

Scantegrity II Municipal Election at Takoma Park: The First E2E Binding Governmental Election with Ballot Privacy

Richard Carback
UMBC CDL

David Chaum

Jeremy Clark
University of Waterloo

John Conway
UMBC CDL

Aleksander Essex
University of Waterloo

Paul S. Herrnson
UMCP CAPC

Travis Mayberry
UMBC CDL

Stefan Popoveniuc

Ronald L. Rivest
MIT CSAIL

Emily Shen
MIT CSAIL

Alan T. Sherman
UMBC CDL

Poorvi L. Vora
GW

Abstract

On November 3, 2009, voters in Takoma Park, Maryland, cast ballots for the mayor and city council members using the Scantegrity II voting system—the first time any *end-to-end* (E2E) voting system with ballot privacy has been used in a binding governmental election. This case study describes the various efforts that went into the election—including the improved design and implementation of the voting system, streamlined procedures, agreements with the city, and assessments of the experiences of voters and poll workers.

The election, with 1728 voters from six wards, involved paper ballots with invisible-ink confirmation codes, instant-runoff voting with write-ins, early and absentee (mail-in) voting, dual-language ballots, provisional ballots, privacy sleeves, any-which-way scanning with parallel conventional desktop scanners, end-to-end verifiability based on optional web-based voter verification of votes cast, a full hand recount, thresholded authorities, three independent outside auditors, fully-disclosed software, and exit surveys for voters and pollworkers.

Despite some glitches, the use of Scantegrity II was a success, demonstrating that E2E cryptographic voting systems can be effectively used and accepted by the general public.

1 Introduction

The November 2009 municipal election of the city of Takoma Park, Maryland marked the first time that anyone could verify that the votes were counted correctly in a secret ballot election for public office without having to be present for the entire proceedings. This article is a case study of the Takoma Park election, describing what was done—from the time the Scantegrity Voting System Team (SVST) was approached by the Takoma Park Board of Elections in February 2008, to the last cryptographic election audit in December 2009—and what was

learned. While the paper provides a simple summary of survey results, the focus of this paper is not usability but the engineering process of bringing a new cryptographic approach to solve a complex practical problem involving technology, procedures, and laws.

With the Scantegrity II voting system, voters mark optical scan paper ballots with pens, filling the oval for the candidates of their choice. These ballots are handled as traditional ballots, permitting all the usual automated and manual counting, accounting, and recounting. Additionally, the voting system provides a layer of integrity protection through its use of invisible-ink confirmation codes. When voters mark ballot ovals using a decoder pen, confirmation codes printed in invisible ink are revealed. Interested voters can note down these codes to check them later on the election website. The codes are generated randomly for each race and each ballot, and hence do not reveal the corresponding vote. A final tally can be computed from the codes and the system provides a public digital audit trail of the computation.

Election audits in Scantegrity II are not restricted to privileged individuals and can be performed by voters and other interested parties. Developers and election authorities are unable to significantly falsify an election outcome without an overwhelming probability of an audit failure [8]. The other side of the issue of integrity, also solved by the system, is that false claims of impropriety in the recording and tally of the votes are readily revealed to be false.¹

All the software used in the election—for ballot authoring, printing, scanning and tally—was published well in advance of the election as commented, buildable source code, which may be a first in its own right. Moreover, commercial off-the-shelf scanners were adapted to receive ballots in privacy sleeves from voters, making the

¹Note that a threat present and not commonly addressed in paper ballot systems is that additional marks could be added to ballots by those with special access. Such attacks are made more difficult by Scantegrity II.

overall system relatively inexpensive.

Despite several limitations of the implementation, we found that the amount of extra work needed by officials to use Scantegrity II while administering an election is acceptable given the promise of improved voter satisfaction and indisputability of the outcome. Indeed, discussions are ongoing with the Board of Elections of the city regarding continued use of the system in future elections.

Another observation from the election is that the election officials and voters surveyed seemed to appreciate the system. Since voters who do not wish to verify can simply proceed as usual, ignoring the codes revealed in the filled ovals, the system is least intrusive for these voters. Those voters who did check their codes, and even many who did not, seem to appreciate the opportunity.

This paper describes the entire process of adapting the Scantegrity II system to handle the Takoma Park election, including the agreement with the city, printing the special ballots with invisible-ink confirmation codes, actually running the election, and verifying that the election outcome was correct.

Organization of this case study The next section provides an overview of related work in this area, summarizing previous experiments with Scantegrity II and other E2E systems in practical settings.

Section 3 describes in more detail the setting for the election: giving details about Takoma Park and their election requirements. Section 4 gives more details of the Scantegrity II voting system, including a description of how one can “audit” an election. Section 5 provides an overview of the implementation of the voting system for the November 3, 2009 Takoma Park municipal election, including the scanner software, the cryptographic back-end, and the random-number generation routines.

Section 6 gives a chronological presentation and timeline of the steps taken to run the November election, including the outcome of the voter verification and the audits. It also gives the results of the election, with some performance and integrity metrics. Section 7 reports some results of the exit surveys taken of voters and pollworkers.

Section 8 discusses the high-level lessons learned from this election. Section 9 provides some conclusions, acknowledgements, and disclosures required by the program committee.

2 Related Work

Chaum was the first to propose the use of cryptography for the purpose of secure elections [5]. This was followed by almost two decades of work in improving security and privacy guarantees (for a nice survey, see

Adida [1]), most recently under the rubric of *end-to-end* voting systems. These voting system proposals provide integrity (any attempt to change the tally can be caught with very high probability by audits which are not restricted to privileged individuals) and ballot secrecy.

The first of these proposals include protocols by Chaum [6] and Neff [19], which were implemented soon after (Chaum’s as *Citizen-Verified Voting* [16] and Neff’s by VoteHere). Several more proposals with prototypes followed: *Prêt à Voter* [10], *Punchscan* [21, 15], the proposal of Kutyłowski and Zagórski [18] as *Voting Ducks*, and Simple Verifiable Voting [4] as *Helios* [2] and *Vote-Box* [24].

Making end-to-end systems usable in real elections has proven to be challenging. We are aware of the following previous binding elections held using similar verification technology: the *Punchscan* elections for the graduate students’ union of the University of Ottawa (2007) and the Computer Professionals for Social Responsibility (2007); the Rijnland Internet Election System (RIES) public elections in the Netherlands in 2004 and 2006; the *Helios* elections of the Rector of Université Catholique de Louvain [3] (2009) and the Princeton undergraduate student government election (2009), as well as a student election using *Prêt à Voter*.

Only the RIES system has been used in a governmental election; however, it is meant for remote (absentee) voting and, consequently, does not offer strong ballot secrecy guarantees. For this reason, it has been recommended that the RIES system not be used for regular public elections [17, 20]. *Helios* is also a remote voting system, and offers stronger ballot secrecy guarantees over RIES. The *Punchscan* elections were the closest to this study, but they did not rise to the level of public elections. They did not have multiple ballot styles, the users of the system were not a broad cross-segment of the population as in Takoma Park, the system implementors were deeply involved in administering the elections, and no active auditors were established to audit the elections. To date, this study is the most comparable use case of E2E technology to that of a typical optical scan election.

The case study reported here is based on a series of systems successively developed, tested, and deployed by a team of researchers included among the present authors originating with the *Punchscan* system. Although it used paper ballots, the *Punchscan* system did not allow manual recounts, a feature that the team recognized as needing to be designed into the next generation of systems. The result was *Scantegrity* [9], which retained hand-countable ballots, and was tested in a number of small elections. With *Scantegrity*, however, it was too easy to trigger an audit that would require scrutiny of the physical ballots. The *Scantegrity II* system [7, 8], de-

ployed in Takoma Park, was a further refinement to address this problem by allowing a public statistical test of whether voter complaints actually reflect a discrepancy or whether they are without basis. Note: in the rest of the paper, “Scantegrity” refers to the voting team or to the Scantegrity II voting system; which one is typically easily determined from context.

As part of the Scantegrity agreement with Takoma Park (see section 3), a “mock election” [26] was held in April 2009 to test and demonstrate feasibility of the Scantegrity system during Takoma Park’s annual Arbor day celebration. Volunteer voters voted for their favorite tree. A number of revisions and tweaks to the Scantegrity system were made as a result of the mock election, including: ballot revisions (no detachable chit, but instead a separate voter verification card), pen revisions (two-ended, with different sized tips), scanner station revisions (better voter flow, no monitor, two scanners), privacy sleeve (no lock, no clipboard, folding design, feeds directly into scanner), and confirmation codes (three decimal digits).

