

Chapter 10

The Scantegrity Voting System and Its Use in the Takoma Park Elections

Richard T. Carback

Draper Labs

David Chaum

Voting Systems Institute

Jeremy Clark

Concordia University

Aleksander Essex

Western University

Travis Mayberry

US Naval Academy

Stefan Popoveniuc

The George Washington University

Ronald L. Rivest

MIT

Emily Shen

MIT

Alan T. Sherman

UMBC

Poorvi L. Vora

The George Washington University

John Wittrock

AppNexus

Filip Zagórski

Wroclaw University of Technology

CONTENTS

10.1	Introduction	239
10.1.1	Key Properties	239
10.1.2	Outline	241
10.1.3	Organization of This Chapter	241
10.2	The Punchscan Voting System	242
10.2.1	Voter Experience	242
10.2.2	Election Set-Up	242
10.2.3	Ballot Printing and Voter Privacy	244
10.2.4	The First Binding E2E Election	244
10.2.5	Lessons Learned	246
10.3	The Scantegrity II Voting System	246
10.3.1	Ballot Features	247
10.3.2	The Scantegrity Back-End	248
10.3.3	Election Day	249
10.3.4	Posting and Tallying	250
10.3.5	Auditing the Results	250
10.3.6	Dispute Resolution	252
10.4	The 2009 Takoma Park Election	254
10.4.1	Requirements	254
10.4.2	Mock Election	255
10.4.3	The Election	256
10.5	The Remotegrity Voting System	259
10.5.1	Voter Experience	260
10.5.2	Security Properties	260
10.6	Auditogegrity	262
10.6.1	The Voting Process with Auditogegrity	263

10.6.2	Dispute Resolution Properties	263
10.7	Takoma Park Election, 2011	264
10.8	Usability Studies	266
10.8.1	Related Studies	266
10.8.2	Methodology	267
10.8.3	Known Limitations	268
10.8.4	Voter Response	269
10.8.5	Observational Results	269
10.8.6	Election Judge Response	271
10.8.7	Study Findings	272
10.9	Current Status of the Project	274
10.10	Conclusions	274
	Acknowledgments	276

10.1 Introduction

The Scantegrity project began with a simple question: is it possible to design a voting system offering the strong security properties of cryptographic end-to-end (E2E) election verification with the intuitive look and feel of a paper optical-scan ballot? This chapter recounts a decade-long research effort toward answering this question, from the design of Scantegrity’s precursor Punchscan, all the way to the first governmental election run by an E2E voting system.

The main focus of this chapter is on the Scantegrity II voting system (hereafter referred to as simply *Scantegrity*) and its use in the municipal elections of Takoma Park, Maryland, in 2009 and 2011. To our knowledge, the Takoma Park election of 2009 was the first use of an E2E-verifiable voting system in an in-person secret-ballot governmental election anywhere in the world, as well as being the first governmental election held in the United States to run on open-source software.

We also describe the *Punchscan* voting system and its use in the 2007 election of the University of Ottawa Graduate Students Association/Association Étudiant(e)s Diplômé(e)s (GSAÉD), which, to our knowledge, is the first time an E2E voting system was used in a binding election.¹ Additionally, this chapter describes the remote voting system *Remotegrity* and accessible Scantegrity variant *Audiotegrity*, and their use in the 2011 Takoma Park election. We also briefly recount a number of smaller, non-binding elections run during the course of this project.²

¹A binding election is one whose outcome is binding on a constituency, though it need not be a political election. It could be, e.g., a primary, party, union, stockholder or student government election. Whereas a non-binding election is effectively an opinion poll.

²These elections are documented on punchscan.org and scantegrity.org.

10.1.1 Key Properties

End-to-End Verifiable

Loosely speaking, a cryptographically end-to-end verifiable (E2E) voting system is designed to provide the following security properties:

- (a) A voter has the means to check that her ballot was included unmodified in the set of cast ballots, and receives evidence that convinces her if it is not.
- (b) Anyone has the means to check that the set of cast ballots were counted correctly, and receives convincing evidence if they were not.
- (c) The preceding checks are performed in a way that upholds the secrecy of every voter's ballot.

Additionally, Scantegrity provides an additional security property:

- (d) If a voter's ballot was recorded incorrectly, she can prove this fact to others without revealing her vote. Similarly, if her vote was recorded correctly, but the voter makes a spurious claim, the election officials can prove this fact to others.

We refer to this last property as *dispute resolution*. This particular property is important in an E2E setting because it allows the electorate as a whole to attribute fault to the actual at-fault party: the elections officials or the voter. It is also special in the sense that not all E2E systems provide it.

Most of the assumptions made by the Scantegrity voting system and its variants are standard to E2E voting systems: a threshold of election trustees is required to be honest for privacy, cryptographic primitives used are assumed secure and an append-only, public bulletin board with authenticated write-access is assumed to exist. In addition, Scantegrity's privacy and dispute resolution properties require that symbols printed in invisible ink or beneath scratch-off are not visible.

Practical Aspects

From a practical perspective, a major design goal of Scantegrity is to provide voters with an intuitive interface. We focused on paper optical-scan ballots for a variety of reasons: it is a predominant mode of ballot casting, and is legally mandated in a number of jurisdictions, such as the United States; it provides a degree of recoverability in the event of an electronic failure; used in a precinct-scan configuration, it does not require electricity at the polling place. Scantegrity itself offers a number of advantages over other optical-scan based E2E systems:

- (a) The verifiability elements are constructed as an overlay on the optical scan ballot, making them compatible with pre-existing layouts.

- (b) Participation in the verification process is opt-in; voters can (in principle) ignore the Scantegrity-related components and vote as they would otherwise on a plain paper optical-scan ballot.

10.1.2 Outline

In this chapter, we trace through a couple of iterations of deploying, redesigning, and redeploying an E2E system. We begin by describing the Punchscan system and its use in an election in 2007. We then describe the lessons learned from that election, which led to the development of Scantegrity and an early variant, *Scantegrity I*. The former improved on the dispute resolution properties of the latter, and was used by Takoma Park in 2009 and 2011. We recount some of the modifications we made to the original Scantegrity proposal for the election, and how feedback from a preliminary mock election further refined the system, as well as the results of the election itself.

Scantegrity was deployed again in the municipal election of Takoma Park in 2011. For this election, we developed and deployed the remote voting system Remotegrity and the accessibility-enhanced Scantegrity variant Audiotegrity. Both of these systems were tested in an open test before the 2011 election, and we document the feedback and how it contributed to changes to the systems finally fielded in 2011.

We end this chapter with a description of the results of usability tests that demonstrate high levels of confidence in Scantegrity, and high levels of support for the verification receipt, even when voters indicated they did not understand the underlying cryptographic mechanisms. While E2E system usability is an understudied subject, the data we collected after the Takoma Park mock and municipal elections currently represents the largest dataset of voter and poll worker reactions from an E2E system.

Following the two Takoma Park elections, there has been considerable activity towards the development of E2E voting systems intended for larger government elections. The state of Victoria, Australia, used vVote for elections in November 2014 (see Chapter 12, this volume), and Travis County, Texas, is considering STAR-Vote as a potential future voting system (see Chapter 14, this volume). There has also been an interest in the use of remote E2E systems.

10.1.3 Organization of This Chapter

This chapter is organized as follows. Section 10.2 describes the Punchscan voting system and the first binding E2E election. Section 10.3 describes the Scantegrity voting system and Section 10.4 its use in the Takoma Park election of 2009. Section 10.5 describes the Remotegrity system, and Section 10.6 describes the Audiotegrity variant of Scantegrity. Section 10.7 describes the use of all three systems in the Takoma Park election of 2011. Section 10.8 summarizes the results of our voter and poll worker surveys in both Takoma Park elections. Section 10.10 concludes.

10.2 The Punchscan Voting System

The Punchscan voting system represented an early attempt to combine E2E verifiability with a paper optical-scan ballot. In this section we describe the Punchscan voting system and our experience fielding it in the 2007 election of the University of Ottawa Graduate Students Association/Association Étudiant(e)s Diplômé(e)s (GSAÉD). We discuss our experiences in this election, and how they went on to motivate the design of the Scantegrity.

10.2.1 Voter Experience

A Punchscan ballot consists of two sheets of paper (see Figure 10.1). The upper sheet consists of a serial number and list of candidate choices. Beside each candidate is an independent and pseudorandomly assigned letter/symbol. Toward the bottom of the ballot, a number of holes are punched in the paper. A second sheet of paper contains letters that are printed such that they show through the holes on the top page. These letters are the same as those found on the top page, but in an independent pseudorandom order.


A bingo dauber is employed as the marking implement, and the voter selects one of the punched holes to daub. The holes are sized such that daubing will distribute ink to both pages. The voter marks her ballot as follows: she (1) chooses her preferred candidate, (2) observes the letter beside her preferred candidate, (3) marks the punched hole showing the corresponding letter, (4) selects one of the sheets (either top or bottom) and shreds it (e.g., using a paper shredder supplied in the voting booth). At this point she exits the polling booth and provides the remaining sheet to the polling staff to be scanned. The voter retains a copy of this sheet as a receipt. Notice that the receipt alone is not sufficient to recover voting intent:

1. If the voter retains the top sheet it shows which letters are associated with which candidates, and the punched hole the voter marked, but not the underlying letter,
2. If the voter retains the bottom sheet it shows which letter was marked, but not which candidate that letter corresponds to.

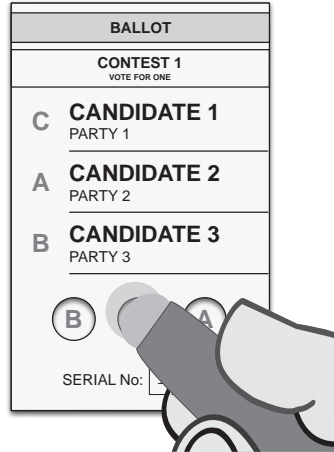
After the election, receipts are posted on an election website. If a voter's receipt is absent from the website or is incorrectly displayed, she is able to file a complaint, and can provide her receipt as evidence. The receipts are tallied in a verifiable manner, as described in [465].

