vVote: a Verifiable Voting System

CHRIS CULNANE, University of Surrey, UK
PETER Y. A. RYAN, University of Luxembourg, Luxembourg
STEVE SCHNEIDER, University of Surrey, UK
VANESSA TEAGUE, University of Melbourne, Australia

The Prêt à Voter cryptographic voting system was designed to be flexible and to offer voters a familiar and easy voting experience. In this paper we present our development of the Prêt à Voter design to a practical implementation used in a real State election in November 2014, called vVote. As well as solving practical engineering challenges, we have also had to tailor the system to the idiosyncrasies of elections in the Australian state of Victoria, and the requirements of the Victorian Electoral Commission. This paper includes general background, user experience and details of the cryptographic protocols and human processes. We explain the problems, present solutions, then analyse their security properties and explain how they tie in to other design decisions.

Categories and Subject Descriptors: H.2.7 [**Database Administration**]: Security, Integrity and Protection; K.6.5 [**Management of Computing and Information Systems**]: Security and Protection

General Terms: Security, Privacy, Verifiability

Additional Key Words and Phrases: Verifiable electronic voting systems, voting protocols

ACM Reference Format:
Chris Culnane, Peter Y. A. Ryan, Steve Schneider and Vanessa Teague, 2015. vVote, a Verifiable Voting System. *ACM Trans. Info. Syst. Sec.* V, N, Article A (January YYYY), 30 pages.
DOI:http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

This paper details a design for end-to-end verifiable voting in the Australian state of Victoria, based on the Prêt à Voter end-to-end verifiable voting system [Ryan et al. 2009]. The system ran successfully in the state election in Victoria (Australia) in November 2014, taking a total of 1121 votes from supervised polling places inside Victoria and overseas.

The proposed protocol is end-to-end verifiable, meaning that there are no human or electronic components which must be trusted for guaranteeing the integrity of the votes (although vision impaired voters must assume that at least one device reads accurately to them). There are probabilistic assumptions about the number of voters who confirm correct printing of some Prêt à Voter ballots, the number who check that their printout matches their intended vote, and the number who check that their receipt appears on the Web Bulletin Board (WBB). It also provides voters with evidence of malfeasance, assuming that they check the signature on their receipt before they leave

Authors' addresses: C. Culnane and S. Schneider, Department of Computing University of Surrey, P.Y.A. Ryan, Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg, V. Teague, Department of Computing and Information Systems, University of Melbourne.

DOI:http://dx.doi.org/10.1145/0000000.0000000

the polling station. Since this is a polling-station scheme, we do not address eligibility verifiability. Prevention of ballot stuffing is by existing procedural mechanisms.

Running an end-to-end verifiable protocol on a subset of the votes does not make the election result end-to-end verifiable, when that result depends on traditional paper processes that are not verifiable. Nevertheless we believe there are significant advantages to using an end-to-end verifiable electronic voting system for part of an election, and we consider that explanation to be a main contribution of this paper.

The team involved in developing the vVote design described in this report were: Craig Burton, Chris Culnane, James Heather, Rui Joaquim, Peter Y. A. Ryan, Steve Schneider and Vanessa Teague.

1.1. Contributions of this work

End-to-end verifiable election protocols are well studied in the academic literature, but (with the notable exception of the Scantegrity II project in Takoma Park MD) have not previously been deployed in binding government elections. This paper contributes new protocols for addressing issues that arise in practice but have not been adequately considered in the literature, and new insights into the important difference between practical requirements and academic security goals. Our main contributions are:

(1) A version of Prêt à Voter usable enough for real people, even for the very complex ballots used in Victoria, with some practical evidence about its use in a real election.
(2) Scalable cryptographic protocols that are fast enough for long preferential ballots (though the details are in [Culnane et al. 2013]).
(3) A clear account of what is achieved by running an end-to-end verifiable system as part of an electoral process that also includes a traditional paper-based system for other votes. The paper elements mean that the whole electoral process is not end-to-end verifiable, but end-to-end verifiability of the subset improves the weakest links in the paper system and hence the security of the overall system. This is substantially better than substituting an unverifiable electronic system in the same place.
(4) An informative account of the challenges of implementing and deploying a verifiable system and some lessons about the distinction between theory and practice.
(5) A comprehensive security analysis of a deployed system, including those attacks that remain and are tolerable, and what trust assumptions remain in practice.
(6) Detailed procedures for achieving some accountability, with a clear description of how the electoral administration should respond to apparent failures and how, and whether, voters who detect errors can demonstrate them.

1.2. End-to-end verifiability

End-to-end verifiability usually consists of three pieces of evidence:

Cast-as-intended verification. Each voter gets evidence that their vote is cast as they intended;

Recorded-as-cast verification. Each voter gets evidence that their vote is included unaltered in the tally;

Universally verifiable tallying. Everyone can check that the list of (encrypted) recorded votes produces the announced election outcome.

This project does not currently achieve verifiability all the way to the announcement of the election result, because it runs alongside an existing paper-based system that relies on scrutineers to check that the cast votes are included unaltered in the final count. See Section 1.3 for more details. In summary, the vVote system provides

— cast-as-intended verification,
— recorded-as-cast verification and

— an output list of decrypted recorded votes, with a universally verifiable proof of the decryption.

An important practical advantage of an end-to-end verifiable election scheme, compared to simpler methods of electronically assisted voting, is that it provides for electronic transfer of ballot information from distant supervised locations, supported by verifiable evidence of correctness. This is particularly important for distant polling places (*e.g.* overseas) and for allowing any voter to vote at any polling place. Since this project commenced, a problem in the transport of West Australian Senate ballots in the 2013 federal election has focused national attention on the security of processes for transporting paper ballots.

1.3. Challenges of combining end-to-end verifiability with traditional Victorian paper voting

A large part of the challenge arises from the special requirements of Victorian parliamentary elections. Victoria, like many other Australian states, runs simultaneous elections for two houses of parliament, the Legislative Assembly (LA) and the Legislative Council (LC), both of which use ranked-choice voting. Each LA representative is elected by Instant Run-off Voting (IRV) with compulsory complete preference listing, with rarely more than 10 candidates. Members of the Legislative Council (LC) are elected in 5-member electorates using the Single Transferable Vote (STV) algorithm.[1]. Voters typically choose from among about 30 candidates—they rank at least 5, and up to all candidates in their order of preference. Because LC voting is complex, voters are offered a shorthand called "Above the line" (ATL) voting, which allows them to select only their favourite political group (usually a party). Each polling place must accept votes for any race, thus serving residents of any district in the state.

This system was not responsible for all of the votes cast in the 2014 state election, so it had to combine with existing procedures for casting and counting ordinary paper ballots. For LC-ATL votes the combination is a straightforward sum. For LA votes it requires careful scrutineering of paper printouts, and for LC-BTL (Below The Line) votes it is complicated. Even those LC-BTL votes cast on paper are tallied electronically—in the existing system they are manually entered first. The authorities then make complete vote data available to allow observers to check the count.[2] The scrutineers who observe the paper count have to check that the publicly verifiable output from vVote matches the votes that are added to the paper count.

This is why the system does not achieve a complete end-to-end verifiable election outcome. The proof that this system produces would be sufficient for end-to-end verifiability if it carried all votes in the election, but it is not possible to do STV tallying (whether verifiable or not) on a subset of votes.

Our system provides privacy and receipt-freeness under reasonable assumptions about the correct randomised generation and careful deletion of secret data, and of course assuming a secure mixnet and that a threshold of decryption key sharers do not collude. It depends on both the electronic ballot marker and the printer protecting their secret data. It does not defend against ballot signature attacks [Di Cosmo 2007] (often called "Italian Attacks,") or other subtle coercion issues, but neither does the current paper-based system. Our system also reveals whether a person voted ATL or BTL. A precise statement about privacy, its assumptions and limitations, is in Section 6.2.

Another challenge is producing an accessible solution for voters who cannot fill out a paper ballot unassisted. This is a primary justification for the project, but producing

[1] The workings of the counting algorithm are outside the scope of this project, but a good introduction is available on the webpage of the Electoral Council of Australia and New Zealand: `http://www.eca.gov.au/systems/proportional/`

[2] These procedures are also under review and improvement, but are out of the scope of this paper.

a truly verifiable solution for such voters is extremely difficult, because many of them cannot perform the crucial check that the printout matches their intention (though see [Chaum et al. 2009] for a verifiable and accessible protocol). We provide a way for them to use any other machine in the polling place to do the check, in which case the cast-as-intended property depends upon at least one of the machines in the polling station not colluding with the others to manipulate the vote.

1.4. Related Work

In the USA, permanent paper records such as Voter Verified Paper Audit Trails (VVPAT) or opscans are a common means of achieving software independence [Rivest 2008]. However, this does not solve the problem of secure custody and transport of the paper trail. Furthermore, performing rigorous risk-limiting audits seems intractable for IRV [Magrino et al. 2011], let alone for 30-candidate STV.

The most closely related project is the groundbreaking use of Scantegrity II in binding local government elections in Takoma Park, MD [Carback et al. 2010]. Our project has very similar privacy and verifiability properties. However, both the overall election size and the complexity of each ballot are greater for our system. Although the Scantegrity II scheme appears to have been highly successful in the context of the Takoma Park elections, Prêt à Voter is more appropriate for our application. Scantegrity II is inherently for single-candidate selections. It has been adapted to IRV in Takoma Park by running a separate single-candidate election for each preference, but would be difficult to adapt to 30-candidate preference lists. Even with computer assistance, a 30 by 30 grid of invisible ink bubbles seems too complicated for most voters.

The STAR-Vote project proposed for Travis County, TX [Bell et al. 2013] represents an interesting combination of end-to-end verification techniques and risk limiting audits. STAR-Vote retains a human-readable paper record for auditing purposes alongside the end-to-end verifiable cryptographic data. Cast-as-intended verification of the end-to-end verifiable part is achieved by a novel interpretation of Benaloh's simple challenge process [Benaloh 2006], in which voters can choose either to cast their ballot into a special ballot box or to spoil it and start again. We hope our observations might be helpful in the final stages of the STAR-Vote design process.

1.5. Prior work and paper overview

A previous paper [Burton et al. 2012a] gave an overview of this project, including the context of Victorian voting and some ideas on implementing the protocol. A followup version [Burton et al. 2012b] gave more details and some preliminary security analysis. The print on demand protocol was presented in [Culnane et al. 2013]—we omit the cryptographic details here. This paper contains a systems-level view of the whole protocol, including how the cryptographic protocol interfaces with the human procedures to be followed in the polling place and at the electoral commission. Our aim is for a comprehensive analysis of the protocol's security, including the assumptions on which privacy depends, a precise explanation of the kind of verifiability achieved, and a clear statement of the issues that remain. State-of-the-art formal and computational security analysis is not yet mature enough to apply to a system the size and complexity of vVote, so we consider the systematic analysis carried out in this paper as the best approach currently available. A complete paper including all the details from all project publications is available on ArXiv [Culnane et al. 2014].

Many of the system's security properties depend on proper procedures in the polling place—these are detailed in Section 3. Each system component is described in Section 4. Section 5 contains mechanisms for achieving robustness in the presence of certain failures. A comprehensive and rigorous threat analysis is in Section 6.

