Human Readable Paper Verification of Prêt à Voter

David Lundin[1] and Peter Y.A. Ryan[2]

[1] University of Surrey, Guildford, Surrey, UK
d.lundin@surrey.ac.uk
[2] University of Newcastle Upon Tyne, Newcastle, UK
peter.ryan@ncl.ac.uk

Abstract. The Prêt à Voter election scheme provides high assurance of accuracy and secrecy, due to the high degree of transparency and auditability. However, the assurance arguments are subtle and involve some understanding of the role of cryptography. As a result, establishing public understanding and trust in such systems remains a challenge. It is essential that a voting system be not only trustworthy but also widely trusted.

In response to this concern, we propose to add a mechanism to Prêt à Voter to generate a conventional (i.e. human readable) paper audit trail that can be invoked should the outcome of the cryptographic count be called into question. It is hoped that having such a familiar mechanism as a safety net will encourage public confidence. Care has to be taken to ensure that the mechanism does not undermine the carefully crafted integrity and privacy assurances of the original scheme.

We show that, besides providing a confidence building measure, this mechanism brings with it a number of interesting technical features: it allows extra audits of mechanisms that capture and process the votes to be performed. In particular, the mechanism presented here allows direct auditing of ballots that are actually used to cast votes. This is in contrast to previous versions of Prêt à Voter, and indeed other verifiable schemes, that employ a cut-and-choose mechanism. The mechanism proposed also has the benefit of providing a robust counter to the danger of voters undermining the receipt-freeness property by trying to retain the candidate list.

1 Introduction

There has been much concern lately as to the trustworthiness of electronic voting systems such as touch screen devices, where the integrity of the count depends heavily on the correctness of the code running on the voting machines. Researchers have pointed out the ease with which the count could be manipulated in virtually undetectable ways [10]. One response to these concerns, originally proposed by Mercury [13], is to incorporate a *Voter Verifiable Paper Audit Trail* (VVPAT), essentially a paper copy of the voter's intent that is printed in the booth and checkable by the voter. Whilst such a mechanism is doubtless an

S. Jajodia, and J. Lopez (Eds.): ESORICS 2008, LNCS 5283, pp. 379–395, 2008.

improvement on the situation in which the count is retained solely in software, with no paper back-up at all, there are still problems:

- Paper audit trails are not invulnerable to corruption.
- It is not clear how any conflicts between the computer and paper audit counts should be resolved.
- Humans are notoriously bad at proof-reading, especially their own material, and hence bad at detecting errors in a record of their choices [3].
- Even if the voter does notice a discrepancy with the paper record created at the time of casting, it may be tricky to resolve, especially without undermining the privacy of the ballot.
- It is not clear under what circumstances the audit trail should be invoked.

An alternative response is to devise schemes that provide high levels of assurance via a high degree of transparency and with minimal dependency on technology. Such schemes provide *Voter-verifiability* in a different sense: voters have a way to confirm that their vote is included in a universally auditable tabulation that is performed on an append-only Web Bulletin Board (WBB) [6].

Prêt à Voter [28,27,1,21,23,24,26,11,12,30] is a particularly voter-friendly example of such high assurance, trustworthy voting schemes. It aims to provide guarantees of accuracy of the count and ballot privacy that are independent of software, hardware etc. Assurance of accuracy flows from maximal transparency of the process, consistent with maintaining ballot privacy.

Verifiable schemes like Prêt à Voter, VoteHere [14], and PunchScan [5], arguably provide higher levels of assurance than even conventional pen-and-paper elections, and certainly far higher assurance than systems that are dependant on the correctness of (often proprietary) code. However, the assurance arguments are subtle and it is unreasonable to expect the electorate at large to understand them. Whether the assurances of *experts* will be enough to reassure the various stakeholders is unclear. This is probably especially true during the early phase of introduction of such systems until a track record has been established. It seems sensible therefore to explore the possibility of incorporating more conventional mechanisms to support public confidence.