3 The Setting

For several reasons, the implementation of voting systems is a difficult task. Most voting system users—*i.e.* the voters—are untrained and elections happen infrequently. Voter privacy requirements preclude the usual sorts of feedback and auditing methods common in other applications, such as banking. Also, government regulations and pre-existing norms in the conduct of elections are difficult to change. These issues can pose significant challenges when deploying new voting systems, and it is therefore useful to understand the setting in which the election took place.

About Takoma Park The city of Takoma Park is located in Montgomery County, Maryland, shares a city line with Washington, D.C, and is governed by a mayor and a six-member City Council. The city has about 17,000 residents² and almost 11,000 registered voters [27, pg. 10]. A seven-member Board of Elections conducts local elections in collaboration with the City Clerk. In the past, the city has used hand counts and optical scan voting, as well as DREs for state elections.

The Montgomery County US Census Update Data of 2005 provides some demographic information about the city. Median household income in 2004 was \$48,675. The percentage of households with computers was 87.4%, and about 32% of Takoma Park residents above the age of twenty-five had a graduate, professional or doctoral degree. It is an ethnically diverse city: 45.8%

of its residents identify their race as “White,” 36.3% as “Black,” 9.7% as “Asian or Pacific Islander” and 8.2% as “Other” (individuals of Hispanic origin form the major component of this category). Further, 44.4% of its households have a foreign-born head of household or spouse, and 44.8% of residents above the age of five spoke a language other than English at home.

Instant Runoff Voting (IRV) Takoma Park has used IRV in municipal city elections since 2006. IRV is a ranked choice system where each voter assigns each candidate a rank according to her preferences. The rules³ used by Takoma Park (and the Scantegrity software) for counting IRV ballots are relatively standard, so we omit further discussion for lack of space.

Agreement with the City As with any municipal government in the US, Takoma Park is allowed to choose its own voting system for city elections. For county, state, and federal elections, it is constrained by county, state, and federal election laws.

Takoma Park and the SVST signed a Memorandum of Understanding (MOU), in which the SVST agreed to provide equipment, software, training assistance, and technical support. The City of Takoma Park agreed to provide election-related information on the municipality, election workers, consumable materials, and perform or provide all other election duties or materials not provided by us. No goods or funds were exchanged.

According to the MOU, if approved by the city council, the election was to be conducted in compliance with all applicable laws and policies of the city. This included using Instant Runoff Voting as defined by the City of Takoma Park Municipal Charter.

The SVST also agreed to pursue an accessible ballot-marking device for the election, but was later relieved of satisfying this requirement. Unfortunately, Scantegrity is not yet fitted with a voter interface for those with visual or motor disabilities, and accessible user interfaces were also not used in Takoma Park’s previous optical scan elections.

Timeline Scantegrity was approached by the Takoma Park Board of Elections in late February 2008, and, after considering other voting systems, the Board voted to recommend a contract with Scantegrity in June 2008. Following a public presentation to the City Council in July 2008, the MOU was signed in late November 2008, about nine months after the initial contact.

³For the exact laws used by Takoma Park, see page 22 of <http://www.takomaparkmd.gov/code/pdf/charter.pdf>. Section (f), concerning eliminating multiple candidates, was used in our implementation for tie-breaking only.

²See <http://www.takomaparkmd.gov/about.html>.

The SVST held an open workshop in February 2009 to discuss the use of Scantegrity in both the mock and real elections. This workshop was held at the Takoma Park Community Center and was attended by Board of Election members, the City Clerk, current members (and a retired member) from the Montgomery County Board of Elections, as well as a representative each from the Pew Trust and FairVote. Following the mock election in April 2009, the SVST proposed a redesigned system taking into consideration feedback from voters and poll workers (through surveys) and the Board of Elections. The Board voted to recommend use of the redesigned system in July 2009; this was made official in the city election ordinance in September 2009.⁴ Beginning around June 2009, a meeting with representatives of the SVST was on the agenda of most monthly Board of Election meetings. Additionally, SVST members met many times with the City Clerk and the Chair of the Board of Elections to plan for the election.

The final list of candidates was available approximately a month before the election, on October 2. The Scantegrity *meetings* initializing the data and ballots were held in October (see Section 6), as was a final workshop to test the system. Absentee ballots were sent out by the City Clerk in the middle of October. The SVST delivered ballots to the City Clerk in late October, and early voting began almost a week before the election, on October 28. Poll worker training sessions were held by the city on October 28 and 31, and polling on November 3, 2009, from 7 am to 8 pm. The final Scantegrity audits were completed on 17 December 2010; all auditors were of the opinion that the election outcomes were correct (for details see section 6).

4 Scantegrity Overview

In this section, we give an overview of the Scantegrity system. For more detailed descriptions, see [7, 8].

Voter Experience At a high level, the voter experience is as follows. First, a voter checks in at the polling place and receives a Scantegrity ballot (See Figure 2) with a privacy sleeve. The privacy sleeve is used to cover the ballot and keep private the contents of the ballot. Inside the voting booth, there is a special “decoder pen” and a stack of blank “voter verification cards.” The voter uses the decoder pen to mark the ballot. As on a conventional optical scan ballot, she fills in the bubble next to each of her selections. Marking a bubble with the decoder pen simultaneously leaves a dark mark inside the bubble and

reveals a previously hidden confirmation code printed in invisible ink.

If the voter wishes to verify her vote later on the election website, she can copy her ballot ID and her revealed confirmation codes onto a voter verification card. She keeps the verification card for future reference. She then takes her ballot to the scanning station and feeds the ballot into an optical scanner, which reads the ballot ID and the marked bubbles.

If a voter makes a mistake, she can ask a poll worker to replace her ballot with a new one. The first ballot is marked “spoiled,” and its ballot ID is added to the list of spoiled ballot IDs maintained by the election judges.

The voter can verify her vote on the election website by checking that her revealed confirmation codes and ballot ID have been posted correctly. If she finds any discrepancy, the voter can file a complaint through the website, within a complaint period. When filing a complaint, the voter must provide the confirmation codes that were revealed on her ballot as evidence of the validity of the complaint.

Ballots The Scantegrity ballot looks similar to a conventional optical scan ballot (see Figure 2 for a sample ballot used in the election). It contains a list of the choices and bubbles beside each choice. Marking a bubble reveals a random 3-digit confirmation code.

Confirmation Codes The confirmation codes are unique within each contest on each ballot, and are generated independently and uniformly pseudorandomly. The confirmation code corresponding to any given choice on any given ballot is hidden and unknown to any voter until the voter marks the bubble for that choice.

Digital Audit Trail Prior to the election, a group of election trustees secret-share a seed to a pseudorandom number generator (PRNG). The trustees then input their shares to a trusted workstation to generate the pseudorandom confirmation codes for all ballots, as well as a set of tables of cryptographic commitments to form the digital audit trail. These tables allow individual voters to verify that their votes have been included in the tally, and allow any interested party to verify that the tally has been computed correctly, without revealing how any individual voter voted.

Auditing After the election, any interested party can audit the election by using software to check the correctness of the data and final tally on the election website. Additionally, at the polling place on the day of the election, any interested party can choose to audit the printing of the ballots. A print audit consists of marking all of the

⁴See <http://www.takomaparkmd.gov/clerk/agenda/items/2009/090809-3.pdf>, section 2-D, page 2.

bubbles on a ballot, and then either making a photocopy of the fully-marked ballot or copying down all of the revealed confirmation codes. The ballot ID is recorded by an election judge as audited. After the election, one can check that all of the confirmation codes on the audited ballot, and their correspondence with ballot choices, are posted correctly on the election website.

5 Implementation

The election required a cryptographic *backend*, a *scanner*, and a *website*. These 3 components form the basic election system and their interaction is described in Figure 1. In addition, Takoma Park required software to resolve write-in candidate selections and produce a formatted tally on election night.

Scantegrity protects against manipulation of election results and maintains, but does not improve, the privacy properties of optical scan voting systems that use serial numbers. To compromise voter privacy using Scantegrity features, an attacker must associate receipts to voters and determine what confirmation numbers are associated to each candidate. This is similar to violating privacy by other means; for example, an attacker could compromise the scanner and determine the order in which voters used the device, or examine physical records and associate serial numbers to voters. The scanner and backend components protect voter privacy, but the website and the write-in candidate resolver do not because they work with public information only.

Each component is written in Java. We describe the implementation and functions of each one in the following sections.