10.2.2 Election Set-Up


Prior to the election, election administrators determine a set of officials and a minimum threshold of officials required to conduct the election tasks. In order to protect

BALLOT	
CONTEST 1 VOTE FOR ONE	
C	CANDIDATE 1 PARTY 1
A	CANDIDATE 2 PARTY 2
B	CANDIDATE 3 PARTY 3
	
SERIAL No:	1234-5678


(a) Unmarked Punchscan Ballot



(c) Marking the ballot (with bingo dauber)

BALLOT	
CONTEST 1 VOTE FOR ONE	
C	CANDIDATE 1 PARTY 1
A	CANDIDATE 2 PARTY 2
B	CANDIDATE 3 PARTY 3
	
SERIAL No:	1234-5678

(b) Marked top sheet

BALLOT	
CONTEST 1 VOTE FOR ONE	
CANDIDATE 1 PARTY 1	
CANDIDATE 2 PARTY 2	
CANDIDATE 3 PARTY 3	
	
SERIAL No:	1234-5678

(d) Marked bottom sheet

Figure 10.1: Punchscan ballot depicting a vote for Candidate 1. The voter retains either the marked top or bottom sheet as a receipt (the other is shredded).

ballot secrecy it is required that any valid subset of election officials contains *at least one* honest party. If all t officials collude, i.e., not even one of them is honest, they can link ballot receipts to the candidate voted for (the integrity of the tally is, however, invariant even with full collusion).

The officials first generate and threshold-share a secret value. This value seeds a pseudorandom number generator (PRNG) from which all the election data can be deterministically produced and reproduced. For each task, the election officials convene and use a trusted computing platform with no persistent memory to generate the seed from their secret shares and to generate the election data. The computer is trusted for privacy: that is, it is trusted to not leak election secrets. Aside from the public output produced during each session, no state is saved in any form. Every session regenerates the state from the shared secret. If the computer performs incorrect computations so as to change the election outcome, this will be detected during an election audit. That is, our trust model does not require that the computer be trusted for election integrity.

In the first meeting, a threshold number of officials input their shares along with the number of candidates, and the number of ballots to generate. Since voters may spoil ballots as part of the audit process, more ballots than voters need to be produced (by some predetermined factor). Cryptographic commitments to the verifiable shuffle and decryption are generated and publicly posted. All intermediate values, internal state and secret values are purged upon powering down the system.

In the second meeting, a threshold number of officials regenerate the shared secret seed which is then used to regenerate the election data. A cut-and-choose proof that the commitments are correctly formed is provided on the bulletin board. The challenges are generated randomly and non-interactively using stock market data and random extraction techniques described by Clark and Essex [156, 157].

10.2.3 Ballot Printing and Voter Privacy

As a pre-election audit, half the (unvoted) ballots, and associated proof-of-shuffle are publicly disclosed. The remaining ballots are printed using a private XML file containing the necessary information. In that sense the entity responsible for ballot printing must be trusted to maintain voter privacy. This *trusted printer problem* is a difficulty with any paper-based E2E election system. Toward addressing the trusted printer problem, Essex et al. [217, 220] proposed a system for *oblivious printing* by combining visual cryptography with invisible ink, and developed a number of secure multiparty protocols for generating and distributing ballot images under encryption. In subsequent work Essex et al. [221, 219] show how oblivious printing could be integrated into a Scantegrity-like voting system with fully distributed trust.

10.2.4 The First Binding E2E Election

In 2007 the University of Ottawa Graduate Students Association/Association Étudiant(e)s Diplômé(e)s (GSAÉD) deployed Punchscan to elect executive positions within the association and to conduct a referendum on a bylaw.³ The election consisted of five polling stations and was held over two days. There were approximately 1000 eligible voters, of which 154 voted. The election consisted of six contests, each of which had two candidates/options. While Punchscan offers support for a distributed election authority, GSAÉD opted to have the chief returning officer (CRO) act as the sole trustee.

To prevent the loss of votes in the event of a power outage, hard drive failure or other unpredictable event, we made a paper record of information on the scanned receipt. While this record does not contain information sufficient for a hand count, it could be used with the underlying cryptographic information to reconstruct the tally.

Ballot Printing. Under the supervision of the CRO we printed batches of ballots using five inkjet printers. We experienced difficulty with the top sheets jamming due to the holes. Printing took about 1 hour and upon completion, the ballots were placed into boxes, sealed and signed along the seal by the CRO.

Vote Casting. Voters were given their double-layer paper ballot on a clipboard, and the clipboard contained a plunger lock that fastened the ballot to the clipboard through aligned holes in one corner (such that removing the ballot would tear it). Voters filled out the ballot in a private voting booth. Upon returning from the booth, the voter was instructed to choose either the top or bottom sheet for shredding. This sheet was ripped out of the locked assembly and shredded in view of the poll workers. The assembly with the remaining sheet still locked in was returned to the poll worker. The poll worker unlocked the clipboard and removed the sheet. The sheet was scanned with an optical scanner and the software's interpretations of the marks were displayed on a computer screen. The voter approved the marks and cast the ballot electronically. The sheet was then placed in the printer, which printed symbols on the sheet according to the scanning software's interpretation of the marks (e.g., overlaying an X on unmarked positions and an O on marked positions), printed a digital signature of the marks, and printed an additional record of the marks in the corner of the ballot. The corner was cut off by the poll worker and retained as a paper backup, while the sheet itself was given to the voter as a receipt.

Vote Tallying. After the polling stations were closed, the Punchscan team, the CRO, and a scrutineer gathered to generate the results. The ballot receipts were uploaded to the Punchscan server and posted. The CRO then entered her passphrase to deterministically regenerate the election data that was committed to prior to the election. Once generated, the tally was computed using this information and the receipt infor-

³Authors Clark and Essex were members of GSAÉD at the time.

mation. The tally and a commitment to its proper computation were also posted on the Punchscan server. Through a web-interface, voters could check that their receipts match the posted information. Our server logs showed that the image files of 83 of the ballots were requested.

As with the pre-election audit, the following day, the tallying procedure was audited using stock market data and a cut-and-choose protocol. This audit is also based on an open specification and is software independent.

10.2.5 Lessons Learned

Over the course of the election the poll workers experienced some technical failures: at times the polling place software became unresponsive, the printers jammed or lost wireless connectivity needed to maintain the electronic pollbook. In these instances, the poll workers manually recorded the serial number and mark positions made by the voters. These were electronically input later on, and any transcription errors would be subject to detection through the online receipt verification check. The ink from bingo daubers also caused problems in the scanners if they were fed in before they were completely dry.

We confirmed that the security mechanisms of Punchscan were not well understood by the voters. In particular, voters did not understand the security benefits of the indirection of Punchscan. Many voters indicated that the voting process was burdensome. We observed voters who appeared to not realize that they were to receive a receipt and wanted to leave immediately after marking their ballots. Furthermore, when it was explained to them that they would receive a receipt, some voters refused it. Further, approximately 85% of the voters chose to shred the top sheet and keep the bottom page as their receipt, indicating that the choice of which sheet is kept is not uniformly random.

Toward Scantegrity. Overwhelmingly the biggest criticism of Punchscan centered around the ballot marking procedure. The voter is required to follow a strategy of indirection: mark the hole containing the letter appearing beside the desired candidate. While at some level this is not more complicated than many administrative tasks the average person encounters in day-to-day life, what eventually became clear to us from this exercise was that we needed to improve the usability of the ballot.

Following the GSAÉD election we formulated a new research question: is it possible to design an E2E-enabled ballot that (1) still allows the voter to simply produce a privacy-preserving receipt of their choice while (2) not requiring any special ballot marking strategy beyond that of conventional optical-scan. This question ultimately led us to develop Scantegrity.

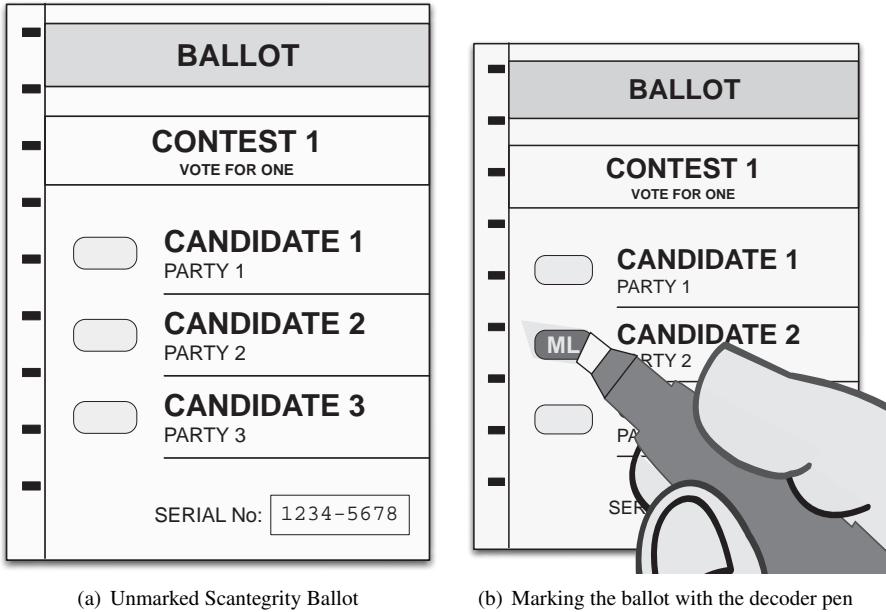


Figure 10.2: Scantegrity ballot depicting a vote for Candidate 2. The voter would retain confirmation code `ML` and serial number as a receipt.

10.3 The Scantegrity II Voting System

In this section, we consider Scantegrity, the successor to Punchscan, which is quite different in many regards from Punchscan and other E2E systems. Relative to Punchscan, Scantegrity is designed to improve upon practically, which makes itself apparent in a few ways. First, it uses a single-sheet paper ballot and a simple-to-perform obfuscation technique: code substitution. Next, Scantegrity is not designed as a replacement to existing voting systems but rather an augmentation. It can be added as a cryptographic layer to any optical scan system without interfering with the two existing methods for counting optical scan ballots: electronically from the scanned images and manually with the paper ballots. Finally, it has essentially no impact on the voter who does not wish to participate in the verification process.

10.3.1 Ballot Features

The Scantegrity ballot is shown in Figure 10.2. It consists of two parts: a **ballot** and a detachable **receipt**. Like an optical scan ballot, the main body of a Scantegrity ballot contains, for each contest, a list of valid selections printed in a canonical order predetermined by polling place procedures (e.g., alphabetical, rotated across precincts, etc.). Next to each selection is a markable region, oval in shape, called a bubble. In

the bubbles associated with each selection, a short alphanumeric code is printed in invisible ink that is not human readable until marked by the voter. We refer to these codes as confirmation codes. In an election with n ballots and m candidates, there are $m \cdot n$ confirmation codes. Each confirmation code is pseudorandomly drawn from a suitable space: in the example, a confirmation code consists of two alphanumeric characters. The voter records their confirmation code(s) on the receipt and will use this information to conduct an important component of the audit.

