflect requirements for a correct conduct of an election. Additionally, our methodology led to the identification of previously unknown deficiencies and defects in the AV-OS logging system, further emphasizing the value of comprehensive audits.

We are currently preparing recommendations on implementing event logging systems for voting terminals that would enable even more comprehensive audit analyses. In our future work we will continue refining our approach and we intend to adapt the language definition for use in other jurisdictions using similar equipment based on their election protocols.

[End]