Backend The cryptographic backend that provides the digital audit trail is a modified version of the Punchscan backend [21]. This backend is written in Java 1.5 using the BouncyCastle cryptography library.⁵ Key management in the Punchscan backend is handled by a simple threshold [25] cryptosystem that asks for a username and password from the election officials.

We chose the Punchscan backend over newer proposals [7] because it had already been implemented and tested in previous elections [13, 28]. At the interface between the Scantegrity frontend and the Punchscan backend, as described in [23], the permutations used by Punchscan are matched to a permutation of precomputed confirmation codes for Scantegrity that correspond to the permutation of codes printed on the ballot.

The Punchscan backend uses a two-stage mix process based on cryptographic commitments published before the election. Each mix, the *left mix* and the *right mix*,

takes marked positions as input, shuffles the ballots, and reorders each marked position on each ballot according to a prescribed (pre-committed) permutation. The result is the set of cleartext votes, where position 0 corresponds to candidate 0, 1 to 1, etc. Between the two mixes, for example, position 0 may in fact correspond to candidate 5, depending on the permutation in the *right mix*.

The Punchscan backend partitions [22] each contest such that each contest is treated as an independent election with a separate set of commitments. In the case of Takoma Park, each ward race and the mayor's race are treated as separate elections. (The announcement of separate mayoral race vote counts for each ward is required by Takoma Park). The scanner is responsible for creating the input files for each individual election.

Election officials hold a series of meetings using the backend to conduct an election. Before the election, during *Meeting 1* (Initialization), they choose passwords that are shares of a master key that generates all other data for the election in a deterministic fashion. After each meeting, secret data (such as the mapping from confirmation codes to candidates) is erased from the hard drive and re-generated from the passwords when it is needed again. In Meeting 1 the backend software creates a digital audit trail by committing to the Punchscan representation of candidate choices and to the *mixset*: the left and right mix operations for each ballot. Later, during *Meeting 2* (Pre-Election Audit), the backend software responds to an audit of the trail demonstrating that the mixset decrypts ballots correctly. At this time, the backend also commits to the Scantegrity front-end, consisting of the linkage between the Scantegrity front-end and its Punchscan backend used for decryption.

After the election, election officials run *Meeting 3* (Results), publishing the election results and the voted confirmation numbers. For the purposes of the tally audit, the system also publishes the outputs of the left and right mixes. In *Meeting 4* (Post-Election Audit), officials respond to the challenges of the tally computation audit. Either the entire *left mix* or the entire *right mix* operations are revealed, and the auditor checks them against data published in *Meeting 3*.

The *Meeting 4* audit catches, with probability one half, a voting system that cheats in the tally computation. To provide higher confidence in the results, the backend creates multiple sets of left and right mixes; in Takoma Park, we created 40 sets for each election, 20 of which were audited. Given 2 contests per ballot and 40 sets of left and right mixes, there are a total of 160 commitments per ballot in the audit trail, in addition to a commitment per contestant per ballot for each confirmation number (15-18, depending on the Ward).

The implementation uses two classes of "random" number sources. The first is used to generate the dig-

⁵<http://www.bouncycastle.org>

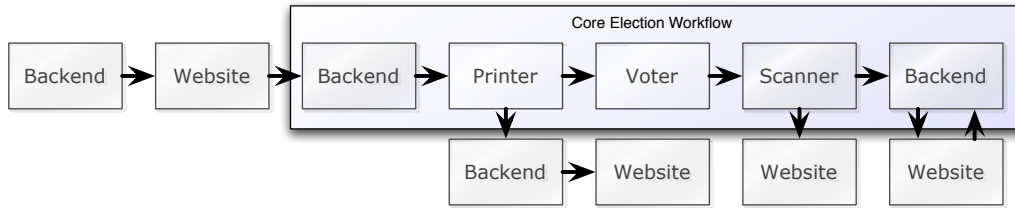


Figure 1: **Election Workflow.** The core election work flow in Scantegrity is similar to an optical scan election: a software backend creates ballot images that are printed, used by voters, and scanned. The results are fed to the backend which creates the tally. The audit capacity is provided by 3 extra steps: (1) create the initial digital audit trail and audit a portion of it, (2) audit the ballots to ensure correctness when printing, and (3) audit the final tally.

ital audit trail, and the second is used for auditing the trail. Both types of sources must be unpredictable to an adversary, and we describe each in turn.

Digital Audit Trail The Punchscan backend generates the mixes and commitments using entropy provided by each election official during initialization of the threshold encryption. This provided a “seed” for a pseudorandom number generator (based on the SHA256 hash function).

We also used this random source to generate the confirmation numbers when changing the Punchscan backend to support Scantegrity. Unfortunately, we introduced an error in the generation when switching from alphanumeric to numeric confirmation numbers as a result of findings in the Mock election (see Section 2). This resulted in approximately 8.5 bits of entropy as opposed to the expected 10 bits. We discovered this error after we started printing and it was too late to regenerate the audit trail.

The error increased the chance that an adversary could guess an unseen confirmation code to approximately one in 360 rather than the intended one in 1000; a small decrease in the protection afforded against malicious voters trying to guess unseen codes in order to discredit the system.

Auditing Random numbers are needed to generate challenges for the various auditing steps (print audit, randomized partial checking). These numbers should be unpredictable in advance to an adversary. They should also be “verifiable” after the fact as having come from a “truly random” source that is not manipulable by an adversary.

We chose to use the closing prices of the stocks in the Dow Jones Industrial Average as our verifiable but unpredictable source to seed the pseudorandom number generator (the use of stock prices for this purpose was first described in [11]). These prices are sufficiently unpredictable for our purposes, yet verifiable after the fact. However, it turns out that post-closing “adjustments” can sometimes be made to the closing prices, which can make these prices less than ideal for our purposes in

terms of verifiability.

Scanner Software The original intent of Scantegrity was to build on top of an existing optical scan system. There was no pre-existing optical scan system in use at Takoma Park, so we implemented a simple system using EeePC 900 netbooks and Fujitsu 6140 scanners.

The scanning software is written in Java 1.6. It uses a bash shell script to call the SANE scanimage program⁶ and polls a directory on the filesystem to acquire ballot images. Once an image is acquired it uses circular alignment marks to adjust the image, reads the barcode using the ZXing QRCode Library,⁷ and uses a simple threshold algorithm to determine if a mark is made on the ballot.

Individual races on each ballot are identified by ward information in the barcode, which is non-sequential and randomly generated. The ballot id in the barcode and the web verification numbers on each ballot are different numbers, and the association between each number type is protected by the backend system. Write-in candidate areas, if that candidate is selected by the voter, are stored as clipped raw images with the ballot scan results. Ballot scan results are stored in a random location in a memory mapped file.

The current implementation of the scanning software does not protect data in transit to the backend, which poses a risk for denial of service. Checking of the correctness of the scanner is done through the Scantegrity audit. The data produced by the scanner does not compromise voter privacy, but—assuming an attacker could intercept scanner data—voter privacy could be compromised at the scanner through unique write-in candidates on the ballot, through a compromised scanner, by bugs in the implementation, or by relying on the voter to make readable copies of the barcode to get a ballot id.

⁶<http://www.sane-project.org/>

⁷<http://code.google.com/p/zxing/>

Tabulator/Write-In Software At the request of Takoma Park we created an additional piece of software, the Election Resolution Manager (ERM), that allows election judges to manually determine for each write-in vote what candidate the vote should be counted toward. The other responsibility of the ERM is to act as part of the backend. It collates data from each scanner and prepares the input files for the backend.

To resolve write-ins with this software, the user cycles through each image, and either types in the name of the intended candidate or selects the name from a list of previously identified candidates composed of the original candidates and any previously typed candidate names. The user is not shown the whole ballot, so he does not know what the other selections are on that ballot, or what rank the write-in was given. We call this process *resolving* a vote because the original vote is changed from the generic "Write-In" candidate to the candidate that was intended by the voter. The ERM produces a PDF of each image, the candidate selection for that image, and a unique number to identify the selection.

Scantegrity handles write-in candidates just like other optical scan systems by treating the write-in position as a candidate. Therefore, the backend does not know how each write-in position was resolved, and two results records are created: one with write-in resolution provided by the ERM, and one without write-in resolution provided by the backend.

To check the additional record generated by the ERM, an observer reduces the resolved results record and verifies that the set of resolved ballots is the same as the set of unresolved ballots. To audit that the judges chose the correct candidates for each write-in, the observer refers to the PDF generated during write-in resolution. The PDF allows the observer to reference each resolved ballot entry in the resolved results file and verify that the image was properly transcribed.