Each Scantegrity ballot ensemble also contains three serial numbers. These are used to identify the ballot and to provide the voter with the ability to report the failure of aspects of the audit under certain scenarios. The ballot contains one of these numbers in the form of a machine-readable barcode that is not easily read or memorized by a human. Since many optical scanners use marksense technology, which only records whether a region is marked or unmarked, a suitable encoding of this number must be used (e.g., a data matrix for marksense scanners). For the i^{th} ballot, we refer to this serial number as α_i . The receipt portion contains two additional serial numbers, β_i and γ_i , which are printed in invisible ink and are individually detachable from the receipt.

10.3.2 *The Scantegrity Back-End*

The Scantegrity backend performs the verifiable tally of the voted confirmation codes. Intuitively, it treats each bubble as a distinct unit. Three values associated with each bubble—its confirmation code, whether it was voted or not, and its candidate—are stored across separate tables:

- Table **P** is not published but is used to print the ballots and to generate the other tables. It contains a row for each ballot and a column for the ballot ID, the confirmation codes for each candidate and two receipt IDs.
- Table **Q** is the same as **P** except the confirmation codes in each row have been independently shuffled. From **Q** alone, it cannot be established which candidate was voted for if a certain confirmation code is revealed. Each cell in **Q** is individually committed to.
- Table **S** contains a placeholder for each bubble in the election, ordered under each candidate. At the end of the election, a bubble may be marked, unmarked or selected for a print audit.
- Table **R** contains a row for every bubble in the election in random order. Each row in **R** connects to a unique confirmation code in **Q** and to a unique placeholder in **S**. Each **Q**-pointer and **S**-pointer is individually committed to.

All codes, serials, permutations, and commitment random factors are generated pseudorandomly. The combination of confirmation code and receipt ID is unique across the election. See Figure 10.3 for an example. Note that while this description is simpler and appears somewhat different from that in [145], the two descriptions are equivalent.

Ballot ID (α)	Adams	Burr	Jefferson	Pinckney	Receipt ID (β)	Receipt ID (γ)
01	EN	PL	JG	VE	8860	3478
02	IV	QU	SA	WC	3813	2448
03	AY	EY	XW	SI	2381	3492
04	BV	SY	ZK	HJ	1337	4592
05	SH	SR	OF	XJ	3492	3472

(a) P

α	κ_0	κ_1	κ_2	κ_3	β	γ
01	JG	EN	PL	VE	8860	3478
02	QU	IV	WC	SA	3813	2448
03	AY	EY	SI	XW	2381	3492
04	HJ	ZK	SY	BV	1337	4592
05	SR	XJ	SH	OF	3492	3472

(b) Q

	Adams	Burr	Jefferson	Pinckney
0				
1				
2				
3				
4				

(c) S

Q-Pointer	Marks	S-Pointer
(05, κ_1)		(2, Pinckney)
(03, κ_3)		(4, Jefferson)
(02, κ_1)		(4, Adams)
(01, κ_3)		(3, Pinckney)
(01, κ_2)		(4, Burr)
(05, κ_3)		(3, Jefferson)
(04, κ_2)		(3, Burr)
(04, κ_0)		(4, Pinckney)
(03, κ_0)		(0, Adams)
(04, κ_3)		(3, Adams)
(02, κ_3)		(1, Jefferson)
(03, κ_1)		(2, Burr)
(05, κ_0)		(1, Burr)
(01, κ_1)		(2, Adams)
(02, κ_2)		(0, Pinckney)
(04, κ_1)		(0, Jefferson)
(03, κ_2)		(1, Pinckney)
(05, κ_2)		(1, Adams)
(01, κ_0)		(2, Jefferson)
(02, κ_0)		(0, Burr)

(d) R

Figure 10.3: Tables P, Q, R and S as generated by the election officials before election day. Table P is kept private. The publicly published versions of tables Q, R and S contain commitments to the information shown above. For example, a vote for Jefferson on ballot 04 would reveal the confirmation code ZK. The corresponding row of table R points to position (04, κ_1) in table Q (containing ZK) and the position (0, Jefferson) in table S, which will be marked if the voter selects Jefferson on ballot 04.

10.3.3 Election Day

The voter marks the ballot as she would any optical scan ballot, using a pen with special ink provided in the booth. As the ink in the pen reacts with the ink printed on the ballot, the background of the bubble immediately turns dark, leaving a confirmation code visible in the foreground (see Figure 10.2). The relative darkness of any marked bubbles to unmarked ones will allow an optical scanner, employing dark mark logic, to register the bubble as marked. The foreground of the marked bubble is human-readable and a voter interested in participating in the election audit may record the code on the receipt portion of the ballot. Uninterested voters may disregard the codes.

When the voter has satisfactorily marked her ballot, it is returned to the poll worker. The poll worker places the main body of the ballot into the scanner, which records the ballot serial number and the marked choices. After a successful scan, the two serial numbers on the receipt, β and γ , are developed by the poll worker and the voter may leave with the receipt. It is expected that public interest groups will make available the possibility of creating a copy of receipts to alleviate the need for concerned but time-constrained voters to personally participate in auditing the election.

Voters may also opt to check that the ballots are printed correctly. In order to do this, the voter requests a ballot to be print audited, reveals all the confirmation codes on the ballot, and chooses one of the receipt serial numbers, β or γ , to retain with the codes. The election authority can keep the other code as a record of both: the ballot that was audited and the code which was kept by the voter.

If the voter makes an error in marking her ballot or wishes to register a protest vote through spoiling their ballot, the ballot is returned to the poll worker. Without seeing the contents of the ballot, the poll worker detaches the right side of the receipt, γ , from the ballot. The main body and left receipt, α and β , are shredded in view of the voter. The right receipt is retained by the poll worker and used in balancing the number of ballots issued with the number of ballots tallied, print audited or spoiled.

10.3.4 Posting and Tallying

After the close of polls, the election officials publish a list of all the voters who voted and the tally given by the underlying optical scan system. The electronic ballot images from the scanner, the list of audited ballots and list of spoiled ballots are entered into the trusted workstation by the trustees. Tables **Q**, **R** and **S** are regenerated and used to translate the votes into the corresponding confirmation codes (these codes were revealed by the voter but not recorded by the scanner, which only records voted bubbles). The commitments in table **Q** to the confirmation codes and serial numbers that were revealed during the election are opened; this demonstrates that the confirmation code was indeed one of those on the ballot. Corresponding marks are posted in **R** and **S** (exemplified in Figure 10.4). Anyone can now compute the tally from **S**. This tally is checked against the one reported by the optical scanners or manual recount. Opening such commitments preserves ballot secrecy: for each confirmation code opened in **Q**, it cannot be determined which candidate that code corresponds to.

Once the results are posted, a voter who has made a receipt of her confirmation codes can go to the election website and look up her voted ballot by either receipt serial number. She checks that, in the row of **Q** corresponding to her ballot, all and only her confirmation codes appear. Anyone can check that the opened commitments match the confirmation codes on the election website. If any of the confirmation codes from the voter's receipt do not appear posted in **Q**, the voter should file a dispute.

Q-Pointer	Marks	S-Pointer
(03, κ_3)	A	(4, Jefferson)
	✓	
(03, κ_0)	A	(0, Adams)
	✓	
(03, κ_1)	A	(2, Burr)
	✓	
(03, κ_2)	A	(1, Pinckney)
	✓	

(a) R

α	κ_0	κ_1	κ_2	κ_3	β	γ
01			PL		8860	3478
02				SA	3813	2448
03	AY	EY	SI	XW		3492
04		ZK			1337	4592
05			SH		3492	3472

(b) Q

	Adams	Burr	Jefferson	Pinckney
0	A		✓	
1	✓		✓	A
2		A		
3				
4	✓		A	

(c) S

Figure 10.4: Tables Q, R and S as published after the close of the election. For each revealed confirmation code in table Q, the row corresponding to that code in table R and the element corresponding to that code in table S have been flagged. For each row of table R, either the commitment to the Q-pointer or the commitment to the S-pointer has been opened and published. Note that the revealed confirmation codes in table Q (other than those for ballot 03, which is a ballot chosen for a print audit), the rows flagged in table R and the flags in table S are in one-to-one correspondence.

10.3.5 Auditing the Results

For those ballots that were chosen for a print audit, election officials open the commitments to all the information associated with the ballot (i.e., the confirmation codes in Q and the Q- and S-pointers in R) *except* the value of β or γ that the voter did not take with her after the audit. Anyone with access to the printed ballot with exposed confirmation codes can check that the revealed values are consistent with the committed and printed values.

To verify that the marks were added correctly to tables R and S, the election officials will be challenged to open either the Q-pointer or the S-pointer in table R; i.e., the first or third column (exemplified in Figure 10.5). If a mark was recorded for the wrong candidate, then either (i) the mark is not consistent with the corresponding mark in R, (ii) the marks in R and S match but the status of the corresponding confirmation code (revealed or unrevealed) in Q is inconsistent or (iii) the tables are consistent but the revealed code in Q does not match the receipt. The receipt check

Q-Pointer	Marks	S-Pointer	Q-Pointer	Marks	S-Pointer
(05, κ_1)					(2, Pinckney)
(03, κ_3)	A	(4, Jefferson)	(03, κ_3)	A	(4, Jefferson)
(02, κ_1)					(4, Adams)
(01, κ_3)					(3, Pinckney)
(01, κ_2)	✓			✓	(4, Burr)
(05, κ_3)					(3, Jefferson)
(04, κ_2)					(3, Burr)
(04, κ_0)					(4, Pinckney)
(03, κ_0)	A	(0, Adams)	(03, κ_0)	A	(0, Adams)
(04, κ_3)					(3, Adams)
(02, κ_3)	✓			✓	(1, Jefferson)
(03, κ_1)	A	(2, Burr)	(03, κ_1)	A	(2, Burr)
(05, κ_0)					(1, Burr)
(01, κ_1)					(2, Adams)
(02, κ_2)					(0, Pinckney)
(04, κ_1)	✓			✓	(0, Jefferson)
(03, κ_2)	A	(1, Pinckney)	(03, κ_2)	A	(1, Pinckney)
(05, κ_2)	✓			✓	(1, Adams)
(01, κ_0)					(2, Jefferson)
(02, κ_0)					(0, Burr)

(a) Heads

(b) Tails

Figure 10.5: Table R as published after the close of the post-election. According to a coin-flip, either the entire correspondence between Q and R is revealed, or the entire correspondence between R and S. In an actual election, many independently permuted copies of R will be generated and can be audited in this fashion.

probabilistically detects (iii), while the present challenge will detect either (i) or (ii). The challenge is generated with a publicly verifiable random number [157].