2. DESIGN OVERVIEW

2.1. Prêt à Voter overview

Prêt à Voter uses a ballot form that is printed before voting, with a list of candidates printed in a random order, and an encrypted version of the same list. Voters select or number the candidates by filling in boxes adjacent to the candidate names (in the Victorian protocol, they have computerised assistance to print out a separate list of marked boxes). They keep the list of marked boxes and the encrypted candidate list, and shred the human-readable candidate list. The two main properties are privacy and end-to-end verifiability. End-to-end verifiability is achieved in Prêt à Voter as follows:

Cast-as-intended verification.

[1. ballot printing confirmation.] Each voter has the opportunity to confirm that the printed candidate lists on some ballots match their encrypted version.

[2. preference printing confirmation.] Each voter checks that their own preferences are correctly written (or writes them in the case of standard Prêt à Voter) on the half of the ballot that they retain.

Recorded-as-cast verification. Each voter gets the opportunity to check that their (encrypted) ballot appears in a public list of recorded votes,

Universally verifiable tallying. Everyone can check the public electronic proof that the list of (encrypted) recorded votes produces the announced (decrypted) output votes.

2.2. System Component Overview

The practical implementation required a number of important details not previously specified in theoretical work. The main idea is that the print-on-demand printer produces a human-readable candidate list, together with a serial number linking it to its encrypted representation on the public Web Bulletin Board (WBB). Then the voter uses an Electronic Ballot Marker (EBM) to print out a list of preference numbers that align correctly with their randomised candidate list. A picture of the ballot is in Figure 3 (p. 10). Details of the voting procedures are in Section 3 and of the system components are in Section 4.

The system has the following main components, pictured in Figure 1. The details of how these components work are deferred until Section 4.

Public Web Bulletin Board (Public WBB). an authenticated public broadcast channel with memory.

Private Web Bulletin Board (Private WBB). a robust secure database which receives messages, performs basic validity checks, and returns a signature. Validly signed messages are guaranteed, under certain assumptions, to appear subsequently on the Public WBB.

Print-on-demand printer. a combination of a computer and printer which generates Prêt à Voter ballots in advance of the election, then prints the candidate list on demand.

Randomness Generation Service. a collection of servers that produce randomness for the print on demand process.

Electronic Ballot Marker (EBM). a computer that assists the user in filling in a Prêt à Voter ballot.

Cancel Station. a supervised interface for cancelling a vote that has not been properly submitted or has not received a valid Private WBB signature.

Cancel Authority. a central server responsible for authorising cancellations and tracking the number of cancellations from each Cancel Station.

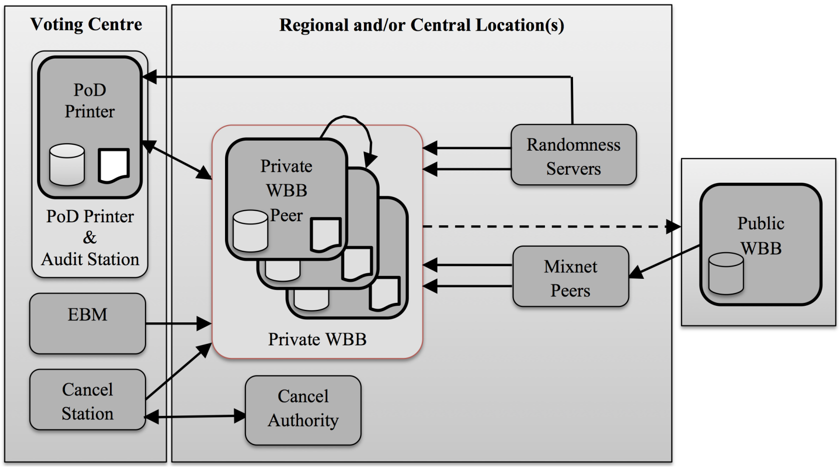

Fig. 1. Architecture of the vVote system

Mixnet. a set of (preferably independently managed and hosted) Mix servers that produces a noninteractive, universally verifiable proof of a shuffle and decryption (of encrypted votes) and posts it to the Public WBB.

Election key sharers. authorities who share the decryption key for the votes.

2.3. Security Assumptions

The intention is to provide a proof of integrity with clearly stated trust assumptions and without any trusted individual people or electronic components, while preserving reasonable privacy. We have several different kinds of security assumptions, which apply at different points:

Computational assumption. A computational problem generally believed to be hard. For example, privacy depends on the semantic security of El Gamal encryption, which relies on the hardness of the Decision Diffie Hellman problem.

Random Oracle assumption. For example, the soundness of the zero knowledge proofs of correct decryption, which use the Fiat-Shamir heuristic, relies on the assumption that the hash function behaves like a random oracle.

Auditing assumption. An assumption that a sufficiently large and unpredictable fraction of a set have been confirmed or checked. For example, proper ballot generation, proper ballot printing, and accurate printing of voter preferences all need to be checked with high enough probability and unpredictability to give us confidence in the accuracy of those that were not checked. Of course, we also need to assume that when someone detects an error they make it public.

Threshold or distributed assumption. An assumption that a known threshold of authorities do not misbehave. For example, votes on the public WBB are private as long as fewer than a threshold of the authorities who share the decryption key collude, and not all the mixers collude. Robustness and reliability of the private WBB are also dependent on threshold assumptions.

Individual trust assumption. Trusting a single device. For example, the EBM a person uses to vote is trusted not to leak the vote.

Obviously, the intention is to minimise instances of trusting a single device. Our system design provides the following properties, which are detailed in Section 6:

Integrity. Based only on random oracle and auditing assumptions (with no need to trust any single device or threshold of authorities).

Non-repudiation. Based on a threshold assumption. Unless more than a threshold of private WBB peers collude, it should be infeasible to produce a properly signed receipt without properly casting a vote. This depends on proper polling-place procedures too.

Robustness. Based on a threshold assumption. If a threshold of private WBB peers are honest, a validly signed receipt is guaranteed to appear on the public WBB.

Privacy. This is the most subtle property, and needs to be discussed separately at several points. We assume that the link between an individual and their receipt is public (though names are not printed on the WBB).

The printer. Trusted not to leak ballot information.

The EBM. Trusted not to leak the vote.[3]

The printer. Prevented from performing kleptographic attacks by the ballot generation confirmation check. The proper generation of randomness for those ballots depends on at least one of the randomness generation servers being honest.

The encrypted votes on the WBB. These remain private under threshold assumptions on the decryption key sharers and assuming there is at least one honest mix server.

The system does not defend against pattern-based coercion attacks (“Italian attacks”), or other subtle coercion techniques such as forced randomisation.

Election integrity depends on a secure human procedure for ensuring that only eligible voters can vote, with at most one vote each. We assume that at some point a ledger of how many people have voted in each division at each polling place is reconciled with the published list of encrypted votes on the WBB.

For verification purposes all the voters and other observers have access to the public WBB, which is broadcast on a reliable channel. They must trust polling place procedures for checking voter eligibility and preventing polling-place ballot stuffing.

There are, however, threshold trust assumptions for liveness, reliability, and non-repudiation. In other words, we rely on certain thresholds to prevent certain kinds of failures, although all those failures would be detectable even if all the authorites misbehaved. (Whether they would in practice be detected might depend on an auditing assumption.) The private WBB peers provide a robust database implementation that distributes trust; the authorities can be confident that if the trust assumption holds then the published information will pass the verification checks.

Another way of looking at it is that the voters themselves do not need to trust individual people, hardware or software for integrity, because they can verify it. The authorities want to be confident that what they publish will indeed verify correctly. The design tells the authorities that, under certain trust assumptions, the system gives them what they need, and hence will satisfy the interested and sceptical members of the public who want to verify the outcome. The confirmation checks involved in verifiability also provide a way of catching bugs or errors in the software: a failed check might also be due to an accidental coding error.

The protocol uses digital signatures to provide evidence of many kinds of failures, rather than focusing on detection alone. This provides two kinds of benefits: voters can prove that a malfunction occurred, but can not persuade anyone that a malfunction occurred when it did not. This is important in defending against the “defaming attack”

[3] Both the printer and the EBM have an official output that is entirely deterministic, *i.e.* their printed and signed candidate or preference lists, which goes on the WBB. They cannot use that data undetectably to encode extra information—all they can do is leak information by some other channel such as Wifi, USB, subtle modifications to the printed lettering, etc.

in which people pretend to have detected a system failure, for instance by fabricating a plausible-looking receipt and claiming that they cast it honestly but it was omitted from the WBB. Of course, there can be no proof that the EBM accurately represented the voter's intention: that step is dependent on the voter's testimony.

2.4. Specific Design choices

The main departure from standard Prêt à Voter is the use of a computer to assist the user in completing the ballot. This is referred to as an "electronic ballot marker" (EBM). This modification is necessary for usability, because a vote can consist of a permuted list of about 30 candidates. It seemed infeasible for a voter to fill in a Prêt à Voter ballot form without assistance. Indeed, simply filling in an ordinary paper ballot with about 30 preferences is a difficult task.[4] Computerised assistance is an important benefit of the project, and trusting the device for privacy seemed an almost unavoidable result of that usability advantage. Hence our scheme depends on stronger privacy assumptions than standard Prêt à Voter.

Other significant departures are print on demand (rather than ahead of time) and printing the two halves separately (rather than overprinting a ballot), and hence the need to commit to the ciphertexts on the bulletin board before they are printed.

2.4.1. Cast-as-intended verification: Why use Prêt à Voter rather than another end-to-end verifiable scheme?. Wombat [Ben-Nun et al. 2012], VoteBox [Sandler et al. 2008] and several other polling-station end-to-end verifiable voting schemes guarantee integrity by using "Benaloh challenges," [Benaloh 2006] which require filling in the vote more than once. This would be time-consuming for 30-candidate STV. It would perhaps be possible to make challenges easier (for example, by letting the device remember the last vote), but the integrity guarantees still depend on the voter performing quite a subtle randomised protocol. We have opted for Prêt à Voter, in which voters may confirm the correctness of the unvoted ballot form. This confirmation process (called "auditing" in older versions) can be completed with assistance without compromising privacy, because it occurs before the person votes. It does not require the voter to redo their (possibly quite complicated) vote. It also provides dispute resolution and some accountability: there is no need to take the voter's word for how they voted. A ballot confirmation check that completes with an invalid proof can be used as evidence; an attempted ballot confirmation check that does not complete at all can have multiple (human) witnesses.

Ballot confirming is separate from voting, so additional ballot confirming by independent observers would be a convenient and practical addition to voter-initiated ballot confirmations. It would be easy for polling-place observers to see that the confirmation process did not involve casting any votes. (Wombat, StarVote and some other systems also separate the process of generating an encrypted vote from casting it.)

These processes are additive in the sense that they do not interfere with each other: the audits and inferences associated with particular trust assumptions are not affected by other audits based on different trust assumptions.