One caveat of this approach is that if a write-in candidate wins, a malicious authority could modify these images to change results, but could not deny that the write-in position had received a winning number of votes. This situation would require additional procedures to verify the write-ins (e.g. a hand count, and/or careful audit of the transcriptions by each judge).

Website Beyond communicating the election outcome itself, the role of the election website is to serve as a "bulletin board" (BB) to broadcast the cryptographic audit data set (i.e., cryptographic commitments, responses to audit challenges, etc). In addition, voters can use this website to check their receipts, and file a dispute if the receipt is misreported. We provided an implementation with these features written in Java 1.6. It used the Stripes

Framework⁸ and an Apache Derby database backend.⁹ In practice, we only used part of this implementation.

Originally, our plan was to have Takoma Park host the website, but officials chose a hybrid approach where they hosted election information and results. That website would link to our server to provide a receipt checking tool and audit data. After the election, officials would provide us with a copy of the public data files to publish. This decision caused a number of changes to our approach.

We decided to only use the receipt checking code from the implementation, and, to make downloading more convenient for auditors, post all election data on our publicly available subversion repository.¹⁰ Additionally, both auditors agreed to mirror the data.

A primary security requirement for the Scantegrity BB is to provide authenticated broadcast communication from election officials to the public. We met this requirement with digital signatures. A team member (Carback) created signed copies of each file with gnupg¹¹ using his public key from May 28, 2009.

Without authenticated communication, it would be impossible to prove if different results were provided to different people. Our specific approach to the website requires observers to verify signatures and check with each other if they receive identical copies of the data (and verify the consistency of the signatures over time). Our auditors, Adida and Zagorski, performed these actions, but we do not know the extent of this communication otherwise. As usual with our approach to Scantegrity, we are enabling **detection** of errors (genuine or malicious).

There are several potential threats to the bulletin board model—we will briefly enumerate some of them. At a high level, threats pertain primarily to misreporting of results, or to voter identification. With regard to results reporting, an adversary may attempt to misreport results by substituting actual election data with false data. In the event that all parties verify signatures of information they receive, and check consistency with the signed files, incorrect confirmation codes on the bulletin board would be detected by voters, and incorrect computation of the tally by anyone checking the tally computation audit. If the voter checking confirmation codes does not check consistency with the rest of the bulletin board (by, for example, downloading the bulletin board data, checking all the signatures and checking that his or her confirmation code is also correctly noted in the entire bulletin board data) he or she may be deceived into believing their ballot was accurately recorded and counted. Similarly, if

⁸<http://www.stripesframework.org/>

⁹<http://db.apache.org/derby/>

¹⁰<http://scantegrity.org/svn/data/takoma-nov3-2009/>

¹¹<http://www.gnupg.org/>

the various signatures are not cross checked across individuals or observed over time, an adversary may replace the confirmation codes after they have been checked, or send different ones to voters and to auditors. An adversary may also attempt an identification attack, whereby the objective is to link voter identities with receipt data, such as by recording IP addresses of voters who check their receipts.

6 The Election

In this section, we describe the election as events unfold chronologically over time.

6.1 Preparations

Preparations for the election include running the first 2 backend meetings, and creating the ballot.

Independent Auditors The Board of Elections requested cryptographers Dr. Ben Adida (Center for Research on Computation and Society, Harvard University) and Dr. Filip Zagórski (Institute of Mathematics and Computer Science, Wrocław University of Technology, Poland) to perform independent audits of the digital data published by Scantegrity in general, and of the tally computation in particular. Dr. Adida¹² and Dr. Zagórski¹³ maintained websites describing the audits and the results of the audits, and Dr. Adida also blogged the audit.¹⁴ Before the election, Dr. Adida pointed out several instances when the Scantegrity information was insufficient; Scantegrity documentation was updated as a result.

The Board of Elections also requested Ms. Lillie Coney (Associate Director, Electronic Privacy Information Center and Public Policy Coordinator for the National Committee for Voting Integrity (NCVI)) to perform print audits on Election Day. Ms. Coney chose ballots at random through the day, exposed the confirmation codes for all options on the ballot, and kept these with her until after the end of the complaint period, when Scantegrity opened commitments to all unvoted and unspoiled ballots (and hence to all ballots she had audited). Ms. Coney then checked that the correspondence between codes and confirmation numbers on her ballots matched those on the website.

Both tasks, of print audits and digital data audits, can be performed by voters. Digital data audits can also be performed by any observers. In future elections, when the general population and Takoma Park voters are more

familiar with end-to-end elections, it is anticipated that voters (and, in particular, candidate representatives) will perform such audits.

Meeting 1 Four election officials (the City Clerk, the Chair, Vice Chair and a member of the Board of Elections: Jessie Carpenter, Anne Sergeant, Barrie Hofmann and Jane Johnson, respectively) were established as election trustees in *Meeting 1*, held on October 12 2009.

It was explained to the trustees that, through their passwords, they would generate the confirmation codes and share the secret used to tally election results. Further, it was explained that, without more than a threshold of passwords, the election could not be tallied by Scantegrity, and that if a threshold number of passwords was not accessible (if they were forgotten, for example, or trustees were unavailable due to sickness) the only available counts would be manual counts. A threshold of two trustees was determined based on anticipated availability of the officials, and it was explained that two trustees could collude to determine the correspondence between confirmation numbers and codes, and hence that each trustee should keep her password secret.

The trustees generated commitments to the decryption paths for each of 5000 ballots per ward (for six wards). Scantegrity published the commitments on October 13 2009 at 12:13am.

Meeting 2 In *Meeting 2*, held on October 14, 2009, trustees used Scantegrity-written code to respond to challenges generated using stock market data at closing on October 14. Half of the ballot decryption paths committed to in *Meeting 1* were opened. Additionally, trustees constructed ballots (associations between candidates and confirmation codes) at this meeting, and generated commitments to them. Scantegrity published the stock market data, the challenges, and the responses.

Ballot Design The ballot used for the 2009 election was based on ballots used for the 2007 election. We made the conscious choice to modify (as little as possible) a design already used successfully in a past election, and not to use the ballot we had designed for the mock election. The main reason for reusing the ballot design was that it would be familiar to voters. The ballot was required to contain instructions in both English and Spanish: marking instructions, instructions for write-ins, instructions for IRV and any Scantegrity-related instructions (see Figure 2).

Printing Ballots We use “invisible” ink to print the marking positions that reveal confirmation codes to voters. We used refillable inkjet cartridges in multiple color