Randell and Ryan [17] explored the possibility of voter-verifiable schemes without the use of cryptography. This tried to achieve similar integrity, verifiablity and privacy goals but using only more familiar, physical mechanisms such as scratch strips. The resulting levels of assurance, in the technical sense, are not as high as for Prêt à Voter.

A more recent proposal is *ThreeBallot* due to Rivest [18]. This does indeed provide voter-verifiability but at the cost of a non-trivial voter interface: voters a required to mark three ballots in such a way as to encode their vote (two votes for their candidate of choice, one for all others) and to retain one ballot, chosen at random. Besides the non-trivial voter interface, a number of vulnerabilities in ThreeBallot have been identified, several in Rivest's original paper. It is probably fair to conclude that ThreeBallot, whilst being a conceptual breakthrough, does not, as it stands, provide a viable scheme for real elections.

Here we explore a rather different route: supplementing a cryptographic scheme with a conventional paper audit trail backup that we refer to as a *Human Readable Paper Audit Trail* (HRPAT). This approach was first explored in [20]. Introducing such a mechanism may introduce certain vulnerabilities not present in the original scheme. However, it may be argued that it is worth introducing such risks, at least during trials and early phases of deployment.

In this paper we propose some enhancements to the scheme [20] that gives rise to a number of additional auditing possibilities. This minimises threats to ballot privacy while maximising the reassurance of having a conventional mechanism as a backup. Once sufficient levels of trust and confidence have been established in a verifiable, trustworthy scheme like Prêt à Voter, we would hope that the scaffolding of an HRPAT could be cast aside.

Besides the confidence building aspects we find that the HRPAT mechanisms proposed here provide a number of unexpected technical benefits. It can provide a robust counter to the danger of voters attempting to leave the polling station with the left hand element of the Prêt à Voter ballot form. This shows the candidate order and so could provide a potential coercer with proof of the vote. A number of possible counter-measures to this threat have been identified previously, for example the provision of *decoy* candidate lists [23,25], but the mechanism here appears to be particularly robust. The procedure we propose here involves the officials verifying that the voter submits the component of the ballot that carries the candidate order at the time of casting.

The approach proposed here enables a number of additional auditing procedures to be introduced that significantly increase the assurance of accuracy, assuming that the integrity of the paper audit trail can be ensured.

The second author previously proposed a Verified Encrypted Paper Audit Trail (VEPAT) mechanism [29]. Whilst this enhances assurance from a technical point of view, the audit trail is not human-readable and so it does not really help with public perception and confidence. It is hoped that the scheme proposed here should be more familiar and understandable.

2 Outline of Prêt à Voter

The key innovation of the Prêt à Voter approach is to encode the vote using a randomised candidate list. Suppose that our voter is called Anne. At the polling station, Anne chooses at random a ballot form sealed in an envelope; an example of such a form is shown in Figure 1.

In the booth, Anne extracts her ballot form from the envelope and makes her selection in the usual way by placing a cross in the right hand column against the candidate of her choice (or, in the case of a Single Transferable Vote (STV) system for example, she marks her ranking against the candidates). Once her selection has been made, she separates the left and right hand strips along a perforation and discards the left hand strip. She is left with the right hand strip which now constitutes her *privacy protected receipt*, as shown in Figure 2.

Obelix	
Idefix	
Asterix	
Panoramix	
	7304944

Fig. 1. Prêt à Voter ballot form

X
7304944

Fig. 2. Prêt à Voter ballot receipt (encoding a vote for "Idefix")

Anne now exits the booth clutching her receipt, registers with an official, and casts her receipt. Her receipt is placed over an optical reader or similar device that records the cryptographic value at the bottom of the strip and records in which cell her X is marked, or the vector of rankings etc. This digital copy of her receipt is posted to a secure Web Bulletin Board (WBB). Her original, paper receipt is digitally signed and franked and returned to her to keep.

The randomisation of the candidate list on each ballot form ensures that the receipt does not reveal the way she voted, thus ensuring the secrecy of her vote. Incidentally, it also removes any bias towards the candidate at the top of the list that can occur with a fixed ordering.