Any interested party can check that the commitments are correct: that each revealed Q-pointer in table R either connects a revealed code in table Q to a marked element in table R or connects a hidden code to an unmarked element, and that each revealed S-pointer in table R either connects a marked element in table R to a marked element in table S or connects an unmarked element to an unmarked element. Essentially, the audit checks that marks are mapped unchanged from table Q through table R to table S.

Currently the soundness of each check is only $\frac{1}{2}$. However by generating and using k independent R tables throughout the election, each with a unique pseudorandom mapping from Q to S, the soundness increases to $1 - \frac{1}{2^k}$ through auditing each independently.

10.3.6 Dispute Resolution

Dispute resolution—between a voter or observer claiming the evidence points to election fraud, and election officials and/or the voting system claiming otherwise—

is one of the key features of Scantegrity. If a dispute cannot be resolved, then a default position must be taken, which will generally either allow a corrupt election authority to get away with fraud or allow a malicious party to prompt a reelection (or at least cast doubt on the result) by filing spurious disputes.

To file a dispute, the voter must be on the registration list as being eligible to vote or as having cast a ballot, and only one dispute may be filed per voter. (If the voter wishes to allow a third party organization to check her receipt, she might sign over the right to her dispute.) Disputes are filed within a period of time. After the period closes, the trustees open **Q** completely, revealing all the confirmation codes and receipt numbers on all ballots (except spoiled ballots), and use this information to address the disputes.

The most likely dispute is that a voter's confirmation code does not match the one appearing on the website. In this case, the voter provides the ballot ID and a claim of what the correct confirmation code(s) should be. Disputes of this type could be (i) the result of the voter making a transcription error, (ii) a mischievous voter attempting to call into question the legitimacy of the election, (iii) an error by the scanner or (iv) evidence of fraudulent behavior.

With a transcription error, the claimant's code is likely close to the revealed code. Once **Q** is opened, the voter can verify that none of the other codes that appeared on the ballot match her code exactly. Since voters do not see the confirmation codes of the unrevealed candidates on the ballot when they file, the probability that a code that a voter claims to have received is actually one of the other codes committed to for her ballot is very small if the voter made a transcription error or is merely guessing. The election officials eliminate from consideration any disputes for which none of the opened codes match the claimed code. We consider the remaining disputes to be plausible discrepancies.

Plausible discrepancies could still be the product of cases (i) or (ii); however, the probability is small (and can be made smaller by increasing the length of the confirmation codes). They are more likely (iii) or (iv). The election officials should set up a statistical trigger, based on various election parameters, such that when a given number of plausible discrepancies is reached, the ballots are rescanned to rule out (iii) and then the election is considered to have strong evidence of fraud.⁴

Another type of dispute that a voter may have is that her ballot is improperly designated as voted, print audited or spoiled. A voter who cast her ballot knows both receipt serial numbers $\{\beta_i, \gamma_i\}$. If a ballot is represented as print audited or spoiled, either one or neither of these codes will be open in **Q**. Similarly the voter retains evidence that a ballot was print audited by knowing all the confirmation codes on the ballot. If a ballot is represented as voted or spoiled, either one or none of the confirmation codes will be open in **Q**. If the voter's claimed receipt code(s) or full

⁴If this were to happen, it would hopefully trigger litigation (as well as a criminal investigation). In Canada, any elector or candidate may apply for a contested election proceeding and if heard, the judge reserves the power to invalidate the election result (upon appeal to the Supreme Court of Canada).

set of confirmation codes match their committed values, the dispute is designated as a plausible discrepancy.

10.4 The 2009 Takoma Park Election

On November 3, 2009, we deployed Scantegrity in its first binding election. This marked the first time an E2E system was used for in-person voting in a governmental election.⁵ The election was run by the city of Takoma Park, Maryland, USA to elect city councilors and a mayor.

The system deployed in the Takoma Park elections is very similar to that described in Section 10.3, except that the cryptographic backend is a modified version of the Punchscan backend [465]. This backend was written in Java 1.5 using the BouncyCastle cryptography library.⁶ We opted to utilize the Punchscan backend because it was already implemented and tested in previous elections as described in Section 10.2.

10.4.1 Requirements

Voters voted for one of six wards, at a single polling station. The mayor race was common to all wards and the councillor race was specific to the ward. Voters were issued one of six ballot styles based on their ward. The city has approximately 17,000 residents and 10,934 voters were registered. The turnout was 1,728 voters.

Ballot. The ballot was required to be worded in both English and Spanish. A bilingual ballot was used instead of different sets of ballots for each language. For each ward, the ballot consisted of two contests. The joint mayor contest contained two named candidates and a write-in candidate. The ward councilor races varied between one and two named candidates and a write-in candidate for each. The Scantegrity system treats the option “write-in candidate” as a single candidate, with a corresponding confirmation code. Hence all write-in votes are grouped together in the initial tally. If the number of write-ins could be relevant to determining the winner (either of the election or the round in a multi-round scoring protocol), the actual candidates voted for are then examined. In this election, write-ins were not a factor in any of the contests.

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 preference. Prior to using Scantegrity, Takoma Park’s IRV bal-

⁵The Rijnland Internet Election System (RIES) deployed in the Netherlands in 2004 and 2006 [272, 314, 264] took a significant step toward the first real-world E2E election.

⁶<http://www.bouncycastle.org>

lots were arranged with a matrix for each contest, with candidates on the rows and preference order on the columns. We adapted this style and put confirmation codes in each cell.

Absentee Ballots. Takoma Park offers the option of casting a ballot by mail. It was too expensive to include a decoder pen with each absentee ballot, so we used Scantegrity ballots without the confirmation codes. This means absentee voters cannot verify the codes (although they can verify that their ballot was received, because a confirmation code would be posted on the website). 70 absentee ballots were cast.

Voter Registration. Registration was handled by the city of Takoma Park. Scantegrity did offer the option of provisional ballots, which would allow voters to vote but the ballots would be escrowed until the voter's eligibility was confirmed. If confirmed, the ballots are added to the tally (this itself may break the secrecy of the ballot if the set of approved ballots is small, however this is a known issue in any election system providing provisional voting).

Disclosed Source Code. 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 overall system relatively inexpensive.

Election Authority. It was decided that the election would have four trustees (the chair, vice chair and a member of the Board of Elections, and the city clerk) and any two were required to regenerate the election data.

10.4.2 Mock Election

In April 2009, we held a mock election to test and demonstrate feasibility of the Scantegrity system, as well as train the poll workers. The mock election was held during Takoma Park's annual Arbor Day celebration at the city hall. Volunteer voters, recruited from people attending the celebrations, voted on a mock ballot with questions relating to trees. Turnout was 95 voters. The election system used in the mock election was very close to the system described in Section 10.3. We uncovered several minor issues and one main issue: the average time-to-vote was unacceptably high at 8 minutes.

We identified two main impedances to voter flow. The first was filling out the ballot. The mock ballot had two IRV questions with four and five candidates, and two single-response questions. A fully marked ballot would consist of marking 11 bubbles and recording 11 confirmation codes. We reasoned that the election ballot would

be quicker to fill out because the ballot was much shorter (at most 6 bubbles/codes) and voters would likely have a preconceived notion of who to vote for.

The second impedance was the scanning station. During the mock election, voters would have their ballots scanned, be shown a representation on a monitor screen of what was scanned, and manually click to accept the interpretation and electronically cast the ballot. Then the receipt would be detached and the receipt codes would be revealed by the poll worker. Finally, the ballot was deposited in the ballot box. In the real election, we eliminated the verification step. Ballots were deposited into a slot, which resulted in the ballot being scanned and deposited into the ballot box. We removed the receipt from the ballot itself and instead provided separate verification cards. Voters who wanted to verify their ballot could copy the ballot serial number and confirmation codes onto the card.⁷ Additionally, we eliminated the receipt codes from the ballots. Instead of using them to provide dispute resolution, we relied on the assumption that the poll book would correctly maintain the number of cast, audited and spoiled ballots, as well as the original copies of any spoiled or audited ballots (for audited ballots, a photocopier was provided, and the voter and an election official signed the copies). Finally, we decided to use a second scanning station in the actual election. Through these modifications, the average time-to-vote in the actual election was reduced to under 3 minutes.

In the mock election, we used locks to prevent chain voting like we did in the Punchscan election however these were unpopular with the poll workers and eliminated. Last, we changed the confirmation codes from two letters (from a reduced set eliminating easily confused characters) to three decimal digits. In retrospect, we should have eliminated 0 and 8 from the set of digits.

10.4.3 The Election

The preparation of the election followed essentially the same steps as the Punchscan election, so we eliminate some details. The ballots were designed to closely resemble the ballots used in Takoma Park's previous municipal election. Each ballot contained some instructions for marking an IRV ballot and for marking and verifying a Scantegrity ballot.

Ballot Printing. We had the invisible ink, developing ink and decoder pens specially manufactured for our use. The ballots were printed on off-the-shelf CMYK inkjet printers. We kept the black ink in the CMYK printer cartridge and replaced the other 3 inks with the invisible ink, a dummy ink, and a fluorescent ink. The invisible ink is actually visible as a light yellow but turns dark gray when developed with the pen (see Figure 10.6). The dummy ink is the same yellow color but does not develop. We print it around the invisible ink to make the inks indistinguishable. Finally, we print

⁷This also makes it difficult to determine if a voter intends to verify their ballot or not. Recall in the Punchscan election, we had the difficulty of uninterested voters refusing their ballot-specific receipts or discarding them at the polling place.




Figure 10.6: Photograph of our invisible ink implementation showing marked oval, confirmation code and decoder pen (foreground).

a random scatter pattern over each bubble with the fluorescent ink to further obfuscate attempts at reading the undeveloped codes. Over many months, the invisible and dummy inks oxidize differently and the codes can become visible but this is not an issue if the ballots are printed shortly before the election.