¹²<http://sites.google.com/site/takomapark2009audits/>

¹³<http://zagorski.im.pwr.wroc.pl/scantegrity/>

¹⁴<http://benlog.com/articles/category/takoma-park-2009/>

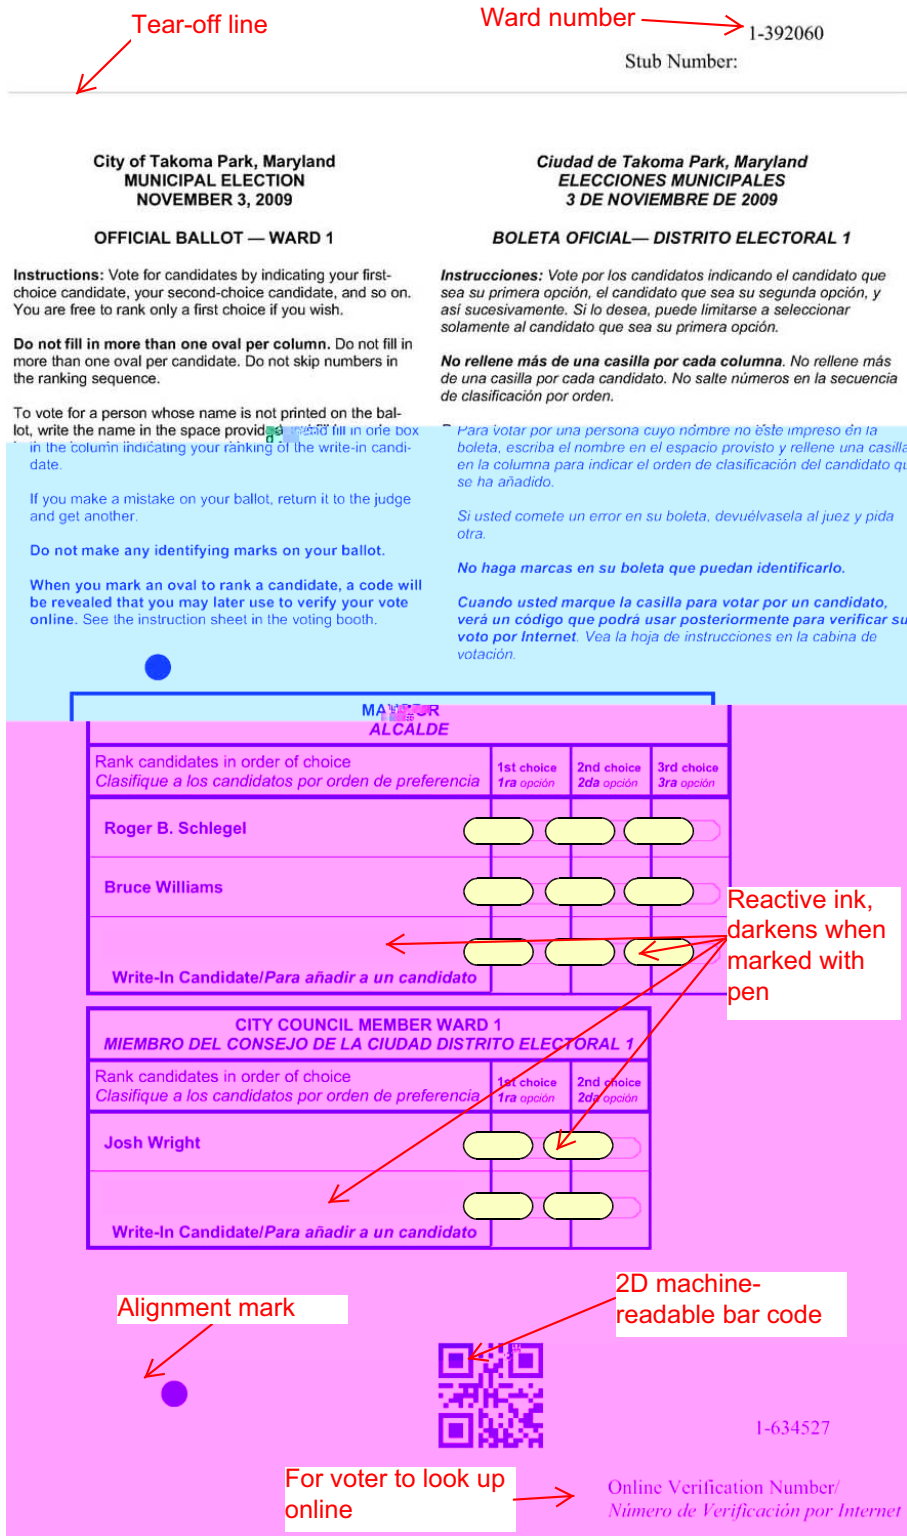


Figure 2: An unmarked Takoma Park 2009 ballot for Ward 1 showing instructions in Spanish and English, the options, the circular alignment marks, the 2D barcode, the ballot serial number (on the stub, meant for poll workers to keep track of the number of ballot used) and the online verification number (for voters to check their codes). The true ballot was printed on legal size paper and was hence larger than shown.

positions of an Epson R280 printer to print confirmation codes. The ink is not actually invisible, but looks like a yellow bubble before marking and a dark bubble with light yellow codes after marking.¹⁵

We initially began printing with 6 printers, but they proved unreliable. It was our expectation that using large amounts of commodity hardware would scale, but it did not. We did not anticipate the number of failure modes we experienced and our printing process was delayed by approximately 1 and a half days.

Ballot Delivery Mail-in (absentee) ballots were delivered to the City Clerk on 16 October. Early, in-person voting ballots were delivered on October 27 for early voting on October 28, and all other ballots a couple of days later on October 30.

Absentee ballots were identical to in-person voting ballots except they did not contain online verification numbers and voters were not given any instructions on checking confirmation numbers online. They were returned by mail in double envelopes and scanned with the early votes. Confirmation numbers for these ballots were, however, made available online after scanning, so that there was no distinction in published data between absentee and in-person voted ballots.

The board decided to issue ballots without confirmation numbers due to the small number of anticipated absentee votes and the costs associated with mailing ballots with special pens. Mailing the ballots with confirmation codes would allow verification of confirmation codes, but opens up new attacks: the possibility of false charges of election fraud by adversaries who might expose confirmation codes and reprint ballots, or use expensive equipment to attempt to determine the invisible codes. Strong verification for absentee ballots is an ongoing research subject within the Scantegrity team.

Early in-person voters used Scantegrity ballots with all Scantegrity functionality, except that the early votes were scanned in after the polls closed on Election Day, and not by voters themselves. Voters were, however, provided verification cards and could check confirmation codes for these ballots online.

Poll Worker Training Several training sessions were held in the weeks prior to the election. Manuals from the previous election were updated and a companion guide was created with Scantegrity-specific instructions. Election judges were given these two manuals, and a member from our team demonstrated the voting process at one session.

¹⁵See <http://scantegrity.org/~carback1/ink> for more information on the printing process

Voter Education Voter education for this election focused on online verification. Articles in the City newspaper before the real election indicated that voters could check confirmation numbers online; this was also announced on the city's election website.¹⁶

Scanner Setup We attempted to minimize, not prevent,¹⁷ the potential for using the wrong software by installing our software on top of Ubuntu Linux on SD flash cards, setting the "read-only" switch on each card, and setting up the software to read and write to USB sticks. We fingerprinted the first card after testing with the sha1sum utility and cloned it to a second card for the other netbook. Each netbook was set to boot from the card and BIOS configuration was locked with a password.

Both flash cards were checked with the sha1sum utility then placed into the netbook which was placed into a lock box and delivered to Takoma Park. The USB sticks were initialized with scanner configuration files. We uniquely identified each scanner by changing the ScannerID field in the configuration files, then we placed the corresponding USB sticks (3 for each netbook) into the lock box.

Upon delivery of the scanners the day before the election, we gave election officials the lock box keys and showed them how to open the lock boxes. We confirmed with election officials the contents of each box and the officials verified, with our assistance, that the USB memory sticks did not contain any ballot data by looking at the configuration file and making sure the ballot data file was blank.¹⁸ To protect against virus infection on the sticks we set them to read-only for this procedure.

6.2 Election Day

On Election Day, November 3, 2009, polls were open from 7 am to 8 pm at a single polling location, the Takoma Park Community Center. Several members of the SVST were present through most of the day in the building in case of technical difficulty. One SVST member was permitted in the polling room at most times as an observer, and a couple of SVST members were present in the vestibule giving out and collecting survey forms through most of the day. Lillie Coney of the Electronic Privacy Information Center, who performed a print audit on the request of the Board of Elections, was present in the polling room through a large part of the day.

¹⁶<http://www.takomaparkmd.gov/clerk/election/2009/>

¹⁷Scantegrity would detect manipulation at the scanner. A better solution would use trusted hardware technology (e.g. a TPM [14]).

¹⁸These were the only 2 files on the disk at this time. Additionally, election officials did not check fingerprints on the flash cards. Since no 3rd party had reviewed the code or fingerprinted it they relied on our chain of custody.

Starting the Election The scanner was the only SVST equipment to set up and it was a turn key system. Election judges needed to plug in the USB sticks and power on the netbooks. The scanner was attached to a scanning apparatus, and cables were run into the lockbox that contained the netbook. When ready, the scanner would beep 3 times. After reading a ballot, the scanner would beep 1 time. During shutdown, the scanner would beep another 3 times. If there were any failure modes the scanner would beep continuously or not beep at all.

Election judges set up the check-in tables, pollbooks, and voting booths. The election started on time.

Voting The election proceeded quite smoothly, with very few (small) glitches. An SVST member was able to assist polling officials in fixing a problem with their poll books (not provided by Scantegrity). Voters had some initial problems with the use of the scanner and the privacy sleeve, some seeking assistance from election judges who also had difficulty. After an explanation to the election judges by the Chair of the Board of Elections, the use of the scanner was considerably smoother. With a few ballots, the privacy sleeve was not letting go of the ballots; one ballot was mangled considerably but scanned fine. Seventeen scanned ballots had lines on them that caused the scanner to be unable to read votes, and one ballot had alignment marks manipulated such that it was also unreadable. Images of all unreadable scans are saved, so we were able to manually enter in these votes. Of the seventeen ballots, many ballots had a line in the same location, which is consistent with there being a foreign substance on a ballot put into the scanner. These problems did not affect our ability to count the votes.