The value printed on the bottom of the receipt, that we refer to as the *onion*, is the key to extraction of the vote during the tabulation phase. Buried cryptographically in this value is the information needed to reconstruct the candidate order and so extract the vote encoded on the receipt. This information is encrypted with secret keys shared across a number of tellers. Thus, only a threshold set of tellers acting together are able to interpret the vote encoded on the receipt.

After the voting has closed, voters (or perhaps proxies acting on their behalf) can visit the secure Web Bulletin Board (WBB) and confirm their receipts appear correctly. Once any discrepancies are resolved, the tellers take over and perform anonymising mixes and decryption of the receipts. All the intermediate stages of this process are committed to the WBB for later audit. Various auditing mechanisms are in place to ensure that all the steps, the creation of the ballot forms, the mixing and decryption etc are performed correctly. These are carefully designed so as not to impinge on ballot privacy. Full details can be found in, for example, [22].

An early version of the Prêt à Voter system used a decryption mix network to break the link between an encrypted receipt and the plaintext vote [1]. We call this configuration of the system Prêt à Voter 2005. When the decryption mix network was exchanged for a re-encryption mix network in Prêt à Voter 2006

[26] this made provisions for a range of measures that protect the secrecy of the election, for example the on-demand printing of ballot forms in the booth. A further extension of the system exchanged the Elgamal encryption for Paillier [22].

2.1 The Security Properties

Cryptographic schemes, like those in the Prêt à Voter class, strive to provide the following properties:

1. Accuracy
2. Ballot privacy and coercion resistance
3. Voter-verifiablity

Accuracy can be thought of as the requirement that all legitimately cast votes should be included in the tabulation. We will assume that a correct register of legitimate voters is maintained and that mechanisms are in place to authenticate voters and ensure that each voter can cast at most one vote.

Ballot privacy requires that, for any given voter, it should be impossible for anyone, other than the voter, to determine how they voted. Coercion resistance requires that even if the voter is prepared to cooperate with a coercer throughout the vote casting protocol, the voter cannot construct a proof of how they voted.

Voter-verifiability requires that voters should have a way to confirm that their votes are accurately included in the tabulation. Clearly this has to be done in a way that does not violate coercion resistance.

Prêt à Voter allows all voters to check that their votes were recorded as intended by the electronic voting system and then the public verifiability allows any interested organisation or individual to check that all recorded, encrypted votes are transformed into countable plain text votes correctly. Thus the tabulation of the receipts is universally verifiable. The assurance arising from the voter checks relies on a reasonable number of voters checking their receipts on a web site.

The goal is to provide high assurance that these properties are guaranteed for any election without needing to trust any component of the system, be it software, hardware or humans. Rivest has coined the term *software independence* to refer to this design requirement [19].

Analysis of the Prêt à Voter schemes indicates that, subject to certain assumptions, they fulfill the above requirements. We refer the reader to the various papers and tech reports for the details.

The scheme that we describe here inherits the security properties of Prêt à Voter 2006. For the accuracy requirement it can be argued that this scheme provides higher guarantees, as long as we assume that the integrity of the paper audit trail can be guaranteed. Regarding the privacy requirements there is a danger that the HRPAT mechanism may undermine the carefully wrought properties of the 2006 scheme. We will discuss the differences in the security guarantees provided by Prêt à Voter 2006 and the scheme of this paper in our conclusions Section 5.

3 Preliminaries

In this section we introduce some of the primitives that we need in what follows.

3.1 Threshold ElGamal

We recall the probabilistic algorithm due to ElGamal, [4]: given a large prime p and a generator α of a q-order subgroup of Z_p^*. A party A chooses a secret key k and computes β:

$$\beta := \alpha^k \pmod{p}$$

The public key is p, α and β. k is the secret key. Encryption of m yields a pair of terms computed thus:

$$c := (y_1, y_2) := (\alpha^r, m \cdot \beta^r) \pmod{p}$$

where r is chosen at random. A decrypts c as follows:

$$m = y_2 / y_1^k \pmod{p}$$

The security of ElGamal rests on the presumed difficulty of taking discrete logs in a finite field. Thus, recovering the secret k exponent from knowledge of p, α and β is thought to be intractable.