2.4.2. Unified Scanner and EBM. We have already described why completing the ballot needs to be assisted by a computer. Our original design [Burton et al. 2012a] included separate steps for filling in the ballot and then scanning the printed receipt. This was designed to separate the information of how the person voted from the knowledge of

[4] Since some people deliberately vote informally, it is difficult to say exactly what percentage of people accidentally disenfranchise themselves by incorrectly filling in their vote. About 2% of votes in the 2006 state election were ruled informal because of "numbering errors" [Victorian Electoral Commission 2007] , but the overall informality rate is closer to 10%, especially when there are many candidates on the ballot. See e.g. `https://www.vec.vic.gov.au/Results/stateby2012distributionMelbourneDistrict.html`

what their receipt looked like: the EBM learnt how the person voted, but could not subsequently recognise their ballot (and hence link it to the individual voter), while the scanner knew the receipt but did not know the corresponding plaintext. However, user studies at the VEC determined that a three-step voting process was too cumbersome for use. Also the necessity of print-on-demand meant that there was already an Internet-connected machine in the polling place that was trusted for maintaining privacy of the information on the printed ballot, including which candidate ordering corresponded to which receipt. For both these reasons, the protocol now unifies the job of the scanner and the EBM, though it retains a separate print-on-demand step. The voter first collects their ballot form, and has an opportunity to perform a confirmation check on it, then goes to an EBM to fill in the ballot, then the EBM sends the receipt electronically and also prints a paper record for the voter to check. This now means there are two online machines in the polling place (the EBMs and the ballot printers) that are trusted for vote privacy.

2.4.3. Print on Demand. This project necessitated a new protocol for verifiable printing on demand, at the polling place, of Prêt à Voter ballot forms [Culnane et al. 2013]. This includes a mechanism for confirming correctness of printed ballot forms. The ballot printer encrypts the vote deterministically using randomness generated by others, a method similar to that of Wombat. This defends against “kleptographic” privacy attacks [Gogolewski et al. 2006; Young and Yung 2004], in which the (public) ciphertexts contain deliberately poorly-chosen randomness that exposes the vote. The protocol’s main properties are:

- ensuring the candidate lists are randomly generated,
- ensuring no single generating entity knows all the (plaintext) candidate lists, and
- ensuring extra information about the candidate list cannot be leaked in the ballot ciphertexts (as in kleptographic attacks).

2.4.4. Randomised Partial Checking. The exact choice of mixnet is independent of other aspects of the protocol, but in this implementation we used Randomised Partial Checking [Jakobsson et al. 2002]. RPC mixes are not zero knowledge—each round anonymizes each vote only within half of the output. However, over multiple rounds this anonymity is significantly improved. Also the likelihood of successful (undetectable) cheating decreases exponentially in the number of substitutions, not, as for many other mixing protocols, exponentially in a predetermined security parameter. RPC was chosen partly for efficiency, and partly for the ease of explaining to the public how the mixnet works. See Section 4.5 for details.

However, improvements in both the implementation and the efficiency of zero-knowledge shuffling proofs [Furukawa and Sako 2001; Neff 2001; Wikström 2012] could make them a reasonable alternative in future versions. In theory they have superior properties, because their privacy and soundness are stronger, can be proven formally, and rely on weaker assumptions than those of RPC. However, they remain computationally intensive and difficult to explain.

3. PROCEDURES FOR VOTING AND VERIFYING

This section details, from the human perspective, how certain important security conditions are enforced by insisting on particular human procedures. The most important procedures in the polling place include authenticating voters, giving each voter the appropriate ballot, allowing them to chose at random whether to perform a confirmation check on their ballot, repeating the process until they choose to vote on one, encouraging them to check their printed vote and its signature, and insisting that they shred their candidate list. The procedures and guarantees for vision impaired voters are

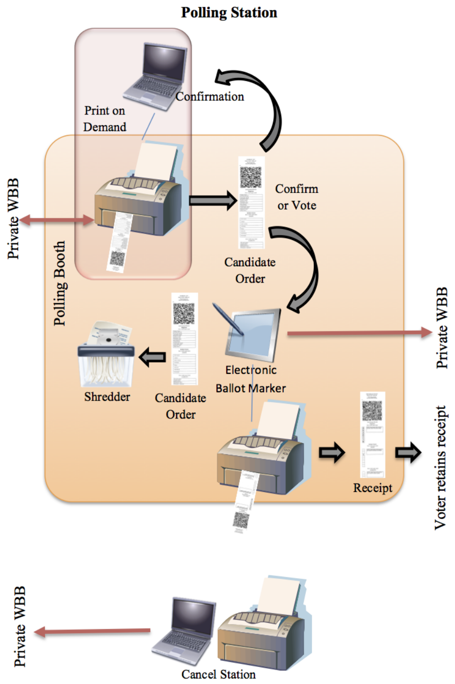

Fig. 2. Station Process.

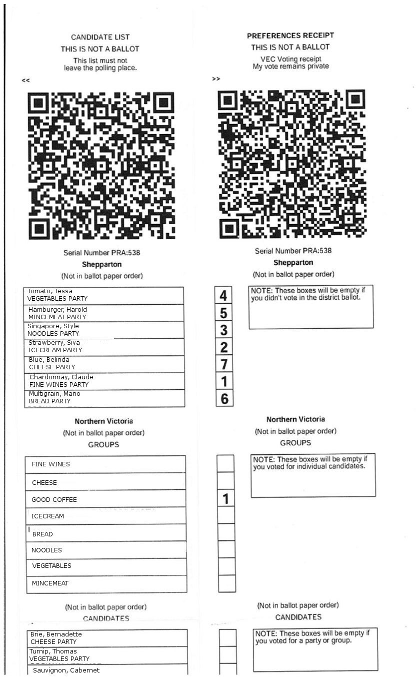
CANDIDATE LIST
THIS IS NOT A BALLOT
This list must not
leave the polling place.

<<

Serial Number PRA:538
Shepparton
(Not in ballot paper order)

Tomato, Tessa
VEGETABLES PARTY
Hamburger, Harold
MINCEMEAT PARTY
Singapore, Style
NOODLES PARTY
Strawberry, Siva
ICECREAM PARTY
Blue, Belinda
CHEESE PARTY
Chardonnay, Claude
FINE WINES PARTY
Multigrain, Mario
BREAD PARTY

Northern Victoria
(Not in ballot paper order)
GROUPS

FINE WINES
CHEESE
GOOD COFFEE
ICECREAM
BREAD
NOODLES
VEGETABLES
MINCEMEAT

(Not in ballot paper order)
CANDIDATES

Brie, Bernadette
CHEESE PARTY
Turnip, Thomas
VEGETABLES PARTY
Sauvignon, Cabernet

PREFERENCES RECEIPT
THIS IS NOT A BALLOT
VEC Voting receipt
My vote remains private

>>

Serial Number PRA:538
Shepparton
(Not in ballot paper order)

4
5
3
2
7
1
6

NOTE: These boxes will be empty if you didn't vote in the district ballot.

Northern Victoria
(Not in ballot paper order)
GROUPS

□
□
1
□
□
□
□
□

NOTE: These boxes will be empty if you voted for individual candidates.

(Not in ballot paper order)
CANDIDATES

□
□
□

NOTE: These boxes will be empty if you voted for a party or group.

Fig. 3. Separate vote printouts.

slightly different from those for sighted voters, because checking the printout requires the use of a device. This section describes what checks should be performed. Recovery from failures is described in Section 5. We begin by describing the details of the ballot.

3.1. The Ballot: Candidate list and preferences receipt

Recall that voters cast an IRV vote for a Legislative Assembly district and then for their Legislative Council region either a full STV vote or an 'ATL' shorthand. Traditionally, ATL group selections are presented on top of a thick line (hence the name); the full STV options are shown below the line (and hence called BTL votes). The idea is that voters use an EBM to help them arrange their preference list with the printed candidate list. The process in the polling place is illustrated in Figure 2 and described below. A printed candidate list consists of:

— a human-readable serial number (shortened to SerialNumber below),
— a human-readable district name (which also determines the region),
— a human-readable randomly ordered list of the candidate names for the LA district,
— a human-readable randomly ordered list of group names (for LC-ATL voting),
— a human-readable randomly ordered list of the candidate names for the LC region,
— a QR code containing all this data, plus a WBB digital signature on it.

Figure 3 shows the printed candidate list on the left side, and also the preferences as printed out by the EBM.

3.2. Typical voters: voting and auditing procedures

3.2.1. Getting a ballot. The voter presents herself to an official at a polling station and her name is marked off a register. Obviously it is essential for integrity that each vote originates with a legitimate voter, who is allowed to vote at most once, on a ballot of the appropriate division. This must be enforced by procedures at the polling place. Reconciling the number of marked-off voters in each division with the number posted on the WBB is essential for detecting ballot stuffing.

The official sends the print station a request for a ballot of the appropriate LA and LC division. The print station prints the ballot with a Private WBB signature. It is essential for privacy that no-one except the voter sees the association between the candidate order and serial number on the ballot.

Check 1: Confirming ballot correctness. Once she has obtained her ballot, the voter should decide whether she wishes to run a confirmation check on it or use it to vote. A confirmation check, called "auditing" in previous versions of Prêt à Voter, means checking that the encrypted list of candidates on the WBB matches the plaintext candidate ordering printed on the ballot. Ballot confirmation ensures that the ballot is well-formed and hence would correctly encode a vote. We describe the ballot confirmation procedure below in Section 3.2.3.

She can repeat the ballot confirmation procedure as many times as she wants in principle, each time obtaining a fresh ballot, until proceeding to vote using the last obtained, unconfirmed ballot. This implements an iterated cut-and-choose protocol: not knowing which option the voter will choose before committing to the printed ballot serves to counter any attempts by the system to manipulate votes by issuing malformed ballots. Confirming ballot construction necessarily reveals encryption information, so a ballot that has been confirmed should not be subsequently re-used for voting. It is essential for integrity that all voters have the opportunity to perform a confirmation check on as many ballots as they wish.

3.2.2. Casting a vote. In order to cast her vote, the voter takes the last obtained ballot to an EBM in a private booth. For each ballot (LA or LC), the EBM scans the the QR code which represents the permutation of the candidate ordering on her ballot, and displays the candidates in legal ballot order. Once the voter enters her choices, she is asked to confirm her choices and when she does so, the EBM prints on a separate sheet of paper:

(1) the district,
(2) the SerialNumber,
(3) the voter preferences permuted appropriately to match the Prêt à Voter ballot,
(4) a QR code with this data, plus private WBB signature.

This is the voter's receipt. Note that the EBM knows the permutation on the ballot and so re-orders the voter's selection accordingly. Note also that the EBM can assist the voter by pointing out syntactic errors, for example, duplicate rankings etc. [5]

Before printing, the EBM submits to the Private WBB exactly the data it will print on the receipt. The Private WBB accepts (and signs) the ballot only if it is accompanied by a WBB-signed serial number and division. Then at printing time the EBM adds the Private WBB signature, now including the voter preferences, as a further QR code, onto the receipt.

Check 2: EBM vote printing. The voter should check that the printed receipt has the same serial number as her ballot form, and that the printed preferences match her intended vote arranged according to the candidate order on her ballot. It is essential for integrity that all voters are encouraged to perform this check.

[5] Voters are currently allowed to spoil their ballots by casting incomplete or invalid preference lists, as long as the user interface warns them. The receipts for such ballots will reveal that they are spoiled. An alternative where the receipt does not reveal this would be to include a candidate called "spoiled ballot" who would be the first preference of any invalid ballot.