During the day, Ms. Coney chose about fifty ballots at random, uniformly distributed across wards, and exposed the confirmation codes for all options for the ballots. A copy of each ballot was made for her to take with her; the copies were signed by the Chair of the BoE. Neither Ms. Coney nor SVST members had any interaction with voters.

Towards the end of the day, after the local NPR station carried clips from an interview with the Chair of the Board of Elections and a voter, the polling station saw a large increase in the number of voters, with the line taking up much of the floor outside the polling room. The SVST prepared to print more ballots, but this was not required. The number of printed ballots ended up being almost twice the number of voted ballots.

Absentee and early voted ballots were scanned in after the closing of polls. Afterward, the scanners were shut down. The chief judge opened each lock box, set all sticks to read only, removed 2 USB sticks (leaving the third with the scanning netbook), and locked the lock

box. Our team was given 1 stick for the ERM system. The other was kept by the city.

In *Meeting 3a*, trustees used Scantegrity code to generate results without provisional ballots at about 10 pm. The Chair of the Board of Election announced the results to those present at the polling place at the time (including candidates, their representatives, voters, etc.); this was also televised live by the local TV station. Confirmation codes and the election day tally were posted on the Scantegrity website.

6.3 After the Election

On the next day, around 2 pm, results including verified provisional ballots were published. Takoma Park representatives had announced a tally without provisional ballots the night before, and followed with the tally that included verified provisionals in accordance with standard Takoma Park procedures. The final *Meeting 3* results were published on November 4th just before midnight.

The number of registered voters were 10,934 and 1728 votes were cast (15.8%). The city-certified final tally for each contest is provided in Table 1. In each race, a majority was won after tallying after the voter's first choice.

Hand Count and Certification Following a hand count performed by representatives from both the SVST and Takoma Park, the Chair of the Board of Elections certified the results of the hand count to the City Council at 7 pm on November 5. The hand count and the Scantegrity count differed because officials were able to better determine voter intent during the hand count. For example, in the mayoral race, the scanner count determined that 646 votes were cast for candidate Schlegel, 972 for Williams, 15 for various write-in candidates, and 90 were not cast. The certified hand count totals were 664 votes for Schlegel, 1000 for Williams and 17 for write-in candidates. Thus 48 of a total of 1681 votes in this race would not have been counted by a scanner count alone. The discrepancy was caused by voters marking ballots outside of the designated marking areas. Such marks, while not read by the scanner by definition, are considered valid votes by Takoma Park law. Similarly, 8 of a total of 447 votes for Ward 1 council member, 8 of 251 for Ward 2, 16 of 431 for Ward 3, 10 of 210 for Ward 4, 2 of 81 for Ward 5 and 11 of 199 for Ward 6 were added to scanner vote totals after hand counting.

Post-Election Audit During *Meeting 4*, held on November 6 at 6 p.m., trustees used Scantegrity-written code to reveal all codes on voted ballots, and to reveal everything for all the ballots that were not spoiled or voted

Mayor	Votes	Ward	Councilor	Votes	Ward	Councilor	Votes
Roger B. Schlegel	664	Ward 1	Josh Wright	434	Ward 4	Terry Seamens	196
Bruce Williams	1000		Write-ins	13		Eric Mendoza	12
Write-ins	17	Ward 2	Colleen Clay	236		Write-ins	2
			Write-ins	15	Ward 5	Reuben Snipper	71
		Ward 3	Dan Robinson	397		Write-ins	10
			Write-ins	34	Ward 6	Navid Nasr	61
						Fred Schultz	138
						Write-ins	0

Table 1: City certified election results for the Mayor’s race and each City Councilman’s race.

upon. Trustees also responded to pseudo-random challenges generated by stock market results at closing on November 6. Scantegrity published all data on November 7th around 9am. While the SVST could have chosen to use closing data on an earlier date, such as November 4 or November 5, which could have been more stable, the team chose to stick to its earlier-announced plan (of using the freshest stock market data) for the sake of consistency.

On November 9, 2009, Dr. Adida and Dr. Zagórski independently confirmed that Scantegrity correctly responded to all digital challenges. In particular, they confirmed that the tally computation audit data was correct. Both made available independently-written code on their websites that voters and others could use to check the tally computation commitments. The Chair of the BoE mentions that several voters have shown an interest in running the code made available by Drs. Adida and Zagórski, and that she expects that Takoma Park voters will use the code to perform some audits themselves in the next few months.

Confirmation Codes and Complaints The period for complaints regarding the election (including complaints about missing confirmation codes) expired at 6 pm on November 6. The Scantegrity website recorded 81 unique ballot ID verifications, of which about 66 (almost 4% of the total votes) were performed before the deadline. The SVST was also told by a BoE member that at least a few voters checked codes on auditor websites. Both Dr. Adida and Dr. Zagórski made the confirmation codes available on their websites after the election.

The number of voters who checked their ballots online before the Takoma Park complaint deadline (66), while not large, was sufficient to have detected (with high probability) any errors or fraud large enough to have changed the election outcome. (Detailed calculations omitted here; these calculations are not so simple, due to the use of IRV.)

Scantegrity received a single complaint by a voter who had trouble deciphering a digit in the code and noted it

as “0,” while the Scantegrity website presented it as “8.” The voter requested that codes be printed more clearly in the future. He also stated that if he were not a trusting individual, he would believe that he had proof that his vote was altered.

All codes for all voted ballots were revealed after the dispute resolution period, and all commitments verified by two independent auditors, Dr. Adida and Dr. Zagórski. Hence, the probability that the code was in error is very small, albeit non-zero. Scantegrity does not believe the code was in error, and there were no other complaints regarding confirmation numbers.

Print Audits Dr. Zagórski provided an interface allowing Ms. Coney to check the commitments opened by Scantegrity in Meeting 4 against the candidate/confirmation-code correspondence on the ballots she audited. In her report [12], she confirmed that the correspondence between confirmation numbers and candidates on all the printed ballots audited by her was correctly provided by the interface.

Followup The Board of Elections and an SVST representative met to discuss the election and opportunities for improvement several weeks after the election. Both sides were largely satisfied with the election. Conversations have begun regarding the use of Scantegrity in the next municipal election at Takoma Park, to be held in November 2011. No decisions have been taken.

7 Surveys and Observations of Voter Experiences

To understand the experiences of voters and poll workers, we timed some of the voters as they voted, asked voters and poll workers to fill out two questionnaires, and informally solicited comments from voters as they left the precinct building. Approved by the Board of Elections and UMBC’s Institutional Review Board, our procedures respected the constraint of not interfering with the elec-

tion process. This section summarizes the results of our observations and surveys.

Timing Data Sitting unobtrusively as official observers in a designated area of the polling room for part of the day, two helpers (not members of the Scantegrity team) timed 93 voters as they carried out the voting process. Using stopwatches, they measured the number of seconds that transpired from the time the voter received a ballot to the time the voter began walking away from the scanner.

Voting times ranged from 55 secs. to 10mins. (the second longest time was 385 secs.), with a mean of 167 secs. and a median of 150 secs. On average, voters who appeared older took longer than voters who appeared younger. Most of the time was spent marking the ballot. The average time to vote was significantly faster than during the April 2009 mock election, when voters took approximately 8 mins. on average due primarily to scanning delays [26].

The observers noted that many voters did not fully use the privacy sleeve as intended, removing the ballot before scanning rather than inserting the privacy sleeve with the ballot into the scanning slot. Two of the 93 observed voters initially inserted the privacy sleeve upside-down, causing the ballot not to be fed into the scanner (even though the scanner could read the ballot in any orientation). A few ran into difficulties trying to insert the sleeve with one hand while holding something else in the other hand.

Election Day Comments From Voters As voters left the precinct building, members of the Scantegrity team conducting the written surveys, and a helper (a usability expert who is not a member of the Scantegrity team) solicited comments from voters with questions like, "What did you think of the new voting system?" The helper solicited comments 1:30-3:00pm and 7-8pm. A common response was, "It was easy."

Quite a few voters did not understand that they could verify their votes on-line and that, to do so, they had to write down the codenumbers revealed by their ballot choices. Some explained that they intentionally did not read any instructions because they "knew how to vote." Others failed to notice or understand instructions on posters along the waiting line, in the voting booth, on the ballot, and in the Takoma Park Newsletter.