Voter Education. Articles in the city newspaper before the election introduced the verification mechanisms provided by Scantegrity and explained how voters could check their confirmation codes online. This also appeared on the city’s election website. During the election, a local radio station covered the election and the use of Scantegrity, prompting many curious voters to attend the polling place. Finally, the ballot itself contained some instruction. Voter education was identified as a main area to improve in future elections.

Election Day Set-Up. The scanning station was the only Scantegrity-specific equipment. It consisted of a netbook, located in a lockbox, to control the scanner. Upon booting, the scanning software ran automatically off an attested, read-only SD card in the netbook. The election data was written onto USB sticks. Scantegrity is software-independent and auditors will detect if the scanner misbehaves in a way that affects the tally, however we strove to prevent this to avoid having to rerun aspects of the election. The scanner itself was embedded into the form factor of the scanning station. An uninterruptible power supply was also used in the case of a power failure.⁸

Voting. Polls were open from 7 am to 8 pm. Ballots in privacy sleeves were issued to registered voters. Voters would mark their ballots in the voting booth, optionally fill out a verification card, and be directed to the scanning station. At the scanner, voters simply dropped the privacy sleeve (with the ballot in it) into a slot. A rivet

⁸If use of the scanner was somehow lost, the Scantegrity election could still proceed without a problem. Ballots would be collected and scanned centrally later. The confirmation codes provide a cryptographic chain-of-custody over the ballots.

in the privacy sleeve prevented it from reaching the scanner, while the ballot itself slipped out and into the scanner's intake. The bottom of the scanner was inside the ballot "box" (which was actually a large bin on wheels making the scanner at waist height). The privacy sleeve could be removed from the slot and returned to the ballot issuing station by the poll workers.

Print Audits. During the day, an independent auditor from Electronic Privacy Information Center (EPIC), Lillie Coney, would visit the polling place at an unannounced time and select a random set of ballots. All the codes on these ballots were revealed and the auditor retained a copy of them for conducting a print audit against the cryptographic information.

Tallying. The trustees used the Scantegrity software to generate the tally at 10 pm. The Chair of the Board of Elections (and one of the trustees) 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. The tally was updated the next day with votes from approved provisional ballots, which did not change the outcome.

Hand Count and Certification. Members of the Scantegrity team and the trustees conducted a hand count of the ballots and certified the results. The hand count and the Scantegrity count differed slightly 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.

Dispute Resolution. On November 6, the dispute resolution period ended. 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 opening of commitments demonstrated that the digit was, indeed, "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.

Post-Election Audit. On November 6, trustees used the Scantegrity software to conduct the cut-and-choose audit (with a public challenge computed from stock market data). By November 9, two independent auditors retained by Takoma Park⁹ had verified the results with auditing scripts they independently wrote and published based on the Scantegrity protocol. The auditor from EPIC later verified that all the ballots she collected were printed correctly.

Receipt Check. The Scantegrity website recorded 81 unique ballot ID verifications, of which about 66 (almost 4% of the total votes) were performed before the dispute resolution deadline. We also were informed that at least a few voters had checked their codes through one of the independent auditor's websites, both of which had made the confirmation codes available. The number of voters who checked their ballots, while not large, was sufficient to have detected (with high probability) any errors or fraud large enough to have changed the election outcome.¹⁰

10.5 The Remotegrity Voting System

Following the 2009 election, Takoma Park requested the use of Scantegrity in the 2011 election. Note that absentee voters were not able to participate in election verification in 2009, and voters unable to independently mark paper ballots had done so with assistance. For the 2011 election, we provided functionality for remote E2E voting, as well as for E2E voting with an audio interface. In this section we describe Remotegrity, the remote E2E system used in 2011. In Section 10.6 we describe the audio voting system Audiotegrity.

Consider the remote voting problem for an election that uses Scantegrity in the polling place. One could mail Scantegrity ballots with the special pens to absentee voters, who could mark the ballots and mail them back, later checking online for their Scantegrity codes as an in-person voter would. Using this approach, however, the absentee voter would not have the in-person voter's ability to prove cheating. If an absentee voter claimed her vote was absent from the bulletin board, and cited a valid combination of code and ballot serial number, an independent observer would not know whether she was lying and had not mailed in her ballot, or the election administrator had dropped her vote. Thus, while we may use Scantegrity ballots for remote voting, we need additional steps that allow an independent observer to determine if a voter's vote had been correctly included in the tally or not. As it would be very cumbersome to conduct a voting protocol with multiple steps over postal mail,

⁹Ben Adida and Filip Zagórski—both cryptographic voting researchers; while Zagórski later drove the design, development and use of Remotegrity in the 2011 election, he was not part of the Scantegrity team for the 2009 election.

¹⁰We omit detailed calculations, noting these calculations become quite complex due to the use of IRV. Also, these calculations make certain assumptions (e.g., independence) that may or may not hold in practice.

we propose a system that is hybrid: the ballot and authentication credentials are delivered over postal mail and all consequent protocol steps are carried out over the Internet.

Remotegrity is a hybrid mail/internet extension to the Scantegrity in-person voting system, enabling secure, electronic return of vote-by-mail ballots [586]. It provides voters with the ability to detect unauthorized modifications to their cast ballots made by either malicious client software, or a corrupt election authority—two threats not studied in combination prior to Remotegrity. Not only can the voter detect such changes, they can prove it to a third party without giving up ballot secrecy; that is, Remotegrity possesses the dispute resolution property. Nothing about Remotegrity requires that it be used with Scantegrity. It can be used with any coded vote system.

10.5.1 Voter Experience

The remote voter receives an envelope containing a Scantegrity ballot and an authentication card. The authentication card bears a visible serial number and acknowledgment code. It also bears several authentication codes and one lock-in code under scratch-off surfaces, such as those used for lottery tickets (see Figure 10.7). At the appropriate step in the protocol, the voter scratches off the required surface to determine the code underneath and enters it online. In the event that the code is used in the protocol but the corresponding surface is not scratched off, one may assume that someone impersonating the voter entered the code.

In order to vote, the voter uses a computer to access the election website and enters her ballot serial number and authentication card serial number to ensure that neither has been used. If either has already been used, she is in possession of unused cards and hence this provides evidence of a problem.

If neither card has been used, she marks the Scantegrity ballot, revealing a confirmation code. She also scratches off an authentication code and enters both codes with both serial numbers at the election website.

A few hours later she accesses the election website again and checks if the election authority has responded with the acknowledgment code, indicating that the election authority has received a valid confirmation code for her ballot. The probability that the computer she voted from unilaterally changed her vote, by correctly guessing another valid confirmation code on her ballot, is small. Hence the election authority has, with high probability, received the correct confirmation code.

The voter now scratches off her lock-in code and enters it in, signalling to all observers, voters and the election authority that she is satisfied that her confirmation code is correctly recorded. The voting system signs the final record, indicating that the lock-in code is valid.

At the end of the election, all valid locked-in confirmation codes entered using Remotegrity are combined with all Scantegrity codes voted in person and the election outcome is tallied as with Scantegrity.

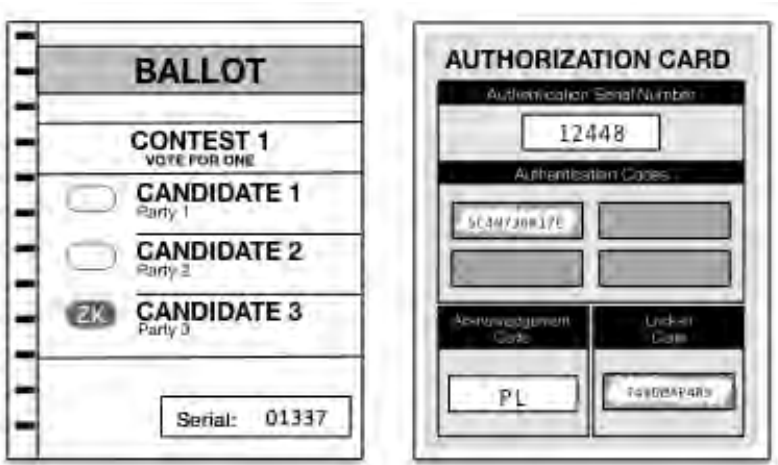


Figure 10.7: A Remotegrity package showing the ballot with one marked position and serial number (left) and the authentication card with one authentication code and the lock-in code scratched off. This figure represents a hypothetical final state of the cards after the voter has completed all steps.

10.5.2 Security Properties

In this section we provide informal arguments for the security properties of Remotegrity.

Privacy. The voter does not directly enter her vote into the computer she votes from. Hence her vote is private. As with mail-in voting, however, it is possible for a coercer to be present while she votes and to ensure she votes as instructed. Additionally, it is possible for the voter to produce her marked ballot and confirmation code to prove how she voted to a coercer who is not present while she votes.

Integrity: Impersonation. If someone has already voted using the voter’s ballot or authentication card, she detects it before she votes on the Scantegrity ballot or scratches-off any fields on the authentication card. She can hence produce the unused cards to prove that there is a problem. There should be mechanisms in place to allow her to vote, with another ballot and/or credentials if necessary.

Integrity: Malware on Voting Terminal—Vote Deletion or Unilateral Replacement. If the computer the voter votes from does not convey her confirmation number at all, the election authority will not respond with the acknowledgment code. If the computer attempts to unilaterally change her vote, it would need to change the code

to another one valid for her ballot. The probability of this is small. Hence, in this case too, the election authority will not receive a valid confirmation code for the ballot, and will not respond with the acknowledgment code. In both cases, the probability that the computer will be able to correctly guess the acknowledgment code and show it to the voter is very small. Hence, with high probability, the voter will know that her confirmation code has not been recorded and she can try to vote again from another computer.

If the voter does not see the acknowledgment code after several attempts, she has experienced a distributed denial of service attack either from all the computers she has attempted to vote from, or the election authority. In either case, she has in her possession her authentication card with an unscratched lock-in field. This proves that she has not completed the voting protocol and should be allowed to vote through another means, such as in person.