A randomising algorithm like ElGamal allows the possibility of re-encryption: anyone who knows the public key can re-randomise the original encryption with a new random value r':

$$(y_1', y_2') := (\alpha^{r'} \cdot y_1, \beta^{r'} \cdot y_2)$$

which gives:

$$(y_1', y_2') := (\alpha^{r'+r}, \beta^{r'+r} \cdot m)$$

Clearly, this is equivalent to simply encrypting m with the randomisation $r + r'$ and decryption is performed exactly as before. We will see the utility of re-encryption when we come to describe anonymising mixes. Note that, crucially, the device performing the re-encryption does not use any secret keys and at no point in the re-encryption process is the plaintext revealed.

In fact we will use *exponential ElGamal*, where m is encrypted as:

$$c := (y_1, y_2) := (\alpha^r, \alpha^m \cdot \beta^r) \pmod{p}$$

Thus the plaintext is carried in the exponent of α. This is convenient when we come to transform the receipts to pure ElGamal terms prior to mixing. It does mean however that we have to limit the plaintext space to avoid having to extract discrete logs to obtain the plaintext. Furthermore, we will use a threshold form of ElGamal. We omit the details and refer the reader to [16], for example.

4 The Scheme

In this section we first present the HRPAT Prêt à Voter ballot form with its onions and how they are created and printed. We then describe the on-demand printing of the candidate list and the method by which votes are cast. Finally we show how the encrypted receipts are decrypted and how the HRPAT can be used to verify the electronic election.

4.1 The Ballot Form and Its Use

The usual Prêt à Voter ballot form is modified to comprise two overlaid pages. The bottom page has the usual two portions: the left hand portion carries an onion and a serial number. The top page overlays the right portion of the bottom sheet and carries another onion value. The top page has a carbon layer or similar on the back to ensure that marks applied to the top page transfer to the bottom

	POST	RETAIN
$onion_L$		$onion_R$
serial		

Fig. 3. The ballot form in two pages

	RETAIN
$onion_L$	$onion_R$
serial	

Fig. 4. The ballot form complete

	RETAIN
$candidate_B$	
$candidate_C$	
$candidate_A$	
$onion_L$	$onion_R$
serial	

Fig. 5. The ballot form with candidates printed

	RETAIN
$candidate_B$	
$candidate_C$	X
$candidate_A$	
$onion_L$	$onion_R$
$serial$	

Fig. 6. The ballot form with marks

	POST	RETAIN
$candidate_B$		
$candidate_C$	X	X
$candidate_A$		
$onion_L$		$onion_R$
$serial$		

Fig. 7. The marked ballot form in two pages

page. The layout of the ballot form is shown in Figure 3. This means that when the top page is aligned over the right column of the bottom page, as is the case when the voter receives the ballot form, the ballot form looks as shown in Figure 4. When the voter makes her mark in the right hand column of this complete form the mark is made on both pages.

The reader will notice that there are no candidate names printed in Figure 3. This is because we are incorporating the on-demand printing of ballot forms introduced in previous papers [26]. When the voter has identified herself to the poll station workers she is allowed to randomly choose a ballot form such as that in Figure 4. At this stage $onion_L$ and $onion_R$ are concealed (for example by a scratch strip) so that they cannot be read by either the poll station worker nor anyone else at the polling station. The other value, *serial*, is noted in the register next to the voter's name.

The voter takes the form into the voting booth where she makes $onion_L$ visible and then allows a machine to read this value. The machine decrypts of the onion, as will be explained later, and from this computes the candidate list, which it now prints in the left column of the ballot form.

The result is depicted in Figure 5.