In response, later in the day, we announced to voters as they entered the building that there is a new system; to verify your vote, write down the codenumbers. These verbal announcements seemed to have some positive effect, and there were fewer voter comments expressing lack of awareness of the verification option after we began the announcements. Nevertheless, some voters still

were unaware of the verification option. It was a humbling experience to see first-hand how difficult it can be to get across the most basic points effectively, especially the first time a new system is used.

Some of the voters complained about the double-ended pen, not knowing which end to use, or having trouble writing in candidates with the chisel-point (the narrow point was intended for write-ins). A small number of voters had difficulty seeing the codenumbers, perhaps largely because repeatedly pressing too hard could erode the paper. A few voters expressed concern about the difficulty of writing down the codenumbers, had the ballot been much longer or had there been a large number of competing candidates.

Many voters expressed a strong confidence in the integrity of elections, while a small minority expressed sharp distrust in previous electronic election technology. These feelings seemed to be based more on a general subjective belief rather than on detailed knowledge of election procedures and technology. Similarly, those expressing strong confidence in Scantegrity seemed to like the concept of verification but did not understand in detail why Scantegrity provides high outcome assurance.

Survey of Voter Experiences As voters were leaving the precinct, we invited them to fill out two one-sided survey forms: a field-study questionnaire, and a demographics questionnaire. The field study asked voters about the voting system they just used, with most answers expressed on a seven-point Likert scale. The last question invited voters to make any additional suggestions or comments. Each pair of forms had matching serial numbers to permit correlation of the field study responses with demographics. 271 voters filled out the forms.

Fifty-one voters wrote comments on the questionnaires, often pointing out confusion about various aspects of the process but with no consistent theme. (1) Some were unaware of verification option. (2) Some did not realize they were supposed to write down codenumbers. (3) Some found the pens confusing to use: they did not realize that the pens would expose codenumbers, and they did not know which end to use. (4) Some found codenumbers were hard to read. (5) Some did not understand how to mark an IRV ballot. (6) Some did not know how to place the ballot into the scanner. (7) One had no difficulty but wondered if seniors or people who speak neither English nor Spanish might have difficulties. (8) One wondered if the government might be able to discern his vote by linking his IP address used during verification with his ballot serial number and noting the time that he was issued a ballot (this may be possible if the cryptography is broken or in other scenarios, but it would be more direct to have the scanner log how he voted). (9) Many

suggested that it would have been helpful to have better instructions, including instruction while they wait in line.

Figure 3 shows how voters responded to four questions from the field study questionnaire. These results strongly show that voters found the voting system easy to use (Question 5), and that they had confidence in the system (Question 13). Question 10 showed that the option to check votes on line increased voter confidence in the election results. Question 9 showed that voters had confidence that the receipt alone did not reveal how they voted; this finding is notable given that it is widely suspected that many people erroneously believe that all E2E receipts reveal ballot choices. We plan to present detailed analysis of our complete survey data in a separate companion paper.

Survey of Poll Worker Experiences Each of the twelve poll workers was given an addressed and stamped envelope with two questionnaires (field study and demographics) to fill out and mail to the researchers after the election. The field study focused on their experiences administering Scantegrity, with most answer expressed on a seven-point Likert scale. This questionnaire also included four open-ended questions. Each pair of forms had matching serial numbers. Five forms were returned.

Poll workers noted the following difficulties. (1) There was too much information. (2) Some voters did not understand what to do, including how to create a receipt. (3) Some voters did not understand how to mark an IRV ballot. (4) The privacy sleeve was hard to use with one hand. (5) The double-ended pens created confusion. (6) Voters, poll workers, and the Scantegrity team have different needs. One wondered if Scantegrity was worth the extra trouble.

They offered the following suggestions: (1) Simplify the ballot. (2) Provide receipts so that voters do not have to copy codenumbers. (3) Develop better pre-election voter education.

8 Discussion and Lessons Learned

Overall, this project should be deemed a success: the goals of the election were met, and there were no major snafus. Many aspects of the Scantegrity design and implementation worked well, while some could be improved in future elections.

Technology Challenges Perhaps the most challenging aspect for future elections is scaling up ballot printing. The printers we used were not very reliable.

Variations on the Scantegrity design worth exploring include the printing of voter receipts (rather than having voters copy confirmation codes by hand)—there are

clearly security aspects to handle if one does this. The design should also be extended for better accessibility. The special pen might be improved by having only a single medium-tip point, rather than two tips of different sizes. The scanning operation and its interaction with the privacy sleeve should be studied and improved.

The website, while sufficient, might utilize existing research in distributed systems to reduce the expectations on observers and voters. The scanner could also be improved with more sophisticated image analysis, and also to better handle unreadable ballots. It only occurred to us after the election that the write-in resolution process could have greater utility if it were expanded to deal with unreadable and unclear ballots.

Real World Deployment of Research Systems As is common with many projects, too much was left until the last minute. Better project management would have been helpful, and key aspects should have been finalized earlier. Materials and procedures should be more extensively tested beforehand.

One of the most important lessons learned is the value of close collaboration and clear communication between election officials and the election system providers (whether they be researchers or vendors).

Another lesson learned is that it is both important to provide voters with clear explanations of the new features of a voting system, and to do so efficiently, with minimal impact on throughput. Resolving the tension between these requirements definitely needs further exploration. For example, it might be worthwhile to have an instructional video explaining the Scantegrity system that voters could watch as they come in. The permanent adoption of Scantegrity II in a jurisdiction would, however, alleviate the educational burden over time, as voters learn the system's features in successive elections.

Comparison with post-election audits It is interesting to compare Scantegrity with the other major technique for election outcome verification: post-election audits. Because these audits do not allow anyone to check that a particular ballot was counted correctly, they do not provide the level of integrity guarantee provided by Scantegrity.

Post-election audits, even those with redundant digital and physical records like optical scan systems, only address errors or malfeasance in the counting of votes and not in the chain of custody.¹⁹ In contrast, end-to-end

¹⁹Having multiple records may make an attacker's job harder, but note that the attacker only has to change the record that will ultimately be used and/or trusted (not necessarily both). Also, redundancy can work against a system, as changing a digital record in an obviously malicious way may allow time for a more subtle manipulation of the physical record.

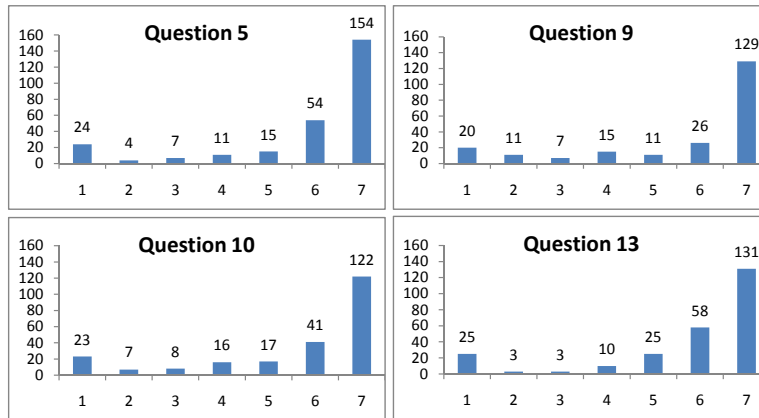


Figure 3: Voter responses to Survey Questions 5, 9, 10, 13 from all 271 voters completing the survey. Using a seven-point Likert scale, voters indicated how strongly they agreed or disagreed with each statement about the voting system they had just used (1 = strongly disagree, 7 = strongly agree). Each histogram shows the number of voters responding for each of the seven agreement levels. The four questions shown are the following: (5) Overall, the voting system was easy to use. (9) I have confidence that my receipt by itself does not reveal how I voted. (10) The option to verify my vote online afterwards increases my confidence in the election results. (13) I have confidence in this voting system.

voting systems such as Scantegrity provide a “verifiable chain of custody.” Voters can check that their ballots are included in the tally, and anyone—not just a privileged group of auditors—can check that those ballots are tallied as intended.

It must be admitted, however, that the additional integrity benefits provided by Scantegrity II come at the cost of somewhat increased complexity and at the cost of an increased (but manageable) risk to voter privacy (since ballots are uniquely identifiable). That said, some jurisdictions and/or election systems require or use serial numbers on ballots anyway, and we have proposed several possible approaches to appropriately destroy or obfuscate serial number information. Furthermore, it can be argued that a voter wishing to “fingerprint” a ballot can do so without being detected in current paper ballot systems simply by marking ovals in distinctive ways.