The voter now *folds her candidate list* to keep it secret, and leaves the booth with both pieces of paper. There should be a public space inside the polling place that allows officials to enforce the following procedures without exposing voters to coercion.

Existing laws preventing voters from photographing their ordinary paper ballots should also apply to the candidate list, for the same reason: a voter who retains evidence of the order the candidates are listed on her ballot can prove later how she voted.

Check 3: Private WBB Signature on receipt. The voter can check the signature on the receipt using a purpose-built smart phone app. This must of course incorporate a check that the data signed by the WBB is the same as the data printed on the paper. It is essential for non-repudiation that the voter checks the signature on her receipt before leaving the polling place. If she fails to check, and does not receive a properly-signed receipt, then she will be able to detect later, but not to prove, that her properly-submitted vote has been excluded.

This is the voter's last opportunity to cancel her vote, for example if Check 2 or Check 3 have failed, or no receipt has been issued. Procedures for vote cancellation are described in Sec 3.3.

Next, the voter shreds the candidate list. This prevents her proving how she voted. It is essential for privacy and integrity that all voters are required to shred their candidate list before leaving the polling place.

The voter should easily be able to produce multiple copies of her receipt, for example using a photocopier or a camera (on a smartphone). This combats the "trash attack," [Benaloh and Lazarus 2011] and also allows others to check her receipt on the WBB. It would also be reasonable for the VEC to retain duplicate copies of receipts, as well as letting the voters take them home. Of course there would have to be a careful procedure for ensuring that the centrally retained receipts were accurate copies of the voters'.

Check 4: Receipt appears on WBB. After a given time period, the voter can use her receipt to check that the information is correctly recorded on the WBB.

Check 5: Checking the mixing and decryption proofs on the WBB. Anyone can verify the single, public, proof that all votes are correctly mixed and decrypted.

These 5 checks provide evidence that the vote is cast as the voter intended, and included unaltered in the count. We now describe the ballot confirmation process in more detail.

3.2.3. Check 1: confirming ballot correctness. Confirming the correctness of the printout includes two steps: checking that the candidates are permuted according to the correct order (Check 1a), and checking that the ballot has been registered for the correct district (Check 1b).

Check 1a: confirming ballot printing. The ballot can be taken back to the printer. The printer prints a proof of correct ballot formation, along with a WBB signature. The WBB must record that the ballot has been confirmed, and therefore not accept any vote cast with that ballot form. As part of the confirmation process, a clear "CHECKED—NOT TO BE USED TO VOTE" message (which must be visible) is printed on the ballot form. The voter can also check the proof of decryption later on any other machine, including at home, so we are not trusting the polling-place machines for confirmation of ballot construction.

When the day's WBB becomes available (see Section 4.1), it shows which serial numbers were confirmed and displays a proof of what the candidate ordering should be. (It also shows which ones were voted and what the preferences were.) See Section 4.2.7 for the cryptographic protocol.

Check 1b: Verifying the WBB signature on the printed ballot. Each ballot is printed with a WBB signature that includes its Serial Number and district, to indicate it is legitimate and has been registered for the correct district. Voters should check this as part of ballot confirmation. (This prevents a corrupt printer from printing candidates for one district onto a ballot paper that is actually registered for a different, presumably more marginal, district.)

Mutual exclusion of confirmed and cast ballots is vitally important. The Private WBB must run a realtime check that the same ballot is not both confirmed and voted. This process is trusted for privacy, though not for integrity because violations are detectable.

Check 1b could be performed on any ballot, including those that will be used for voting on. However, it is difficult to allow this while also enforcing procedures for preventing voters from recording their candidate list. In the absence of such procedures, Check 1b is only part of the ballot confirmation procedure.

In Section 5 we describe what to do when some of these checks fail.

3.2.4. Procedures for defeating chain voting. In the chain voting attack, which applies to conventional voting, a coercer smuggles a (partially) completed ballot out of a polling place, then gives it to a voter with instructions to cast it and bring an unmarked ballot back out. The coercer uses the new ballot to repeat the attack with a new voter.

In Prêt à Voter, the chain voting attack is to smuggle out a printed ballot form, record the candidate order, then send a voter back into the polling place with instructions to vote in a particular way and return with both a receipt and a new, unmarked, ballot form. Since the coercer has already recorded the voter's candidate order, the receipt reveals the vote. The new ballot form is used to repeat the attack with a new voter.

vVote includes some technical measures to defend against chain voting. Printed ballot forms expire after 5 minutes if they have not been used to start a session, and the private WBB refuses to allow the same ballot form to be used to start another voting session once it has been used to start one. This means someone who sneaks an unused printed ballot form out of the polling place has only 5 minutes to send it in with another voter. If someone sneaks one out having used it to start a session (and the EBM sits there with session active), then attempting to sneak this back in will not work as the ballot cannot be used to start a fresh session, and the abandoned session "locks."

3.3. Ballot cancellation: individual quarantine

There may be legitimate circumstances when a voter finds a check is not successful and wishes to instruct that the vote should not be cast, including:

— when a voter claims the printed preference list differs from their intention
— when a printed preference list does not include a valid WBB signature
— when the EBM fails to produce a printout
— when an attempted ballot generation or ballot printing confirmation check fails, either because it times out or because it does not produce a valid proof of correctness or a valid signature.

A cancel request overrides any other request, such as confirming or voting. When a vote is cancelled, the cancellation is recorded against the Serial number on the (private and public) WBB. The voter must provide their candidate list in order to request cancellation. A cancellation request is allowed only if the voter presents the candidate list, and never after the voter has left the polling place.

It is important to emphasise that a vote is never cancelled except according to the following procedure. The process is:

(1) The voter requests a cancellation and provides the candidate list. If the candidate list is already shredded or missing, then cancellation is refused.
(2) Polling official scans the Serial Number on the ballot and requests a cancellation.

(3) VEC HQ provides permission for the cancellation to occur. This authorisation of cancellation is uploaded to the WBB, which replies to the printer with a receipt. Printer prints a signed cancellation onto the ballot.
(4) Voter checks signature on cancellation.
(5) Polling officials make a paper log of the cancellation, which is signed by the voter and retained by the electoral commission.[6]

The intention and expectation is that this process is used rarely, and with the explicit observation by at least two officials at the polling place. Cancellation requests should be independently recorded on paper at the polling place, and should require approval from senior officials. Those paper records should be publicly reconciled with the electronic cancel requests on the WBB.

3.4. Vision impaired voters

Vision impaired voters may need to have special procedures to help them collect their ballot privately, and insert it into the EBM, without revealing the printed candidate list. Voting uses the same software as everyone else, with adaptable audio support.

This voter is unable to perform by sight the crucial check that the printed values match her intended vote. Hence she may take both her candidate list and printed preferences to another EBM, which scans the QR code and the printed preferences, and reads her vote back to her. It can also read back the preference list or candidate list separately. This cast-as-intended verification mechanism depends on the voter finding at least one EBM in the polling place that does not collude with the first one she used. It is essential for integrity that vision-impaired voters are encouraged to check their printed preferences using an independent machine.

An alternative design would be to allow voters to bring their own devices in to perform this check, but this would violate vote privacy because the device might record the data, hence telling someone else how the person voted.

She must now destroy the candidate list. It is essential for privacy that all voters are required to shred their candidate list before leaving the polling place.

The only steps that need to be private are the ballot marking by the EBM and the check with a second EBM. All the other verification steps: confirmation of the ballot, confirmation of the receipt signature and of correct posting of the receipt to the public WBB, are exactly the same as those for typical voters, and can be performed with assistance without jeopardising ballot privacy.

Confirmation. If she has performed a confirmation check on a ballot, the voter can still go home and use her screen- or print-reader, with the same confirmation-checking software as everyone else, to make sure her candidate list matches the encrypted list on the WBB. The only important detail is that she has to make sure she knows what the cleartext candidate order is. She must either ask several people or use (a) print reader(s). This has no impact on privacy, since the confirmed ballot was not voted on.

3.5. Observing that the vVote output matches what is input into the count

VEC procedures require vVote ballots to be printed out before being incorporated into either the manual tally of paper votes (for LA ballots) or the manual data entry of

[6] We would like to be able to guarantee that people cannot walk out of the polling place with validly signed receipts that have nevertheless been cancelled; unfortunately, this cannot be enforced—voters can always pocket their valid receipt and claim they never got one. We need to be careful that they cannot cancel it and then use their preference printout to claim that their vote was incorrectly cancelled. The insistence that they sign a paper log of their cancellation request is designed to defeat this attack.

paper votes into the electronic STV count (for LC ballots). The scrutineers who observe the manual tally must reconcile these printouts with vVote's output on the public WBB.

The printouts are visually distinct from ordinary paper ballots. All the vVote votes bear a unique number on their footer which aligns with a verifiable output vote on the WBB so that they can be checked independently later. Note that these unique numbers are added to the votes *after* they have been shuffled and hence disassociated from the voter who cast them.

4. SYSTEM COMPONENT DETAILS

4.1. The Web Bulletin Board

The design presented here separates the public WBB from the private WBB, which manages the system transactions (including ballot generation, confirming/auditing, and voting) and stores them in a distributed secure database.

4.1.1. Public Web Bulletin Board (Public WBB). An authenticated public broadcast channel with memory. It is updated infrequently (*e.g.* daily). It guarantees

- that every observer gets the same information, and
- that the data written to it cannot be changed or deleted without detection.

It publishes a static digest of the day's transcript. We assume some genuine public broadcast channel[7] that can be used to send a small amount of information, specifically a signed cryptographic hash of the transcript. The signed hash of the prior commit step is included in each commit along with the other election data. When someone checks for their data on the public WBB, they also re-hash the contents and check the result against the broadcast hash. The public WBB could be replicated in the cloud.

4.1.2. Private Web Bulletin Board (Private WBB). A robust distributed database which:

- accepts items to be posted after performing basic validity checks,
- issues receipts (which are signed accepted items), and
- periodically publishes what it has received on the public WBB.

An item "clashes" if it is a vote or audit on a ballot that has already been voted or audited, unless the request is identical. The key properties of the Private WBB are:

- only items that have been posted to the bulletin board may appear on it;
- every item that has a signed receipt issued must appear on the public WBB;
- two clashing items must not both appear on the bulletin board;
- items cannot be removed from the bulletin board once they are published.

It follows from the second and third properties that if two items clash then receipts must not be issued for both of them.

Robustness is achieved through the use of several peered servers which cooperate on accepting items, issuing receipts, and publishing the public WBB. They use a deterministic threshold signature scheme which allows a sufficiently large subset of the peers to jointly generate signatures on data. The peers collectively provide the private WBB as long as a threshold of them are honestly involved in handling any item posted to the private WBB. A malicious threshold could collude to misrecord and expose votes. This is detectable by observing the public WBB, but may not be provable. The protocol with proofs of the key properties is given in [Culnane and Schneider 2014].

[7] For the 2014 State Election VEC used the "Public Notices" classified ads in the Herald Sun newspaper

4.2. Print-on-demand printers and Randomness Generation Service

The processes for printing and confirming correctness of ballot forms are vital components of Prêt à Voter. This project necessitated a completely new scheme. Here we give a "voter's eye" overview which suffices for understanding how verification works.