Integrity: Vote Replacement. If the election authority colludes with the terminal, it knows valid confirmation codes from the voter's ballot and can post an incorrect but valid confirmation code on the website, and return the valid acknowledgment code. The voter can respond by neither scratching off nor entering her lock-in code. If the election authority does lock in the incorrect confirmation code by entering her lock-in code, the voter can show her unscratched lock-in field to prove cheating and should be allowed to vote through another means, such as in person. If the election authority does not lock in the incorrect confirmation code, the voter cannot prove cheating, but her unscratched lock in field does prove that she has not completed the protocol and she should be allowed to vote through another means.

Summary. Thus, in the worst case, the voter experiences (a) a denial of service: either from all the computers she has access to or from the election authority which either does not respond with the acknowledgment code or enters an incorrect confirmation code or (b) an attempt to change her vote by the entry of an incorrect but valid confirmation code that is locked in. In (a) the voter can demonstrate that she has not completed the voting process and can vote in person, but cannot prove a denial of service. In (b) the voter can prove cheating and the changed vote is deleted and she is allowed to vote.

10.6 Audiotegrity

Audiotegrity is a voting system which enables voters with differing physical abilities to cast votes in the Scantegrity model, while preserving Scantegrity's dispute resolution and verification properties. It does so through the use of an audio interface provided by a computer, which reads options to a voter and records their votes without the intervention of a voting official. Votes are recorded on a ballot that is designed to look exactly like a marked Scantegrity ballot. Audiotegrity provides nearly

all of the dispute resolution and verification properties of Scantegrity while making it possible for a large subset of the population who previously could not independently cast a secret ballot to do so. Audioteegrity was used along with Scantegrity in the 2011 election in Takoma Park.

10.6.1 The Voting Process with Audioteegrity

Audioteegrity can be thought of as an extension to Scantegrity, and it shares many properties with the original system. Audioteegrity consists of a computer with monitor and large-format keypad with Braille stickers on the keys, a set of headphones with an attached microphone, and a printer. The interface supports multiple race formats, inclusive of all formats with which Scantegrity has been used.

The Audioteegrity voting station is slightly separated from the main voting floor by a privacy screen, as selections are displayed on a screen in large print for the benefit of those with visual disabilities who can read large type. Selections are made on a large format keypad in a process similar to that used with a touch-tone telephone menu. The printer is next to the voting station, and will later print both ballot and confirmation card, which are of different sizes so as to be distinguishable by touch. “Speak-in” votes are also allowed by the system, facilitated by a microphone on the voting headset. The voter can speak their desired vote for any race in the election, and their voice will be interpreted later by election officials, as with handwritten votes. These votes are identifiable by a pseudorandom number which is printed on the ballot in the space for the write-in candidate.

The system produces a marked Scantegrity ballot and confirmation card face down. The voter can choose whether to cast or audit their ballot before looking at the printed ballot (this prevents a coercion attack described in [341]). The voter can then check that the machine has marked their ballot correctly, and either cast it, or make a copy to take home for later ballot audit steps (which are identical to those for Scantegrity).

10.6.2 Dispute Resolution Properties

The main attack vector that Audioteegrity needs to protect against is sabotage of the computer (in this case used as an assistive marking system), partially described in [116]. A voter who does not examine the printed ballot (such as, for example, a voter with visual disabilities) cannot know if the computer has marked their ballot correctly. If, however, all voters use the same voting system, a computer marking the ballot incorrectly will be caught by voters who do examine the ballot. The voter cannot prove that the system marked the ballot incorrectly, so voters are allowed to spoil a ballot and vote again as many times as they desire. The number of spoiled ballots for a given voter is not made public, which eliminates a particular coercion threat. If we assume that the system cannot know if a voter will examine her ballot or

not, and if all voters who do examine it have correctly printed ballots, then those who do not check can be confident that, with high probability, their ballots were printed correctly.

Once the ballot is printed, and if it is printed correctly, it now functions as a Scantegrity ballot, with all of the Scantegrity dispute resolution and verification properties. If the system provides incorrect confirmation numbers for the voter's choice of candidate, it is caught during a Scantegrity print audit. The voter can prove that the system provided the wrong confirmation numbers because their vote is marked on the ballot, and Scantegrity has previously committed to the confirmation numbers for the serial number of that ballot. If the voting system posts a confirmation number to the bulletin board which is incorrect, this can also be resolved with Scantegrity.

Audiotegrity's dispute resolution properties are superior to those of fully electronic systems, and are only slightly weaker than those of Scantegrity. In particular, the Scantegrity voting system allows for a voter to be able to prove all attempts to cheat. Using Audiotegrity, a voter is able to prove all attempts to cheat except when the assistive device prints an incorrect vote.

10.7 Takoma Park Election, 2011

In late 2010 we were approached by the Takoma Park Board of Election regarding the possibility of becoming involved in their upcoming municipal election of 2011. We provided demonstrations of prototypes of Remotegrity and Audiotegrity in a couple of election board meetings in the first half of 2011. The city of Takoma Park held an open test of Audiotegrity and Remotegrity on June 8, 2011, in the Takoma Park Community Center. The test was publicized in the local news media and election officials sent announcements to various special-interest listservs. The test was not restricted to Takoma Park residents, and all who were interested were allowed to test the system. About 25–30 individuals tested the system and 17 individuals filled out a survey. The purpose of the survey was not usability research, but to obtain feedback on the systems in an informal manner, and to make potential users of the systems aware that Takoma Park might choose to use them in the election. Because we collected the data informally and interacted considerably with participants while they were testing the system, and the number of participants was very small, we do not present the data from our surveys. To obtain a qualitative, independent, albeit brief, assessment of the test, the reader may refer to a blog article [348].

We presented the results of the surveys to the Board of Elections (BOE) (many of whom had participated in the test) in the June meeting. The BOE outlined a number of concerns, centered around the usability and security of Remotegrity (because of the protocol's use of the Internet, and the problems Washington, DC, had had with an Internet voting trial—see Chapter 7, this volume). In the July meeting, the Board members communicated to us that they had confidence in the technology, but they were concerned about the procedures, which appeared ad hoc, about potential secu-

rity mishaps, and that the system had not been peer-reviewed. In this meeting, they communicated that they were leaning towards not using Remotegrity, but would go ahead with a mail-in Scantegrity ballot.

In the August Board meeting, we proposed (and they agreed) that the city provide voters with the option to use Remotegrity in addition to mailing back marked ballots. Only marked ballots would be counted, but voters using Remotegrity could test/audit the system, and, if they chose to lock in their vote, could communicate that the system was accurately recording their vote. Instructions for voting and auditing would be sent in separate envelopes in the same package, with appropriate marking, so as not to overwhelm voters not interested in the audit. Thus the system we finally used had some major differences with the protocol described in Section 10.5. Voters were not required to lock in (this means that, in practice, an election authority colluding with the voting terminal to change the vote could not be distinguished from the voter by a third party). Second, the Remotegrity system included ballots with visible codes (these ballots correspond to Scantegrity I [146]). This avoids the requirement of mailing invisible ink development pens. Third, voters needed to submit marked paper ballots by mail for votes to be counted; this eliminated any dependence on the Internet, but made it possible for the EA to ignore a mailed-in ballot. However, the Scantegrity codes of the votes were posted on the election website and voters could check if their votes made it in the count.

We made some changes to Audiotegrity as well, based on the criticisms and concerns of some participants: we provided variable speed and volume for the audio and obtained a professional recording for the real election. We also changed the instructions to make them more understandable. The Audiotegrity protocol used in Takoma Park was different from that described in Section 10.6 in a few aspects. No public declaration was required to cast or audit, and the ability to audit the ballot was not publicized widely. This was to simplify the process for the first use of the system. We chose to give audio confirmation codes to the voter before the printing began. Again, this was a consequence of the fact that we were not planning on many voter audits in this election and we wanted to provide voters with visual disability some of the information that sighted voters got. A better way to do this would be to provide digital media with confirmation codes on it.

Both systems were deployed on November 8, 2011.

The Remotegrity bulletin board contains 123 entries which correspond to 119 voters. Only 5 ballots were submitted online, and two of these were not counted as the corresponding paper ballots were not mailed in. While the number of voters who used the online system was small, full preparation and a complete implementation were required to deploy the system. Additionally, even voters who did not use the online system to enter or lock in their votes could check their confirmation codes online.

Audiotegrity was used to cast a few votes including by poll workers and auditors. Audits were made on the system by the election auditor, Neal McBurnett. This election marked one of the first times (if not the first) where the voting system design did

not prevent a voter with visual disability from independently casting an E2E ballot in a secret ballot precinct-based public election.

We are not able to provide information on how Audioteegrity votes were audited. We intentionally do not keep information on Audioteegrity ballot IDs after the election, in order to reduce the ability to distinguish between Audioteegrity and Scantegrity ballots.

Votes from both systems were combined with the Scantegrity votes cast in person, and a combined verifiable Scantegrity tally with the details of a tally audit was provided.

At the election certification meeting, Audioteegrity was called out as a valuable contribution by the chair of the board of elections and a council member.

10.8 Usability Studies

We surveyed voters and election judges during the 2009 and 2011 Takoma Park Municipal elections. These survey studies [135, 134] provide insight into the election experience with Scantegrity. They are among a handful that study voters using E2E voting systems, and are unique in being conducted for a binding election. While binding elections constrain research methodologies, these observational studies survey real voters under actual conditions that yield observations of the true voter and election judge experiences.

In the first election, we hypothesized that voters would have increased confidence but otherwise have no significant effect over traditional optical scan systems due to the cryptographic protocol. In the second election, we expected voters to show increased comfort with Scantegrity.

As hypothesized, we did not find evidence that the cryptographic protocol negatively affected voter perception in either election. The majority of voter reactions to Scantegrity were positive and we saw no consistent strong demographic effects between the two elections. Negative reactions correlated highly with problems in the voting experience, which pointed to problems with implementation (voting process issues, trouble reading the confirmation numbers, difficulty with filling out the ballot for instant runoff, cumbersome receipt process, etc.).

In this section we briefly summarize related work, our methodology and the highlights of these results. Readers interested in deeper analysis of the data should refer to the data reports available online [135, 134].

10.8.1 Related Studies

Sherman et al. [526, 525] report on focus groups and a survey similar to ours conducted at a preparatory mock election. In the mock election, Scantegrity team mem-

bers worked side-by-side with election officials to demonstrate capabilities of the system, and the surveys from voters during this election were positive. Carback et al. [133] focus primarily on the technical and administrative aspects of the deployment of Scantegrity at Takoma Park, including lessons learned and briefly touching on survey findings.