The voter now makes her mark(s) on the form in the privacy of the voting booth and the result is exemplified in Figure 6. She then detaches the top page from the bottom and the result is shown in Figure 7. The voter places the page marked *POST* into an envelope through which only the serial number is visible and then leaves the booth carrying the envelope and the top page, which will constitute her receipt. She now presents herself to the vote casting desk and hands over the envelope and receipt. The poll station worker checks that *serial* is the same as the one previously assigned to the voter. Once this is done, the serial number is detached and discarded and the envelope containing the lower

page is placed in the ballot box. The page marked $RETAIN$, which acts like a conventional Prêt à Voter receipt, is scanned, a digital copy posted to the WBB and handed back to the the voter to keep as her *protected receipt*.

The serial number serves a dual purpose here: firstly it counters chain-voting attacks as suggested by Jones [8]. Secondly, it serves to verify that the voter does not retain the lower layer of their ballot form. This is a useful spin-off of the HRPAT mechanism: in the standard Prêt à Voter, there is the possibility of the voter retaining the LH portion of the ballot form, along with her receipt, to prove to a coercer how she voted.

4.2 Cut-and-Choose

Early versions of Prêt à Voter used preprinted ballot forms and so, for the election to be guaranteed accurate and to instill trust in the voters, randomly selected ballot forms are audited before, during and after the election. That is to say they are decrypted and shown to have been correctly printed [2,23]. Such random selection is performed by suitable auditing authorities but may also be supplemented by the voters themselves. One mechanism to provide such a cut-and-choose protocol to the voter while maintaining control on the number of ballots issued to each voter, is to have a double sided form, one side of which (selected at random by the voter) is used to cast the vote and the other is automatically audited [25,26]. However, any such "cut-and-choose" mechanism only allows forms that are not used to be audited.

In the scheme presented here, we add a paper audit trail to Prêt à Voter. As has been described above, the candidate list is printed on the bottom page of the ballot form and this page is placed in a ballot box and provides the human readable paper audit trail. Because of the properties of the relation between the two pages as described in this section, it is possible to audit the printing of the candidate list of any number of forms that were actually used for voting after the close of the election. The device or authority printing the form would thus be caught with a probability proportional to the number of forms audited. Hence the HRPAT method shown in this paper has this further audit application. This auditing mechanism can be used with either pre-printed or on-demand printed forms.

4.3 Generation of the Encrypted Ballot Forms

We describe a distributed, parallel construction of the onion pairs, analogous to the Paillier construction presented in [22]. Suppose that we have L *clerks*. They will be responsible for generating I onion pairs, where each onion pair will carry the same seed/plainext.

We further suppose that we have an ElGamal public key for the tellers PK_T and public keys for the Booths PK_{B_k}, where k indexes the booths. Both of these public keys will have the same modulus. We provide the construction for a single booth key; we simply replicate the construction for other booth keys. Denote the public key of the booth in question as PK_B.

The jth clerk generates I *sub-onion* pairs:

$$\{\theta^T_{j,i}; \theta^B_{j,i}\}$$

Where:

$$\theta^T_{j,i} := \{s_{j,i}, x_{j,i}\}_{PK_T}$$

and

$$\theta^B_{j,i} := \{s_{j,i}, y_{j,i}\}_{PK_B}$$

The first term is an encryption of the j, ith seed under the Teller's public key. The second term is the encryption of the same seed value under the booth's public key. The randomisations x, y, used for these two encryptions should be independent.

All of these sub-onions are all posted to a WBB in cells of an $L \times I$ matrix (L columns, I rows) — one pair in each cell. To audit these, an independent auditing entity chooses for each row a randomly selected subset of the cells in the row, say half. For these selected cells the clerks reveal the s, x and y values. The auditor can check that the encryptions match the posted sub-onion values and that the two seed values are equal for each pair. The auditor can also check that the s values are consistent with the required distribution.