The main idea is that the printer generates a permuted list of candidate ciphers using randomness values generated by a distributed set of peers. The printer undertakes the expensive crypto operations, but does not have any influence over the values used in those operations. This prevents the printer from mounting "kleptographic" attacks [Gogolewski et al. 2006; Young and Yung 2004] or otherwise having any influence over the ciphertexts. See [Culnane et al. 2013] for details of the algorithms.

We post on the public WBB values that are encrypted with a threshold key, or perfectly hiding commitments. We do not post values that are encrypted with a non-thresholded key. Our system does not achieve everlasting privacy, but it does guarantee that no single entity's data (apart from the printer's or the EBM's) is enough to break ballot privacy, even given WBB data.

4.2.1. Protocol overview. Our protocol has two roles. The "randomness generation servers," $RGen_1, RGen_2, \ldots$, of which a threshold of at least one are trusted for privacy, send randomness to a "printer". The "printer" uses only that randomness in an otherwise deterministic ballot generation. When we refer to the "printer" throughout this paper we are in fact referring to the computer that is connected to the printer.

Before the voting period:

(1) Each randomness generation server generates some randomness, commits to it publicly, and sends the opening secretly to the printer. See Section 4.2.3.
(2) The printer uses the randomness combined from all the servers to generate the encrypted ballot, which it publishes. See Section 4.2.4.

During the voting period:

(3) When required, the printer prints the next ballot in sequence, with human-readable candidate names. See Section 4.2.6.

There are thus two important points for public confirmation checking:

A) A confirmation check of the encrypted ballot produced in Step 2, to check that the candidate ciphertexts are valid and that the printer used the proper randomness. This is described in Section 4.2.5.
B) A standard Prêt à Voter confirmation check of the printed ballot from Step 3, to check that the printed human-readable candidate names match those of the encrypted ballot. This is described for our scheme in Section 4.2.7.

We now describe each important step in more detail.

4.2.2. Pre-Ballot Generation. Before ballot generation starts the following must occur:

i) The election public key sharers jointly run a distributed key generation protocol to generate a thresholded private key and joint public key PK_E.[8]
ii) A list of candidate identifiers is generated and posted on the public WBB.
iii) For each printer, a list of serial numbers of the form "PrinterID:index" is deterministically generated and posted on the public WBB. These serve as row indices.
iv) Each randomness generation server and each printer establish an authenticated, encrypted channel between them.

[8] We keep the key sharers and the randomness generation servers conceptually separate

4.2.3. Randomness Generation. The randomness generation consists of each server $RGen_i$ generating a large table of secret random values, posting (public, but perfectly hiding) commitments to them on the WBB, then sending the values and the commitment openings privately to the printer. Each peer $RGen_i$ posts its commitments, and checks all others' commitments are posted, before sending the values and openings privately to the printer.[9]

4.2.4. Ballot Generation. For every i, the printer checks that every value and opening received privately from $RGen_i$ is a valid opening of the corresponding commitment on the WBB. It is important that the printer raise an alarm on any commitments that are not correctly opened. Note that the issue does not affect public verifiability, because the absence of proper commitment opening would be detected by a confirmation check of this ballot. It does, however, affect accountability: if we insist that the printer performs this check, then we can be certain that a failed confirmation is the printer's fault.

For each ballot serial number, for each candidate identifier in the pre-committed table, the printer encrypts it with a deterministic algorithm which computes a probabilistic encryption algorithm (El Gamal) by using the combined randomness from the *RGen* servers from the rows with the right SerialNo.

The resulting ciphers are then sorted into canonical order to produce a random permutation. These ciphers are posted on the public WBB. Note that the output of the encryption is pseudo-random and as such sorting the encrypted ciphers gives a pseudo-random permutation π. The printer retains this permutation so that it can print the plaintexts in the appropriate order when requested to print that ballot.

The intention is that only the printer knows which ciphertexts correspond to which candidates, but its algorithm for generating those ciphertexts is deterministic. Hence it cannot use the ciphertexts to leak information without detection. Of course, the printer could always leak that information via a side channel, but this is unavoidable and occurs with any form of electronic ballot printing or marking.

4.2.5. Ballot Generation Confirmation Checking. Generated ballots must be confirmed ("audited") to prevent a printer from generating invalid ballot ciphertexts or performing a kleptographic attack by using randomness other than that specified by the protocol. A suitable percentage of ballots are chosen at random for confirmation checking.[10] For each ballot selected, the printer posts on the WBB the randomness it used during the generation, *i.e.* to open the commitments for that SerialNumber posted on the WBB by every *RGen* peer. Anyone can verify the commitment openings and reconstruct the ballot ciphertexts from them. Thus anyone can check that the ballots were correctly constructed and that the printer used the appropriate randomness.

4.2.6. Print on Demand. This section describes the cryptographic protocol for printing ballots on demand and confirming that they have been correctly printed. The following assumes that the number of candidates on the pregenerated ballot is exactly the number required. In practice we generate a generic ballot with many candidates and use only as many as we need. The details are contained in [Culnane et al. 2013].

There is a risk that a misbehaving printer might print a completely invalid ballot, *i.e.* one that has not been part of the generation process described above, or might print the correct ballot permutation on the wrong set of candidates. Although these

[9] This is to stop the last peer choosing their randomness when they know the others'. If this was not enforced, then one bad randomness generation server colluding with a printer could determine the randomness values for each of that printer's ballots, thus breaking privacy. The bad server would wait until the printer told it all the other servers' random values, then generate its own to produce a particular final value.

[10] The questions of who chooses, how they choose, and how it can be guaranteed that they choose well enough to engender confidence in a particular election result are discussed in Section 6.1.2.

are detectable by confirmation checks, it is better to prevent this altogether. Hence the printer must obtain a signature from the WBB in order to create an authentic ballot. The WBB is attesting to those ciphertexts matching what the printer has already committed to, and to that serial number having been assigned to that electoral division.

Upon receiving a print request for a particular division,

Printer. retrieves the next available ballot
Printer. sends to the WBB:
- the SerialNumber,
- the division,

WBB. signs the serial number and division and returns the signature to the printer.
Printer. checks the WBB signature of the serial number and division and, if valid, prints the ballot and signature. If there is no valid available ballot or no valid signature, it returns an error message. The printer knows the permutation and plaintexts so does not need to do any crypto to print the ballot

Now the ballot is ready for either a ballot confirmation check or for voting.

4.2.7. Ballot Printing confirmation-Check 1a. This describes the cryptographic protocol when a voter wants to confirm a printed ballot, *i.e.* to check that the printed candidate list matches the ciphertexts on the WBB as described in Section 3.2.3, Check 1a.

Voter. requests a confirmation check from the same printer that printed their ballot,
Printer. sends to the WBB the randomness to open the commitments to the randomness used to generate the ballot,
WBB. checks the serial number has not already been voted on or confirmed and if not, opens the commitments, reconstructs the ballot, computes the permutation π, posts all the data on the public WBB, and sends a jointly signed copy of π (or candidate names in permuted order) to the printer
Printer. prints the signature
Voter. checks the signed order of candidates π against the order printed on the ballot
Voter. takes their confirmed ballot home and checks that the value provided on the WBB matches the candidate order that was signed.

After the ballot has been used to cast a vote, a confirmation check is not allowed and so the randomness no longer needs to be retained and should be deleted. Note that the randomness values are sent directly (over an encrypted channel) to the printer and not posted on the WBB. As such, there is no publicly available information that could be combined with a stolen, but previously honest, printer to reveal used ballots.

4.3. Electronic Ballot Marker (EBM)

The EBM is a computer that assists the user in filling in a Prêt à Voter ballot. This is already described in the Introduction.

4.4. Cancel Station

The Cancel Station is a supervised interface for cancelling a vote (by Serial Number). This is implemented on the same devices as the print-on-demand printers, but remains conceptually distinct and could easily be implemented on a separate device.

To request cancellation of a vote, a voter presents the printed candidate list to the election staff. The candidate list contains the signed serial number, which is scanned by the cancel station and passed to the Cancel Authority to run centrally. That cancel authority constructs its own signature on the cancellation request and sends it back to the Cancel Station. That cancel authorisation is then forwarded to the Private WBB, which returns a (signed) receipt. This “cancellation receipt” is returned for the voter to retain. A precise description of the surrounding procedures was given in Section 3.3; details of the electronic protocol are in the full version on ArXiv [Culnane et al. 2014].

4.5. Mixnet

The system uses a re-encryption mixnet which produces a noninteractive, universally verifiable proof of a shuffle and decryption (of encrypted votes) and posts it to the Public WBB. The main advantage of the re-encryption (as opposed to decryption) mixnet is that it separates the processes of shuffling and decryption. If one mix server fails, it can simply be removed and the mix re-run. Decryption works as long as a threshold of key share holders remains available.

We currently use Randomised Partial Checking [Jakobsson et al. 2002] with the modifications proposed by Khazaei and Wikström [Khazaei and Wikström 2013], including requiring all mixnet peers to contribute to the randomly generated challenges.

Each division is shuffled separately. We pad all votes going through the same mix to the same length, so all votes appear to have the maximum number of preferences. The padding is a zero/null value that is agreed and published in advance. That prevents trivially tracking the data through the mix. However, it does not prevent matching a receipt of 27 preferences to a decryption of 27 preferences. That data, along with the division name, will be revealed on the mix inputs, so each category needs to be mixed separately.

There is a privacy problem if the number of votes cast in a particular division (for ATL or BTL) is too small to be an acceptable anonymity set. This is an unavoidable problem that has nothing to do with the choice of mix protocol, though it is affected by the decision to allow receipts to reveal the choice of ATL or BTL. It is up to the administrators to decide how small an anonymity set must be before the disadvantages of privacy compromise outweigh the advantages of public verifiability.

In the recently completed deployment, some categories were determined by the electoral commission to be too small for publication. See Section 6.2.4 for a discussion of the options and implications.

The joint holders of the decryption key must check the mixing proof before decrypting the results, otherwise a cheating (initial) mix server could substitute a vote, record all the decrypted votes, and then wait for the substitution to be detected and fixed. The updated (correct) list of decrypted votes would differ by one from the incorrect list, thus revealing the vote.

5. ROBUSTNESS AND RECOVERY FROM FAILURES

Section 3 described a series of checks that voters and others can perform to ensure integrity, but did not specify exactly what happens when any of these checks fail. Some checks, such as confirming ballot correctness (check 1) produce evidence of malfeasance if they fail. It is less clear how seriously to regard a failure of Check 2, in which the voter checks their printed preferences. Unfortunately there may be some rate of false alarms, in which voters claim their vote was misrecorded when they simply misremembered it or changed their minds. Hence a zero-tolerance policy is unworkable, even though any tolerance increases the chances for vote manipulation.

The obvious fallback for most technical failures (*e.g.* a printer or EBM malfunctioning irretrievably) is to revert entirely to paper, which most voters are using anyway.