There are few user studies on E2E systems. Most studies are preliminary usability studies such as the student projects at UMBC (on Punchscan), MIT (on 3Ballot) [333], and the University of Surrey, England (on Prêt à Voter). These studies focus on user interface. We are not aware of other studies which focus on public acceptance, public reaction, and administrative challenges.

Acemyan et al. [40] attempted a comparative study of Scantegrity with Helios and Prêt à Voter and did not find fault with Scantegrity's verification mechanism, although they did point to problems with voter completion rates attributable to a poor optical scan implementation. Their system differs in critical ways to the one deployed in the municipal elections, requiring voters to scan the ballot and then separately drop it into a ballot box as well as providing mismatched voting instructions which likely confused subjects in their study. A detailed criticism of their study is provided by McBurnett et al. [386].

Using expert review, laboratory studies, and a field experiment with 1540 participants, Herrnson et al. [305] found that voting system interface and ballot styles had an impact on voter satisfaction, the need for help and voter's abilities to cast their ballots as intended. He also found that verification technologies typically had a negative impact on voter experience. Results of this experiment varied by voter demographics and voting experience.

Examining social issues, Newkirk [412] found that public opinion remained remarkably stable between 2004 and 2008. During that time, Direct recording electronic (DRE) systems were the top-rated systems for voter trust, followed closely by precinct count optical scan (pcos) systems. Voters rated vote-by-mail, central count optical scan and Internet voting less trustworthy.

Norris [420] describes a telephone survey of registered voters in Maryland in which voters provide strongly positive opinions about the usability and accuracy of touchscreen voting. Voters were also positive about the reliability, trustworthiness and count-accuracy of touchscreen machines, while admitting that the systems could be corrupted by malware.

10.8.2 Methodology

Our research protocols and questionnaires were approved by UMBC's Institutional Review Board, as required for experiments with human subjects. The study participants comprised election judges who administered the election and voters who voted in the election. Election judges operated the system, but select members of our re-

search team filled vendor roles, working separately from the survey and observation teams.

Our survey team surveyed voters in the exit polling area. In the first election, every voter was asked to fill out a survey. In the second, every 3rd voter was asked to minimize the effect of selection bias. After the election, judges were handed a questionnaire to fill out and return by mail.

Observers were stationed in each polling room. For each observation they filled out a standard form to indicate how long the voter took to vote and any incidents (e.g., problems with scanner, asking for help, etc.). Observers and surveyors were separated and did not switch roles throughout the day.

In 2011, our team was allowed to observe election day events, but we were not permitted to serve as election judges nor to interfere with the elections process. Additionally, only two representatives were permitted in each of the two voting areas at once. Observers were treated separately, and interacted minimally with our team throughout the day. Similarly, our surveyors stayed outside the polling place and did not interact with voters outside of when they exited the polling place.

Two members of our team acted as technical support when needed, fulfilling the role that a vendor would during election day.¹¹ Before the election the technical support team was directed to set up the scanning stations under supervision by an election judge. After the election they were asked to disconnect the scanning stations and to collate the memory sticks for tabulation by the election night tabulation software. The technical support team members did not interact with participants in the survey.

10.8.3 Known Limitations

Outside of avoiding selection bias in the second election, both studies share the same limitations. Our data cannot be construed as fully representative, but should be considered informative and indicative of how E2E verifiable audit mechanisms will be received in other jurisdictions at the municipal level.

Takoma Park does not have a history of election fraud, “dirty politics,” or similar concerns that would make its residents distrustful of the election process there. There was voter turnout even in uncontested wards, indicating a strong sense of civic duty in at least the subset of the population that showed up to vote. While the population of Takoma Park is very diverse, the exit survey data indicates that many of the people who came out to vote were highly educated and frequent computer users. Also, Takoma Park uses Instant Runoff Voting, which is not used in many jurisdictions in the United States.

The small number of questions on each ballot worked in the system’s favor. Increasing the number of questions on the ballot would likely decrease impressions of

¹¹In Maryland, technical support representatives from the election vendor are available to election judges at each polling site on election day.

usability, although this is, to varying degrees, the case for any system. It is unknown if a ballot as long as the typical Maryland state ballot could achieve this level of satisfaction from voters, and it would be difficult to change the Scantegrity system to support such a ballot.

Because these studies are observational, we are unable to address the question of the effects of the confirmation number receipts as well as we would have liked. Respondents to the questionnaire reported that the presence of a receipt increased their confidence in the results, but how many would have high confidence in the results if they had also used a system without a receipt? A comparative study which looks closely at this issue would be needed to address this question.

There were several miscellaneous technical problems throughout the voting day during both elections. A respondent who was the victim of, or witnessed, any of these issues was likely to have a negative response, and it is impossible to control for these types of issues in a real-world environment. Lastly, our survey sampled few voters with disabilities.

10.8.4 Voter Response

The primary feature of our surveys were Likert questions designed to capture voter satisfaction. Both surveys skewed toward positive voter satisfaction, although the 2009 survey had 2 low-response rate questions with poor wording. Results of selected comparable Likert questions are combined in Figure 10.8.

To understand if voter demographics affect voter experience we used ordinary least squares (OLS) regression of the combined dependent voter satisfaction variable against the demographics factors we collected. The cronbach's α of the satisfaction variable was .97 ($N = 142$) in 2009 and .84 ($N = 435$) in 2011.

The resulting dependent satisfaction variables had means of 5.69 (StdDev = 1.7, $N = 271$) and 5.84 (StdDev = 1.04, $N = 463$), respectively. Because the data was negatively skewed (-1.94 and -2.08) and had high kurtosis (5.92 and 8.9) we analyzed the cubes (x^3) of the values (skew = -.86 and -.44, kurtosis = 2.69 and 2.93).

The regression analysis did not agree on statistically significant effects between 2009 and 2011. Table 10.1 shows the statistically significant effects and their magnitudes each year.¹² Refer to the individual publications for the full tables and additional discussion.

10.8.5 Observational Results

In the 2011 Takoma Park election our passive observers tracked 314 voters as they carried out the voting process. One voter did not complete the voting process. Twenty-six voters appeared confused, but did not ask for help.¹³

¹²Please note the regression models are slightly different between 2009 and 2011.

¹³We did not have the resources in the 2009 election to collect this observational data.

Voter Response to Likert Questions in 2009 and 2011

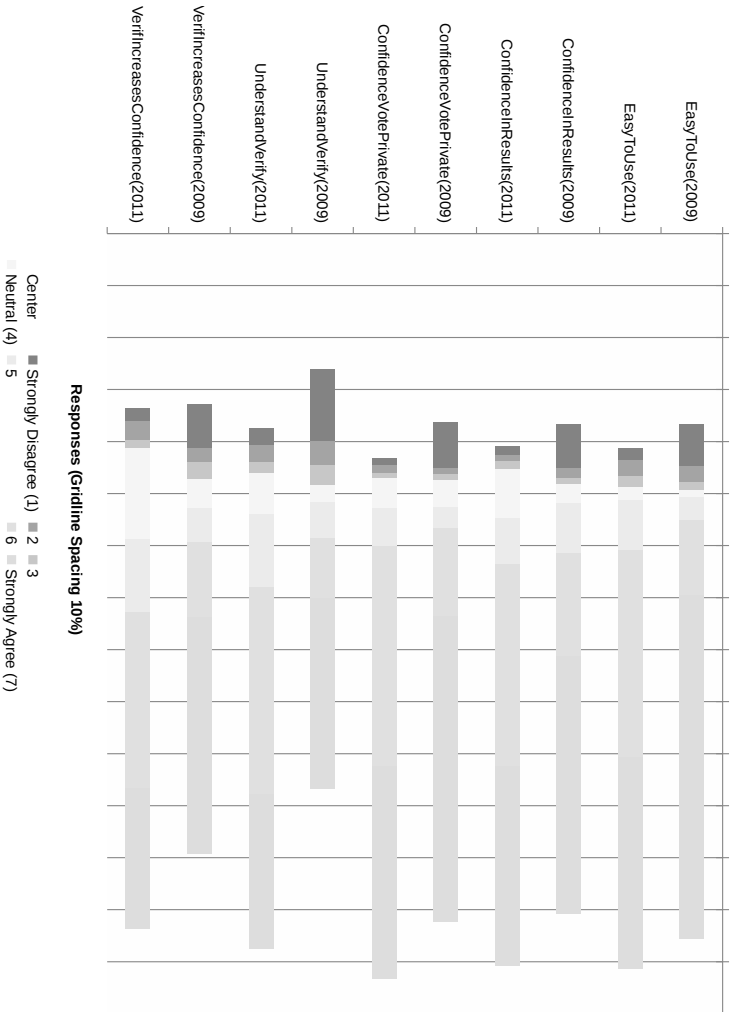


Figure 10.8: Voters were more positive towards Scantegrity in 2011 compared to 2009. This figure is a horizontal bar chart showing distributions around the neutral position of the Likert questions. The width of each bar represents the total number of respondents for that question, and each bar is divided into subsections whose width represents the respondents in that category.

Table 10.1: Significant factors affecting voter satisfaction with Scantegrity in 2009 and 2011. Entries are OLS coefficients, standard errors, t-values and probabilities. All entries are one-tailed. No factors showed as significant between both 2009 and 2011, and 2011 data featured significant factors dominated by voter experiences (e.g., problems voting, voting in 2009, etc.).

-	2009				2011			
Category	Est	StdErr	t	P> t	Est	StdErr	t	P> t
Female	33.53	15.76	2.13	.02	11.81	8.06	1.47	.14
Income	-6.89	5.10	-1.35	.09	4.27	4.30	0.99	.32
T. Screen	-47.52	25.96	-1.83	.03	-1.78	14.52	-0.12	.90
P. Card	48.32	18.43	2.62	.01	8.31	8.19	1.014	.31
No-Eng.	-39.99	33.04	-1.21	0.11	86.60	40.64	2.13	.03
Eyesight	NA	NA	NA	NA	-33.58	18.70	-1.80	.07
Op. Scan	NA	NA	NA	NA	21.23	8.64	2.46	.02
Und. Ver.	NA	NA	NA	NA	49.99	10.43	4.79	.00
Vote 09	NA	NA	NA	NA	-19.98	9.74	-2.05	.04
Ver. 09	NA	NA	NA	NA	23.41	13.74	1.70	.09
New Vot.	NA	NA	NA	NA	77.84	41.01	1.90	.06
Prob. Vot.	NA	NA	NA	NA	-68.97	14.02	-4.92	.00

Voting times ranged from 43 seconds to 19 minutes and 36 seconds (the second longest time was 9 minutes and 54 seconds, the skew was highly positive at 2.81 for timing data), with a mean of 3 minutes 27 seconds, a median of 3 minutes, and a standard deviation of 1 minute 56 seconds. Most of the time was spent marking the ballot. The average time to vote was slightly longer than during the November 2009 election, when the median voting time was 2 minutes 30 seconds. None of the voters performed a print audit.