Assuming the posted material passes the audits, the "full" onions are formed by taking the product of the remaining, un-audited pairs row-wise. This step is universally verifiable. Let A_i denote the set of indices of the pairs selected for audit in the ith row. Then the "full" onions for the i th row are computed as:

$$\Theta^T_i := \prod_{j \in \bar{A}} \theta^T_{j,i}$$

$$\Theta^B_i := \prod_{j \in \bar{A}} \theta^B_{j,i}$$

To create the proto-ballots, suppose that we have paper ballots forms that initially just carry index values from I, each form will carry a unique index value. We now introduce two new processes P_1, P_2. P_1 takes a form with index i, looks up Θ^T_i on the WBB, re-encrypts it and prints the result on the RH portion of form. This now constitutes the $\Theta_{R,i}$ for the ballot form. It then covers this with a scratch strip. Once it has finished a batch of these, they are shuffled and passed on to P_2.

P_2 looks up the appropriate Θ^B_i, re-encrypts this and prints the resulting value, $\Theta_{L,i}$on the LH portion of the ballot and covers it with a scratch strip.

We perform audits on a randomly selected subset of the resulting proto-ballots. For the selected ballots, the onions are revealed and P_1 and P_2 are required to prove the re-encryption link back to the onion pair on the WBB. Audited forms are marked are discarded.

Our construction ensures that it would take corrupt booth or access to the paper audit trail and a two-way collusion, of P_1 and P_2, to link the R (receipt) onions to the candidate lists. The index value on the ballots can serve as the serial number, and is removed at the time of casting.

4.4 Anonymising Tabulation

Anonymising tabulation proceeds as for Prêt à Voter 2006. We outline it here for completeness. The encrypted receipts scanned in the polling station are published on the web bulletin board and all voters are able to check that their receipts appear there. When all tellers are satisfied that the election has ended and all electoral rules have been followed they start the decryption process, which is shown in Table 1. The first teller, T_1, takes all encrypted receipts and injects the voter's choice(s) into the $onion_R$, using the homomorphic properties of exponential ElGamal. We call the onion with the injected choice(s) $onion_I$. Suppose:

$$onion_R = (\alpha^r, \alpha^s \cdot \beta^r) \pmod{p}$$

Then:

$$onion_I := (\alpha^r, \alpha^v \cdot \alpha^m \cdot \beta^r) \pmod{p}$$

The index number v indicates the position of the X on the receipt. In effect, we are multiplying $onion_R$ by the encryption of v with randomisation $r = 0$. The result is:

$$onion_I = \{v + s, t\}_{PK_T}$$

Thus, the I onion is the encryption of the v index plus the seed value. The offset ϕ of the candidate list printed on the ballot form is computed as $\phi := s \ (mod\ n)$, where n is the number of candidates. The candidate order is cyclically shifted upwards from the canonical ordering by ϕ. Thus, $v + s \ (mod\ n)$ gives the index of the candidate chosen by the voter in the canonical numbering of the candidates.

No mixing is performed at this step: the I and R onions are posted side-by-side on the WBB. That each $onion_I$ is correctly formed w.r.t. $onion_R$ is thus universally verifiable.

Table 1. Decryption of the encrypted receipts

$onion_R$	Inject choices	$onion_I$	Re-encryption mix network	$onion_{I_n}$	Decryption	Plaintext vote
O_{R_2}	$\Rightarrow$	O_{I_2}		O_{I_5}	$\Rightarrow$	V_5
O_{R_1}	$\Rightarrow$	O_{I_1}		O_{I_2}	$\Rightarrow$	V_2
O_{R_4}	$\Rightarrow$	O_{I_4}		O_{I_3}	$\Rightarrow$	V_3
O_{R_5}	$\Rightarrow$	O_{I_5}		O_{I_4}	$\Rightarrow$	V_4
O_{R_3}	$\Rightarrow$	O_{I_3}		O_{I_1}	$\Rightarrow$	V_1

We now perform a sequence of re-encryption mixes, performed by a set of mix tellers. Each mix teller takes the full batch of $onion_I$s, re-encrypts each onion, shuffles the batch and outputs to the next mix teller. The output batch from each teller is published onto the web bulletin board. The last output batch we call $onion_{I_n}$.