Alternatively, the EBM could run in “plain EBM mode”, in which it prints an ordinary paper ballot that the voter puts in an ordinary ballot box. The voter can check the printed vote before casting it, thus providing simple cast-as-intended verification. This would be a reasonable fallback in the event that vVote is unavailable, particularly to provide accessibility for voters who would require assistance to complete a paper ballot manually.

Officials should keep detailed records of device and procedural failures of any kind. Repeated and/or widespread occurrences, *e.g.* in the device failing to sign receipts, may necessitate de-commissioning and replacing a device, and forensic examination to determine the possible cause. The authority may need to specify a limit for each type of failure, after which further investigation is carried out.

We assume there is an established procedure for registering and authenticating voters at the polling station, for example marking off the electoral roll on production of some recognised form of identification or claim of identity. If the register is electronic, it too may be affected by a network or power outage, but this is out of scope.

We will walk through the vVote voting ceremony, following the checks described in Section 3, and consider the various error states that may arise, and the possible remedial procedures to ensure robustness of the system. Many failures require a vote to be *cancelled*, which has its own special set of procedures, described in Section 3.3.

5.1. Potential failures in ballot generation or ballot printing confirmation

Print station does not respond correctly to a confirmation or a print request. A printer that occasionally fails to produce a proof of correct opening in response to a confirmation request should be treated with suspicion. Without a good reason for expecting network failures (preventing the print station from contacting the private WBB), the assumption should be that the ballot is not properly formed, and hence this is like a failure of Check 1a. The threshold for replacing such a printer should be low. A printer that had printed a malformed ballot would not necessarily fail by proving that the ballot was malformed—more likely, it would fail to provide a proof at all.

Check 1a: Ballot is incorrectly formed. A printer that fails Check 1a has demonstrably malfunctioned and should not be used.
A printed ballot with a proper WBB signature but a candidate list that is inconsistent with the cryptographic information opened in a confirmation check, represents a serious failure of the system. The affected device should not be used.

Check 1b: Ballot has no valid WBB signature. This is also a demonstrable printer malfunction. The printer should not be used.

5.2. Potential failures in the voting ceremony

There are several different potential failures in the voting phase, but they may be difficult to distinguish and hence have similar responses.

5.2.1. Demonstrable EBM malfunctions

The preference receipt is invalid. For example, it is blank or it contains repeated numbers. This is a demonstrable EBM error.

Check 3: EBM receipt does not contain a valid digital signature. This is also an EBM error. An invalid or absent signature may be the result of simple device malfunction, but it may also have more serious implications such as failure to upload the vote to the WBB, or incorrect recording on the WBB.

Recommended action in both cases. — The vote should be cancelled, then

- — The voter should be issued with a fresh ballot and allowed to cast another vote.
- — The EBM should be removed from use until the problem is resolved.

These steps are auditable, *i.e.* voters and third parties can check that cancelled votes are recorded as such on the WBB. See Section 4.4 for details of procedures.

5.2.2. Apparent malfunctions that may be due to voter, EBM, or private WBB error. In all of the following cases, it is possible that the EBM malfunctioned but also possible that a voter did not follow the instructions, or claimed that the EBM malfunctioned when it did not. When the voter is issued a fresh ballot, it is up to them and the pollworkers whether it is a candidate list for the electronc system or a traditional paper ballot.

The EBM does not print a preference receipt. The voter cannot know if the VEC captured their vote as they intended it because they cannot verify it.

Check 2: Preferences receipt is different to what the voter expects. Note that this is indistinguishable from an voter who is mistaken, frustrated or dissatisfied claiming that their receipt is different when in fact the device functioned correctly.

Recommended action in both cases.

— The vote should be cancelled, then

— The voter should be issued with a fresh ballot, and allowed to cast another vote.

It is not possible to tell whether the failure was due to the EBM failing to send the correct vote, or a network failure, or the private WBB failing to sign and return what it received. Vote cancellation and re-voting are the most appropriate actions as the system may have failed to record the vote. However, it would be premature to remove the EBM from service following only a small number of accusations of this sort of behaviour. The authority may decide to treat sporadic, or a small number of errors as simple anomalies or voter mistakes. However, the authority should investigate the cause of an abnormal number of reported errors. Whatever the level of tolerance, it is important that ballots spoiled in this way remain secret, or the process can introduce opportunities for coercion.

5.3. Potential failures in post-election checking

Check 4: Vote is incorrectly recorded or absent from WBB. On discovering such an error, a voter can lodge a complaint. The first step would be for an official to check the reported error and the WBB signature on the receipt. If the vote is indeed wrongly recorded, and the voter has a validly signed receipt, this represents a serious system failure which implies that more than the assumed maximum number of private WBB peers have been compromised. This is a demonstrable failure.

A voter who makes this complaint without a validly signed receipt does not have a strong case. Like misprinting of preferences, this is potentially an opportunity for someone to pretend the system has malfunctioned when it has not.

Final tally proof does not verify. An error found post-publication is more likely to be caused by a technical problem *e.g.,* in uploading, than in the actual calculation. Clearly, the error would have to be investigated, the problem located and a solution sought. This is a publicly demonstrable error. While it is potentially very serious, possibly indicating failure in the “back-end” processing, it could also be the result of a minor error that is easily fixed by, for example, uploading missing data.

If no simple error is found, it should be possible to identify the mix servers whose proofs were not valid (or not present) and rerun the mix without them, eventually producing a valid, verifiable proof.

6. SECURITY CLAIMS AND ANALYSIS

Security requirements for voting systems fit into two main categories: integrity properties and privacy properties. We give a brief overview here and then more detail on each set of properties.

The protocol requires typical voters to place no trust in individual people or in hardware or software, apart from trusting that they can find at least one honest device to view the bulletin board, check signatures and verify ballot confirmations. It does of course rely on voters to perform some checks (an auditing assumption), which are detailed in Section 3. They must trust polling-place procedures for ensuring that each eligible voter is allowed to cast at most one vote, and that only eligible voters can vote.

Invalid ballots, in which the candidate list does not match the encrypted ciphertexts on the WBB, are detected at ballot confirmation by Check 1. Check 2 detects incorrect vote printing by the EBM. Incorrect vote submission by the EBM before submission to the private WBB is detected by Check 3. Check 4 detects vote substitution or removal by the WBB. Incorrect mixing or decryption would be detected by Check 5.

The vision-impaired voter is unable to do Check 2, that the EBM printed the correct ballot. She cannot ask for human assistance without destroying privacy. This leads to a distribution of trust over the machines in the polling place: she can check her vote on as many machines as she likes, and must assume that at least one is honest.

Some vision impaired voters have good enough vision to check their printout directly, just like ordinary voters, without using a second EBM. The harder it is for a cheating EBM to predict who will check directly, the harder it is to get away with cheating.

The protocol not only detects, but also provides evidence of, many kinds of failures, including the failure of ballot generation confirmations and the failure of a properly produced receipt to appear on the bulletin board. In both cases, a voter who detects such a failure has a WBB signature that proves that the system malfunctioned. This provides two benefits: voters can demonstrate to a court that a real malfunction occurred, but it is not feasible to pretend that a malfunction occurred when it did not. This is important in defending against the "defaming attack" in which people pretend to have detected a system failure which did not actually happen. Of course, there can be no proof that the EBM accurately represented the voter's intention: that step is dependent on the voter's testimony and hence is vulnerable to the "defaming" attack.

Privacy of the contents of each receipt depends on the assumption that at least two *RGen* Server generate randomness correctly and keep it secret. Further, that a threshold set of those who share the keys is honest. If these assumptions hold, then the receipt itself does not leak information about the voter's preferred candidates (though it does show how many preferences they listed, and whether they voted ATL or BTL).

Provided that these two assumptions holds, the system has some defence against kleptographic attacks on the receipt [Gogolewski et al. 2006]. This is because the receipt's random data is generated in a distributed way, and the entities that do the printing (the printer and the EBM) are deterministic. Thus information cannot be leaked in the ballot data itself without some chance of detection, though it could be subtly leaked in slight font changes or other printing effects.

Privacy of the votes also depends on the privacy of the mixing protocol. If the mix is secure, then the tallying protocol does not add any information about the link between a receipt and its vote. RPC challenges are constructed so that overall anonymity across the mix is preserved.

The system is receipt-free, meaning that (except for some major procedural violations described below) voters cannot construct a proof of how they voted. However, there are coercion attacks on this protocol, including the "Italian attack." These, along with other important details, are described below. We concentrate first on integrity properties and then discuss the subtleties of privacy.

6.1. Integrity properties

vote integrity. This means that all attempts to manipulate the votes are detectable by confirmation checks or other audits.[11]

non-repudiation. This means that failures can not only be detected, but (in most cases) demonstrated. In particular, failures of the private WBB to post something it has accepted on the public WBB can be proven by producing the signed accepted item

[11] Of course this does not imply that they will always be detected, if the appropriate checks are not performed on the manipulated ballot. The claim is that any manipulation can in principle be detected

(whether a signed ballot confirmation or a signed receipt for a submitted vote). This also defends the system against people falsely claiming to have detected an error.

prevention of ballot stuffing. This means that (under a threshold assumption, and a procedural assumption for voter markoff) only votes entered via the legitimate interface are included in the count.

Defences against ballot stuffing using the legitimate interface (e.g. by unauthorised people gaining access to a legitimate polling place) are not part of the system and must be defended against by procedural mechanisms.

Procedures must prevent voters from taking someone else's ballot off the printer and hence voting in the wrong division.

6.1.1. Justifying Integrity claims.

Vote integrity based on confirmations. An informal argument for the integrity of each person's vote is:

— The ballot-generation confirmation checks that the ballot is a permutation of properly-encrypted candidate identifiers.
— The ballot-printing confirmation checks that the printed list of candidate names matches the encrypted candidate identifiers on the WBB.
— The voter's check of the EBM's printout confirms that the correct numbers (or other marks) are recorded against the correct candidate names.
— The signature check confirms that the printed preference numbers match what was submitted to the WBB.
— The check of the vote on the WBB confirms that the correct ciphertexts were used (in the case of a larger-than-necessary generated ballot) and that the vote submitted to the WBB was posted.
— verifying the shuffling and decryption proof from the mixnet confirms that the announced output votes match those posted (encrypted) on the WBB.

Of course the first two confirmations are performed only on ballots that are *not* subsequently voted on. The argument is that any attempt to manipulate the vote by generating or printing invalid votes will be detected with high probability depending on the confirmations being numerous and unpredictable.

6.1.2. Selecting Ballots for Confirmations and Audits. Clearly it is important that we use a suitable source of randomness for the selection of the ballots to do the confirmation checks. Some combination of public confirming (with officials, scrutineers, observers, and a public source of randomness such as dice or lotto balls) with voter-initiated confirmation checks (in which any voter may choose to confirm ballot construction or printing) would be ideal. This will require further investigation to see what procedures are possible in practice. The argument about the integrity of the results of the election depends on all steps of the confirming process having been performed diligently, including those that the voters have to do themselves.

Of course, it is difficult to compute the appropriate amount of confirming for an IRV/STV election, especially in advance [Magrino et al. 2011]. This question will have to be addressed for the project, but is out of scope for this paper. We expect that most of the IRV (*i.e.* single-seat) contests will have a relatively easy margin computation in practice. However, the full STV contests are a different matter and require significant further thought. One possibility is to announce the result, explain what quantity of cheating might have been possible given reasonable estimates of the amount of confirming that was done, and ask any election challenger to demonstrate a set of votes in which they win a seat and the number of changed votes is reasonably probable given the rate of checking.