The number of times a voter spoiled a ballot, asked for help or a judge intervened during the voting process is shown in Table 10.2. Note that, per process, the judge is supposed to intervene to explain the voting process. This happened in most but not all cases.

The verification website recorded 81 unique IP addresses in 2009 and 119 in 2011. Most addresses submitted multiple requests to the submission form, but our system did not record if different ballot confirmation numbers were checked between submissions. If each address only checked 1 ballot, this represents verification rates of 4.7% (1723 voters) and 6.1% (1951 voters), respectively.

Table 10.2: Frequency of different events observed in the polling place by four different observers (O1 to O4) across 314 observations at random time intervals during the election.

Event	O1	O2	O3	O4
Voter spoiled a ballot	4	0	0	0
Voter asked for assistance	30	8	1	0
Judge intervened with voter	156	136	16	2

10.8.6 Election Judge Response

In both elections, the judges identified as their main challenge the need to provide more instructions to voters, in comparison with non-Scantegrity elections. Specifically, voters needed more instructions on how to fill out the ballot and how to use the verification card. Many voters also needed instructions on how to mark the ballot using IRV.

Overall, election judges felt that the revised and increased educational efforts concerning verification paid off. By contrast, many voters in 2009 did not understand that they could verify their votes and that, to do so, they needed to write down the exposed code numbers. The judges noted that some people who voted in 2009 already knew how to verify from that election. The process worked more smoothly than in 2009.

Election judge recommendations focused on improving instructional materials, separating voter types to better serve their needs (e.g., special attention for voters with disabilities), and improvements to the voting equipment (e.g., automated receipts, improved scanners, etc.). Judges appeared to see the value of the verification option and agreed with voters that it improved confidence in election results.

10.8.7 Study Findings

Scantegrity is not too complex to use and administer in the context of the elections at Takoma Park. The verification system’s existence does not appear to impact voter experience negatively. Voter’s experience showed high levels of confidence in Scantegrity. While it is non-trivial to address if certain voting populations will be disadvantaged by this system, we do not have evidence that there are disadvantaged groups over a traditional optical scan system.

Voters appeared to accept the system even if they did not understand it. This study shows high levels of support for the verification receipt, even when voters indicated they did not understand the cryptographic mechanisms behind it. While voters appreciated the extra security as the presence of the verification increased their confidence, this contrasts with the number of people who reported intent to verify and more so with the 3–5% of voters who actually checked their ballot data. It

appears that, even though voters appreciated the technology, they did not necessarily care to use it.

In both elections, there were no negative comments about the verification option itself. We also had enough verifiers to ensure election security. Many voters expressed deep gratitude for the system and judges reported less confusion about verification in the second election.

The general process of creating a receipt remains a problem absent automation, and it does impact voter experience. This election experience inspired us to develop an automatic receipt printer [527]. Judges and observers noted increased time to vote due to the verification mechanism. Many voters expressed frustration with a number of practical problems (codes being hard to read, not writing down all necessary information, problems understanding how to use the pen, etc.), but note, however, that the overwhelming majority of surveyed respondents agreed that the system was easy to use. Most of these issues can be fixed or streamlined with better equipment. For example, individuals with vision disabilities reported problems despite our efforts to improve code legibility.

We paid attention to the ostensibly small details. Instructions for IRV and how to record a receipt were checked carefully for clarity. Takoma Park used an official staff member to create translations of these instructions. The scanner recognized ballots being placed into the system with a beep. More work can still be done to modify the voting process and technology to improve the voter and election judge experiences.

We believe that the development of practical, easy-to-use and well-accepted E2E verifiable audit technology will greatly improve election results assurance, voter confidence, election transparency, and independent election verifiability, and provide solutions to chain of custody issues. Takoma Park's diverse voting population offers insight into the acceptance of such systems.

The findings of this study increase the knowledge and understanding of how to successfully implement E2E verifiable post-election audit technology and allows other jurisdictions to use our report and project data to implement similar systems in their districts. It indicates that voters are positive about the system, that they value the security provided, that the extra work of optionally noting down confirmation codes does not significantly impact the voter experience negatively and that they accept it in spite of not understanding its inner workings completely. Election judge responses indicated that the system was not too hard to administer, and we did not observe significant relationships between demographic data and responses while voters were highly satisfied with the system.

Most voters and judges were positive about the system, and very few people expressed any doubts about the post-election audit mechanism. Even voters who did not plan to verify had few negative reactions, and most either did not trust computers at all or did not understand the audit system.

For future work, a clear measure of the confidence increase (or decrease) the receipt provides is necessary. A comparison study between Scantegrity and a com-

mercial optical scanning system is an obvious next step. Another area to explore is whether enough voters will use the receipts.

More work needs to be done for voter education. Despite significant efforts, educating voters before the election turns out to be a significant challenge, and we believe that most voters learned about it through the video while waiting in line or when interacting with the election judges. Education ahead of the election requires getting voters' attention on a matter that they think they already understand ("I show up and check a box"). It is likely that there will continue to be a need for informing voters close to the moment of voting (waiting in line, as they are handed their ballot and verification card). For populations that need more time to assimilate information, such as senior citizens and non-native English speakers, what can we do to educate them ahead of time so that when it comes time to vote, they can make an informed decision rather than simply skipping it because it just seems too confusing?

If vote verification becomes an expected part of voting systems, it will be important to ensure that all voters who want to verify their vote have the means to do so. For voters who don't have access to a computer and want to verify their vote later on, is there a way to help them get access to the means to verify their vote?

These conclusions should of course be considered in view of the setting: a municipality without a history of election fraud, civil contests rather than vitriolic or heavily partisan contests, a voter turnout comprised primarily of people who are highly educated and comfortable with computers. Responses might have been different in a situation that had more voters who were distrustful of their election process, less comfortable with computers or lacking computer access, or voting in a hotly contested and/or highly charged election.

10.9 Current Status of the Project

As an academic project, our main interest was in proving the feasibility of running E2E elections in the real world. The logistics and expenses of continuing to run elections, however, became difficult as all of the students on the project eventually graduated and began working in new areas in industry and academia. We decided to suspend active development on the Scantegrity project after the 2011 election at Takoma Park, leaving on good terms with the city.

10.10 Conclusions

In 2005, just after the introduction of the first human verifiable E2E voting systems, a list of recommendations were proposed by Karlof, Sastry and Wagner [340] for realizing the full deployment potential of cryptographic voting techniques. To summarize, they recommended the following:

- **Certification:** a new framework for evaluating E2E systems and criteria for their certification by an independent body,
- **Usability evaluation:** review of the systems by usability experts and running trials with questionnaires,
- **Recoverability:** the ability to fall back to a reliable underlying system, such as hand-countable paper ballots
- **Transparency:** full disclosure of source code and documentation of the systems.

As of 2005, these were all open problems. Through our work, and the work of others in the community, progress has been made on all four.

Certification

In the United States, the Election Assistance Commission (EAC), with guidance from the National Institute of Standards and Technology (NIST), provides vote system guidelines (VVSG). These guidelines now include E2E voting systems. In 2010, NIST held a workshop on E2E voting systems and we reported our experiences with the mock election at Takoma Park (the workshop preceded the actual election).

Usability

While independent usability evaluation of all E2E systems is significantly deficient from the literature, the data we collected after the Takoma Park mock and municipal elections currently represents the largest dataset of voter and poll worker reactions from an E2E system.

Recoverability

The introduction of Scantegrity, which simply added E2E verification to an otherwise normal optical scan system, represents a major step forward for recoverable E2E systems. Most paper-based E2E systems are like Punchscan, without a recoverable paper audit trail. This has benefits from a privacy perspective—the E2E system is the only interface to the real tally (even the scanner does not know how you voted). But it also creates a concentrated point of failure. We believe systems like Scantegrity are a first step toward verifiability, and increased voter privacy should follow after time.

Transparency

Both Punchscan and Scantegrity are fully open-source projects (although they license patented technology; see disclosure below). In fact, to our knowledge, the Takoma Park election also has the distinction of being the first open-source government election held in the United States.

The successful E2E voting pilot at Takoma Park demonstrates that voters and election officials can use advanced cryptographic techniques within an election, and, with reference to our polling data, be satisfied with its usability. We believe the Scantegrity system and the 2009 and 2011 elections demonstrate a significant advancement in the technical maturity of E2E voting. We can now say it is ready for real binding governmental elections. The remaining hurdle for the acceptance of cryptography in elections is when voters stop asking why we are using it, and start demanding why we are not using it.¹⁴

Acknowledgments

Ben Adida and Filip Zagórski performed pre- and post-election audits for the 2009 election. Lillie Coney audited paper ballots. Robert Araujo and Neil McBurnett performed pre- and post-election audits for the 2011 election. Neil McBurnett did in-person ballot-print audits and also audited the Audiotegrity audio interface. Noel Runyan provided expertise for the design of Audiotegrity. The Board of Elections of the City of Takoma Park and the City Clerk displayed infinite patience and shared with us their considerable experience with elections.

We were not compensated by the City of Takoma Park for either election. Clark was/is supported by NSERC. Vora was/is supported in part by NSF awards 0505510, 0831149, 0937267, 1137973 and 1421373, and a grant from the NSA. Zagórski is partially supported by Polish National Science Centre contract number DEC-2013/09/D/ST6/03927.

Sherman was supported in part by the Department of Defense under IASP grants H98230-09-1-0404 and H98230-10-1-0359 and by the National Science Foundation under SFS grant 1241576. Survey work at the 2011 Takoma Park election was supported in part by the U.S. Election Assistance Commission under grant EAC110150I.

¹⁴Adapted from Stu Feldman's roadmap for technical maturity (as quoted in [250]): (i) You have a good idea. (ii) You can make your idea work. (iii) You can convince a gullible friend to try it. (iv) People stop asking why you are doing it. (v) Other people are asked why they are not doing it.