9 Conclusions

Traditional opscan voting systems have the clear benefit that “votes are verifiably cast as intended”—the voter can see for herself that the ballot is correctly filled out. Yet once her ballot is cast, the voter must place her trust in others that ballots are safely collected and correctly counted. With end-to-end voting systems these last two operations (collecting ballots and counting them) are verifiable as well: voters can verify—using their receipt and a website—that their ballot is safely collected with the others, and anyone can use the website data to verify that the ballots have been correctly counted. The Scantegrity

II voting system provides such end-to-end verification capability as an overlay on top of traditional opscan technology. Further development should improve scalability (esp. printing), usability (e.g. with printed receipts) and accessibility of the Scantegrity II system.

The successful use of the Scantegrity II voting system in the Takoma Park election of November 3, 2009 demonstrates that voters and election officials can use sophisticated cryptographic techniques to organize a transparent secret ballot election with a familiar voting experience. The election results show considerable satisfaction by both voters and pollworkers, indicating that end-to-end voting technology has matured to the point of being ready and usable for real binding governmental elections. This paper thus documents a significant step forward in the security and integrity of voting systems as used in practice.

Acknowledgments The authors would like to acknowledge the contributions of the voters of Takoma Park, the City Clerk, the Assistant City Clerk, all Board of Elections members since 2008 when this project was first proposed, and the independent auditors—Lillie Coney, Ben Adida and Filip Zagórski—to the success of the election. Vivek Relan and Bhushan Sonawane timed voters as they voted and helped assemble the privacy sleeves. Lynn Baumeister interviewed some voters as they left the precinct. Cory Jones provided general assistance and Alex Florescu and Jan Rubio assisted with ink creation.

Alan T. Sherman was supported in part by the Depart-

ment of Defense under IASP grants H98230-08-1-0334 and H98230-09-1-0404. Poorvi L. Vora was supported in part by The National Science Foundation under grant CNS 0831149. Jeremy Clark and Aleksander Essex were supported in part by Natural Sciences and Engineering Research Council of Canada (NSERC).

Disclosure Portions of the Scantegrity system may be covered by pending patents under applications US 2008/0272194, and US 2009/0308922. All source code was released under the GPLv2 software license.²⁰

References

- [1] ADIDA, B. *Advances in Cryptographic Voting Systems*. PhD thesis, MIT EECS Dept., 2006.
- [2] ADIDA, B. Helios: web-based open-audit voting. In *Proceedings of the 17th USENIX Security Symposium* (2008), pp. 335–348.
- [3] ADIDA, B., DEMARNEFFE, O., PEREIRA, O., AND QUISQUATER, J.-J. Electing a University President using Open-Audit Voting: Analysis of real-world use of Helios. *Proceedings of the Electronic Voting Technology Workshop / Workshop on Trustworthy Elections* (August 2009).
- [4] BENALOH, J. Simple verifiable elections. In *Proceedings of the 2006 USENIX/ACCURATE Electronic Voting Technology Workshop* (Berkeley, CA, USA, 2006), USENIX Association.
- [5] CHAUM, D. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM* 24, 2 (1981), 84–90.
- [6] CHAUM, D. Secret-Ballot Receipts: True Voter-Verifiable Elections. *IEEE Security and Privacy* 2, 1 (2004), 38–47.
- [7] 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. In *Proceedings of the 2008 USENIX/ACCURATE Electronic Voting Technology Workshop* (2008), pp. 1–13.
- [8] CHAUM, D., CARBACK, R. T., CLARK, J., ESSEX, A., POPOVENIUC, S., RIVEST, R. L., RYAN, P. Y. A., SHEN, E., SHERMAN, A. T., AND VORA, P. L. Scantegrity II: End-to-End Verifiability by Voters of Optical Scan Elections Through Confirmation Codes. *IEEE Trans. on Information Forensics and Security, special issue on electronic voting* 4, 4 (Dec. 2009), 611–627.
- [9] CHAUM, D., ESSEX, A., CARBACK, R., CLARK, J., POPOVENIUC, S., SHERMAN, A. T., AND VORA, P. Scantegrity: End-to-End Voter Verifiable Optical-Scan Voting. *IEEE Security and Privacy Magazine* 6, 3 (May/June 2008), 40–46.
- [10] CHAUM, D., RYAN, P. Y., AND SCHNEIDER, S. A. A practical, voter-verifiable, election scheme. Technical Report Series CS-TR-880, University of Newcastle Upon Tyne, December 2004.
- [11] CLARK, J., ESSEX, A., AND ADAMS, C. Secure and observable auditing of electronic voting systems using stock indices. In *Proceedings of the 2007 IEEE Canadian Conference on Electrical and Computer Engineering* (2007).
- [12] CONEY, L. Report on the Manual Ballot Audit: Takoma Park, Maryland, November 3 2009 Election, 19 November 2009. Electronic Privacy Information Center, http://epi.c.org/privacy/voting/takoma_park_audit.pdf.
- [13] ESSEX, A., CLARK, J., CARBACK, R. T., AND POPOVENIUC, S. Punchscan in Practice: an E2E Election Case Study. In *Proceedings of the 2007 IAVoSS Workshop on Trustworthy Elections* (2007).
- [14] FINK, R. A., SHERMAN, A. T., AND CARBACK, R. Tpm meets DRE: reducing the trust base for electronic voting using trusted platform modules. *Trans. Info. For. Sec.* 4, 4 (2009), 628–637.
- [15] FISHER, K., CARBACK, R., AND SHERMAN, A. T. Punchscan: Introduction and System Definition of a High-Integrity Election System. In *Proceedings of the 2006 IAVoSS Workshop on Trustworthy Elections* (2006).
- [16] HOSP, B., JANSON, N., MOORE, P., ROWE, J., SIMHA, R., STANTON, J., AND VORA, P. Citizen-Verified Voting. Presentation at DIMACS Workshop on Electronic Voting – Theory and Practice, May 2004, <http://dimacs.rutgers.edu/Workshops/Voting/slides/vora.ppt>.
- [17] HUBBERS, E., JACOBS, B., SCHOENMAKERS, B., VAN TILBORG, H., AND DE WEGE, B. Description and Analysis of RIES, June 2008.
- [18] KUTYLOWSKI, M., AND ZAGÓRSKI, F. Verifiable Internet Voting Solving Secure Platform Problem. In *Advances in Information and Computer Security, Lecture Notes in Computer Science* (2007), vol. 4752, pp. 199–213.
- [19] NEFF, C. A. Practical high certainty intent verification for encrypted votes, 2004.
- [20] OFFICE FOR DEMOCRATIC INSTITUTIONS AND HUMAN RIGHTS. The Netherlands Parliamentary Elections 22 November 2006 OSCE/ODIHR Election Assessment Mission Report, March 12 2007. 28 pages.
- [21] POPOVENIUC, S., AND HOSP, B. An introduction to punchscan. In *Proceedings of the 2006 IAVoSS Workshop on Trustworthy Elections* (2006).
- [22] POPOVENIUC, S., AND STANTON, J. Undervote and Pattern Voting: Vulnerability and a mitigation technique. In *IAVoSS Workshop On Trustworthy Elections (WOTE 2007)* (University of Ottawa, Ottawa, Canada, June 2007).
- [23] POPOVENIUC, S., AND VORA, P. L. A framework for secure electronic voting. In *Proceedings of the 2008 IAVoSS Workshop on Trustworthy Elections* (2008).
- [24] SANDLER, D. R., DERR, K., AND WALLACH, D. S. VoteBox: a tamper-evident, verifiable electronic voting system. In *Proceedings of the 17th USENIX Security Symposium* (2008).
- [25] SHAMIR, A. How to Share a Secret. *CACM* 22, 11 (Nov 1979), 612–613.
- [26] SHERMAN, A. T., CHAUM, D., CLARK, J., ESSEX, A., HERRNISON, P. S., MAYBERRY, T., POPOVENIUC, S., RIVEST, R. L., SHEN, E., SHERMAN, A. T., AND VORA, P. L. Scantegrity Mock Election at Takoma Park. In *Proceedings of the 4th International Conference on Electronic Voting (EVOTE 2010)* (2010).
- [27] City of Takoma Park, Maryland City Election November 3, 2009 Certification of Election Results, November 2009. <http://www.takomaparkmd.gov/clerk/election/2009/results/2009cert.pdf>.
- [28] VoComp Voting System Competition. July, 2007. Portland, Oregon. <http://www.vocomp.org>.

²⁰<http://www.gnu.org/licenses/gpl-2.0.html>