6.1.3. Hash functions and the Random Oracle assumption. The random oracle assumption is due only to the use of the Fiat-Shamir heuristic to choose challenges for the proof of correct shuffling and decryption. An alternative method of generating unpredictable challenges based on some other assumption could also be used, in which case there would be only auditing assumptions for integrity.

The method of providing public broadcast of the bulletin board contents by hash chaining also requires a collision-resistant hash function.

6.1.4. Properties of different kinds of mixnet. This design is largely independent of the kind of mixnet chosen to shuffle and decrypt the votes. However, it inherits the privacy properties of the mixnet it uses. The current implementation uses Randomised Partial Checking [Jakobsson et al. 2002], adapted according to recommendations by Khazaei and Wikström [Khazaei and Wikström 2013]. As described in Section 4.5, RPC mixnets allow a small probability of successful cheating, which decreases exponentially with the number of substituted votes. Our implementation is designed to prevent cheating unless all the mixers collude to cheat, because they all work together to compute their challenges. Even if they all do, a given probability of detection requires an amount of precomputation work that is exponential in the number of substituted votes.

We intend to continue to evaluate whether a mixnet based on zero knowledge proofs such as Verificatum [Wikström 2012] will be feasible to use in future.

6.1.5. Other integrity properties.

Serial Number uniqueness and defence against the "clash attack". The "clash attack" [Küsters et al. 2012] is a vote dropping technique that applies to many cryptographic voting schemes. An attacker (as a server or ballot generator) gives several different voters identical receipts. All affected voters see their receipt appear on the public WBB, but only one vote is counted. In our protocol, the serial numbers are carefully generated to guarantee their uniqueness (See Section 4.2.2), but this does not prevent a corrupt printer from printing off exactly the same ballot, with the same Serial Number, for many different voters. The printer would have to collude with a corrupt EBM that merely reused the private WBB signature, without resubmitting multiple instances of the same vote to the WBB. There are two reasons that this is an acceptable risk.

- The attack works only if the voters subsequently cast identical votes—otherwise the cheating EBM will be unable to produce a valid signature on the receipt, and unable to post it to the WBB. (The attack is in general harder for Prêt à Voter than for direct-encrypting schemes such as Helios and Wombat, because the attacker must commit to the identical ballot before learning the person's vote.)
- The attack is detectable by ballot printing audit, which would fail because the ballot has already been voted on.

This attack is no more effective, requires more conspirators, and has a higher probability of detection than incorrectly ordering the candidate names on the printed ballot. Hence it is appropriate (*i.e.* conservative) to include this attack implicitly in computations quantifying the extent of ballot misprinting, without explicitly counting it.

Non-repudiation and defence against the "defaming attack". Every legitimate receipt includes a WBB signature on the SerialNumber, preferences and division. Voters should check the signature before leaving the polling place. If a voter can produce such a receipt without it appearing on the WBB, then this demonstrates that the Private WBB has malfunctioned. Conversely, accusations that a particular receipt was properly submitted remain unconvincing without a valid WBB signature.

Ballot generation and printing confirmations can demonstrably fail, or demonstrably succeed, or fail by simply stopping. If a voter claims to have received an invalid opening of a printed ballot, then that should be demonstrable because a printed ballot should have a valid WBB signature. If a voter claims to have attempted to confirm a ballot but not received a result, this cannot be checked.

Misprinting of voter intention by the EBM cannot be demonstrated, because only the voter knows what they truly entered. Consequently, accusations of incorrect printing cannot be repudiated—the evidence that a particular machine is misbehaving needs to consist of a series of observations and comparisons with other machines.

Prevention of ballot stuffing. The private WBB enforces that votes may only be uploaded to the WBB by a legitimate EBM casting a vote that has been properly printed and signed. In other words, a colluding printer and EBM could stuff the ballot, but this would be detected by the reconciling of markoff data with the number of submitted ballots described in Section 3. Also (a threshold of peers of) the private WBB could stuff the ballot, by pretending to have received a legitimate signed vote, but again this is detected by reconciling with markoff data.

The EBM signature is not included on the public WBB because the protocol does not guarantee that the same signature has been sent to all private WBB peers.

6.2. Privacy properties

privacy. The system hides how each person voted, under assumptions stated below. It does reveal whether the person voted above or below the line, and how many preferences they expressed. This could potentially be used to coerce certain types of voting (such as a vote for a particular number of preferences), but it would not be possible to coerce a particular political effect. Although this could have been avoided, the opportunities for coercion are very limited, and hence did not justify the extra difficulty for voters of a more complicated protocol that kept it secret.

receipt freeness. Even a voter who deliberately colludes with a coercer cannot prove after voting how they voted (except by major violations of enforced procedures, such as taking a photo of their candidate list). Note that this implies that a voted ballot cannot also be confirmed.

resistance to kleptographic attacks. A printer attempting to leak information via the WBB data will be detected with some probability.

It is not intended to defend against “Italian attacks” (in which the voter is coerced into producing a vote matching a detectable pattern) or randomisation attacks (in which a voter is coerced to produce a receipt of a particular form, which has a random effect on the actual vote). The coercion attack described in Kelsey *et al.* [Kelsey et al. 2010] Sec 4.3 is also possible, but so complicated for this scheme that it would be very far from the most effective way to coerce voters. It is not worth defending against a coercion attack that is harder to execute and no more effective than the “Italian attack” already possible in the existing paper system.

Defences against coercion associated with failing to shred the candidate list must be enforced by procedural mechanisms. Deliberate uses of out-of-band recording technology (such as taking a photo of the candidate list before shredding it) or side-channel information leakage (from the EBM or printer) must be defended against by mechanisms outside the vVote system. Section 3.2.4 describes technical and procedural measures, which are unrelated to Prêt à Voter, for defending against chain voting.

6.2.1. Justifying privacy claims. These coalitions can violate individual ballot privacy:

— The printer that generated and produced that ballot,
— All but one of the *RGen* servers,
— The EBM the voter used,

— The number of mix servers necessary for breaking shuffling privacy, depending on the mixnet being used,
— A threshold of key sharing authorities,

The crucial claim is that smaller coalitions cannot. This is expanded into several specific claims below. Clearly if the printer leaks its information it can violate vote privacy for everyone who used a ballot it printed. This means that practical opportunities for compromising the printer must be reduced as much as possible, *e.g.* turning off the wireless connection. Apart from the printer and an electronic ballot marker, no other single entity can violate vote privacy. This is summarized in the first two claims below, which are justified in [Culnane et al. 2013].

CLAIM 1. *A collusion of all but two randomness generation authorities does not have sufficient information to recover the ballot permutation (in polynomial time with non-negligible probability).*

CLAIM 2. *The posted ballots on the WBB reveal, for each receipt, whether the vote was ATL or BTL and how many preferences were cast, but no other information, unless a threshold of key sharers colludes.*

The following claim is really an assumption about the mixing process, whose accuracy depends on which mixnet is used. For a justification of this claim for RPC, see [Jakobsson et al. 2002; Khazaei and Wikström 2013].

CLAIM 3. *The mixing process anonymises votes within anonymity sets defined by their division, ATL/BTL choice, and length of preference list. The assumptions about mixnet collusion for privacy violations depend on the mixnet being used.*

For example, in RPC, if all but one pair of mix servers exposes their permutation, then each anonymity set is half the total being mixed. In a mixnet based on zero knowledge proofs, votes are anonymised within the whole set if at least one mixer is honest.

Our system is not susceptible to the replay attack described in [Khazaei and Wikström 2013] because all ballots are pregenerated in a distributed fashion.

6.2.2. Receipt freeness. Receipt freeness is a subtle property and we do not have a formal argument for it. However, the main idea behind Prêt à Voter is to provide the voter with either a proof of the contents of a ballot's encrypted values, or an opportunity to vote on the ballot, but never both for the same ballot. A ballot that is allowed to be voted on should never have revealed the random values used to produce it. The threat of using the confirmation process to expose the contents of a voted ballot is ameliorated by the electronic locking process described in Sec 4.2.7.

A voter could attempt to collude with a corrupt printer to produce a receipt, and could promise not to perform a ballot printing confirmation check, but the incorrect formation of the ballot necessary to produce such a receipt would be caught by a ballot construction confirmation check with some probability.

6.2.3. Privacy Threats Ameliorated By Procedural Controls. As the voter inputs her choices into the EBM, the device necessarily "learns" how she voted. The potential for the EBM to leak vote information clearly raises privacy issues. Any data stored in the EBM's memory should be deleted, ideally after each session.

Prêt à Voter introduces a privacy threat that does not exist for either standard paper voting or for DRE's with VVPAT: someone may discover and record an unvoted ballot's candidate order and Serial Number, then learn the vote choices when they are later posted on the WBB. Therefore there procedural controls must protect both the printout and the electronic data on the printer from observation by anyone but the voter.

As for any voting system, voters may ask for assistance at a point that potentially violates their privacy simply because the assistant sees what the voter has already written or entered. This threat to privacy however, exists in the current system.

6.2.4. Privacy issues arising from small populations and complex ballots. Sometimes vote privacy is unachieveable, for example if everyone in one polling place votes the same way, and results for that polling place are observed or announced alone. Small populations exacerbate problems caused by the complexity of Victorian ballots. It is easy for a voter (or a coercer) to choose a BTL vote that is highly likely to be unique. If the BTL votes are made public, this allows a voter to prove how they voted. Since vVote receipts expose whether the person voted ATL or BTL, and how many preferences they expressed, this problem is exacerbated again: it may be possible to identify a vote uniquely based on its division and number of BTL preferences. If only one person cast a particular number of preferences below the line in a division, that person's vote would have to be withheld from publicly verifiable decryption.

There are some techniques for verifiable privacy-preserving tallying STV tallying using (more) mixing and homomorphic sums [Benaloh et al. 2009; Heather 2007]. However, these work only on a complete list of votes, while the vVote votes need to be input into an (unencrypted) existing VEC counting system. This leaves us with ad hoc approaches based on how many people vote in each division, how close the election result is, and what anonymity thresholds are deemed acceptable. The problem could be mitigated by doing on-demand decryption of the packed candidate IDs, so we only decrypt the next pack when the previous one has been eliminated. However, that still does not guarantee that privacy will be preserved, since if all preferences are counted the full information will be made public anyway.

In the recently completed deployment, some divisions had very few vVote voters. 149 divisions, including all the LC-BTL votes, were determined (by the electoral commission) to be too small for publication. This left 88 for proper mixing and publication. Fortunately all the small sets were too small to affect an electoral outcome.

6.2.5. Other possible attacks. "Psychological" attacks are always possible, for example a coercer could convince voters that he can decrypt their receipts and find out how they voted. Voter education could mitigate this attack; however psychological attacks will be a problem for virtually any end-to-end verifiable system.

7. CONCLUSION, REPORT ON THE DEPLOYMENT, AND FUTURE WORK

For the vVote system we have taken the original design concept of Prêt à Voter, and extended and customised it to the requirements of the State of Victoria, Australia. This practical deployment entailed a whole suite of technical and procedural problems not previously addressed in the academic literature. We have developed novel solutions to make the design applicable to the particular aspects of the target election while maintaining end-to-end verifiability, notably: the introduction of electronic ballot markers for capturing and casting votes; the requirements of preference voting whereby candidates are ranked in order of preference rather than simply selected; the design of a secure and robust web bulletin board; the requirement to print ballot papers on demand at polling places while preserving the assurances that their cryptographic construction is correct; and designing the system around procedures that are straightforward for voters and pollworkers to follow.

This paper has described the general background and motivation for the system, and its human processes and cryptographic protocols, motivated how these tie in to other design decisions. Finally the paper has considered the system from a robustness point of view, and has provided a security analysis of the system with particular emphasis on the privacy and integrity provided by the system.

7.1. Report on the deployment

The system successfully took 1121 votes. During the deployment, six e-votes were cancelled. This seems to have been due to networking problems. All the affected voters subsequently cast a paper ballot.

Many divisions had too few e-votes for publication, and hence were not verifiable. This is a result of limited deployment, not a protocol flaw, and would disappear if more people used the system.

This project ran on a constrained budget on a very firm timetable dictated by the election cycle. Certain tradeoffs were forced by timing, and by the current availability of certain tools, but do not necessarily represent the bestmin the long term. For example, the unified scanner and EBM could be revisited with better tools for printing and better optical character recognition on the scanner. Also the choice of mixnet could be reviewed if a more efficient zero-knowledge open source mix becomes available.

Numbers were undetermined until the last minute and so the system was developed to handle a much larger number of voters than it eventually actually did. Many of the design tradeoffs and developments meant for larger numbers, such as the agreeement algorithm for the private WBB, were not necessary for the numbers we actually had.

7.2. Reflections on putting it into practice, with suggestions for future improvements

7.2.1. Usability. There is a tension between verifiability, privacy and usability. For professional administrators who actually have to run an election, usability is the overwhelming priority. This includes usability for voters and for the temporary staff hired to conduct the election on the day. In practice any working system has to be usable. Although we had put significant effort into a design that was easy for voters to use, the polling-place procedures need to be simpler if this is to be more widely deployed. The challenge is to retain the usability, while also providing verifiability and privacy.

There is an important link between verifiability, usability, and the practicalities of organising a polling place. For example, the electoral commission was able to conduct quite extensive ballot construction audits, because this could be performed before polls opened, at a time when there was not enormous pressure from other tasks. However, ballot printing confirmations were much more problematic, requiring a voter to pause just at the critical point when they were about to go and vote. An important item of future work is to redesign the ballot printing confirmations so as to make them easier to perform without getting in the way of the main voting process.

The protocol uses signature checking to achieve non-repudiation in many circumstances, particularly to allow voters to prove that their receipts are properly accepted by the private WBB. In practice signature checking is complicated, both the procedures in the polling place and the prior work necessary to download the app. Alternative weaker but simpler methods of preventing the defaming attack, such as anti-counterfeiting paper, or franking or stamping of receipts in the polling place, might offer better usability for only slightly reduced security.

7.2.2. Vision impaired voters and assumptions. If the system is used more widely, this would reduce the importance of the current assumption that vision impaired voters can find a non-colluding EBM in the same polling place as the one they used to vote. The original intention was that the system would be used for a wide class of voters, including disabled ones, and that the effect of some able-bodied voters conducting audits would thus be shared by those who were unable to conduct them themselves. However, this first deployment was not extended to non-disabled voters within the state, which left us with the choice between inviting such voters to verify with their own device (a serious practical risk to privacy) or to rely on another device in the same polling place (which is far from ideal) or not at all (and rely on being indistinguishable from sighted voters). It was decided to rely on the second device in the polling place, mainly because this is a harmless and additive assumption. Furthermore, even in this deployment some eligible voters were perfectly able to read their own printout—the possibility of some such voters checking their printout directly might deter an attack by colluding EBMs in the same polling place.

ACKNOWLEDGMENTS

We are grateful to the following for contributions and discussions during the formulation of this design: Craig Burton, Matthew Casey, James Heather, Rui Joaquim, Thea Peacock, Olivier Pereira, Sriramkrishnan Srinivasan, Roland Wen, Jason White, Douglas Wikström, Zhe (Joson) Xia and Karen Young. We are also grateful to Ron Rivest for the suggestion in Section 6.1.2, and to the anonyous reviewers for their careful and extensive comments and suggestions on this paper.

The name "vVote" is a trademark of the Victorian Electoral Commission.

This work was supported by the EPSRC Trustworthy Voting Systems project EP/G025797/1, and by the Fonds National de Recherche (FNR) Luxembourg SeRTVS project.

References

Susan Bell, Josh Benaloh, Michael D. Byrne, Dana DeBeauvoir, Bryce Eakin, Gail Fisher, Philip Kortum, Neal McBurnett, Julian Montoya, Michelle Parker, Olivier Pereira, Philip B. Stark, Dan S. Wallach, and Michael Winn. 2013. STAR-Vote: A Secure, Transparent, Auditable, and Reliable Voting STAR-Vote: A Secure, Transparent, Auditable, and Reliable Voting System. *USENIX Journal of Election Technology and Systems (JETS)* 1, 1 (August 2013).

Jonathan Ben-Nun, Niko Fahri, Morgan Llewellyn, Ben Riva, Alon Rosen, Amnon Ta-Shma, and Douglas Wikström. 2012. A new implementation of a dual (paper and cryptographic) voting system. In *5th International Conference on Electronic Voting (EVOTE)*.

Josh Benaloh. 2006. Simple Verifiable Elections. In *Proc. 1st USENIXAccurate Electronic Voting Technology Workshop*. http://www.usenix.org/events/evt06/tech

Josh Benaloh and Eric Lazarus. 2011. *The Trash Attack: An Attack on Verifiable Voting Systems and a Simple Mitigation*. Technical Report MSR-TR-2011-115. Microsoft.

Josh Benaloh, Tal Moran, Lee Naish, Kim Ramchen, and Vanessa Teague. 2009. Shuffle-sum: coercion-resistant verifiable tallying for STV voting. *IEEE Transactions on Information Forensics and Security* 4, 4 (2009), 685–698.

Craig Burton, Chris Culnane, James Heather, Thea Peacock, Peter YA Ryan, Steve Schneider, Sriramkrishnan Srinivasan, Vanessa Teague, Roland Wen, and Zhe Xia. 2012b. Using Prêt à Voter in Victorian State elections. In *Electronic Voting Technology Workshop / Workshop on Trustworthy Elections*.

Craig Burton, Chris Culnane, James Heather, Thea Peacock, Peter Y. A. Ryan, Steve Schneider, Sriramkrishnan Srinivasan, and Zhe Xia. 2012a. A supervised verifiable voting protocol for the Victorian Electoral Commission. In *Proc. 5th International Conference on Electronic Voting*.

Richard Carback, David Chaum, Jeremy Clark, John Conway, Aleksander Essex, Paul S. Herrnson, Travis Mayberry, Stefan Popoveniuc, Ronald L. Rivest, Emily Shen, Alan T. Sherman, and Poorvi L. Vora. 2010. Scantegrity II Municipal Election at Takoma Park: The First E2E Binding Governmental Election with Ballot Privacy. In *Proc. USENIX Security*.

David Chaum, Benjamin Hosp, Stefan Popoveniuc, and Poorvi L. Vora. 2009. Accessible Voter-Verifiability. *Cryptologia* 33, 3 (2009), 283–291.

Chris Culnane, James Heather, Rui Joaquim, Peter Y. A. Ryan, Steve Schneider, and Vanessa Teague. 2013. Faster Print on Demand for Prêt à Voter. *USENIX*

Journal of Election Technology and Systems 2, 1 (2013).

Chris Culnane and Steve Schneider. 2014. A Peered Bulletin Board for Robust Use in Verifiable Voting Systems. In *IEEE Computer Security Foundations Symposium*. extended version at arXiv:1401.4151.

Chris Culnane, Steve Schneider, Peter Y A Ryan, and Vanessa Teague. 2014. vVote: a verifiable voting system. (2014). ArXiV eprint: arXiv:1404.6822.

Roberto Di Cosmo. 2007. On Privacy and Anonymity in Electronic and Non Electronic Voting: the Ballot-As-Signature Attack. (2007). http://hal.archives-ouvertes.fr/hal-00142440/en/

Jun Furukawa and Kazue Sako. 2001. An efficient scheme for proving a shuffle. In *CRYPTO 2001*. Springer, 368–387.

Marcin Gogolewski, Marek Klonowski, Przemyslaw Kubiak, Miroslaw Kutylowski, Anna Lauks, and Filip Zagórski. 2006. Kleptographic Attacks on E-Voting Schemes. In *International Conference on Emerging trends in Information and Communication Security*. 494–508.

James Heather. 2007. Implementing STV securely in Prêt à Voter. In *IEEE Computer Security Foundations Symposium*. 157–169.

Markus Jakobsson, Ari Juels, and Ronald Rivest. 2002. Making Mix Nets Robust for Electronic Voting by Randomized Partial Checking. In *USENIX Security Symposium*. 339–353.

John Kelsey, Andrew Regenscheid, Tal Moran, and David Chaum. 2010. Attacking Paper-Based E2E Voting Systems. In *Towards Trustworthy Elections, New Directions in Electronic Voting*. 370–387.

Shahram Khazaei and Douglas Wikström. 2013. Randomized partial checking revisited. In *Topics in Cryptology, CT-RSA 2013*. Springer, 115–128.

R. Küsters, T. Truderung, and A. Vogt. 2012. Clash Attacks on the Verifiability of E-Voting Systems. In *IEEE Symposium on Security and Privacy (S&P 2012)*. IEEE Computer Society, 395–409.

Thomas R. Magrino, Ronald L. Rivest, Emily Shen, and David Wagner. 2011. Computing the Margin of Victory in IRV Elections. In *USENIXAccurate Electronic Voting Technology WorkshopWorkshop on Trustworthy Elections*.

C. Andrew Neff. 2001. A verifiable secret shuffle and its application to e-voting. In *Conference on Computer and Communications Security*. ACM, 116–125.

Ronald L Rivest. 2008. On the notion of 'software independence' in voting systems. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 366, 1881 (2008), 3759–3767.

Peter Y. A. Ryan, David Bismark, James Heather, Steve Schneider, and Zhe Xia. 2009. Prêt à Voter: a voter-verifiable voting system. *IEEE Transactions on Information Forensics and Security* 4, 4 (2009), 662–673.

Daniel R. Sandler, Kyle Derr, and Dan S. Wallach. 2008. VoteBox: A tamper-evident, verifiable electronic voting system. In *Proc. 17th USENIX*.

Victorian Electoral Commission. 2007. Report to Parliament on the 2006 Victorian State election. (July 2007).

Douglas Wikström. 2012. Verificatum. http://www.verificatum.org/verificatum/.

Adam L. Young and Moti Yung. 2004. *Malicious cryptography - exposing cryptovirology*. Wiley.

Received XXXX; revised YYYY; accepted ZZZZ