When all mix tellers have performed their re-encryption mixes, the independent auditors confirm that the mixes have all been performed correctly. This might be done using partial random checking [7], or perhaps Neff's proofs of ElGamal shuffles [15]. If the auditors confirm that the mixes are correct, we can proceed to the decryption stage. If problems are identified with the mixes, corrective actions can be taken. Thus, for example, if one of the mix tellers is identified as having cheated, it can be removed and replaced. The mixes can be re-computed from the point onwards and re-audited. We might routinely re-run the mixes and audits in any case for additional assurance.

Once we are happy that the mixes have been performed correctly, a threshold set of the decryption tellers take over and cooperate to decrypt each $onion_{I_n}$. No mixing is required at this stage and each step of the decryption can be accompanied with a ZK proof of correct (partial) decryption. The final, fully decrypted values can be translated into the corresponding candidate values using:

$$candidate_i = (s + v) \ (mod \ n))$$

Such re-encryption mixes are known to provide anonymity against a passive attacker. Against an active attacker, who might have some capability to inject or alter terms entered into the mix, we have to guard against ballot doubling attacks: to identify a particular voter's choice, he injects a term that is a re-randomisation of the voter's receipt. If unchecked, this will result in two decrypted receipts with the same adjusted seed value. We will in any case have procedures in place to guard against ballot stuffing that will help counter such dangers. An additional measure is to run (threshold) plaintext equivalence checks against the terms in the mix prior to decryption, see [9].

4.5 Audit of the Paper Trail

We now have a number of possible strategies for auditing the election. One scenario is to perform a full, manual recount of the election using the HRPAT and simply compare this with the cryptographic count. In practice, due to inevitable errors with manual counting, this will differ from the electronic count, even if the latter is exact and correct. If the difference is small and well within the winning margin, this could probably be disregarded.

An alternative is to take a random subset of the HRPAT ballots and, for each of these forms, the auditor requires the appropriate booth to decrypt the onion and so reveal the seed s. The tellers are required to provide ZK proofs of the correctness of their decryption steps. From the seed value s it computes the candidate order and checks that this agrees with the list printed on the ballot.

This audit serves to catch any cheating by booths that might not have been detected earlier during any cut-and-choose audits. The advantage of these audits

is that we are checking the candidate orders on ballot forms actually used by the voters to cast their votes rather than just on unused ballots.

We can now perform some checks of correspondence between the paper audit trial and the decrypted ballots posted from the tabulating mixes. For each selected paper audit ballot, the auditor now computes the adjusted seed value:

$$\bar{s} := v + s$$

It should now be able to find a matching value amongst the decrypted outputs of the tabulation process on the WBB. Failure to find a matching value casts doubt on the conduct of the election. If the auditor finds an adjusted seed value in the tabulation that differs slightly (i.e. by less than n) from the closest seed value from the paper audit trial this may be indicative of corruption. This might be due to some manipulation of index values in the paper audit trial or the electronic records. Further investigation would now be required, firstly to establish that the paper ballot has not been manipulated.

For ballots selected for audit for which the above check fails, we can perform a diagnostic check: we perform PET checks of the paper ballot onion against the posted receipt onions. If a match were found, and the corresponding index posted against this onion on the WBB agrees with the index of the paper copy, this would indicate that this receipt had been corrupted in the mix/tabulation phase not detected.

We can also compute amended onions from the paper audit trail by folding the index into the LH onion in the same way that we formed the I onions. We refer to these as J onions.These J onions will have different randomisations from the corresponding I onions computed previously. However, as long as all computations have been performed correctly, the sets of $onion_I$s, $onion_{I_n}$s and $onion_J$s contain the same plaintexts. In other words, The J onions should be related to the I by a re-encryptions and shuffles. We could test this hypothesis by performing a full PET matching of the I and J onions or, perhaps more realistically, performing some spot checks on a random selection.

5 Analysis

Rather than attempt a full analysis of the present scheme, we will discuss the respects in which it differs from Prêt à Voter 2006. In terms of the accuracy guarantees we will see that this scheme provides stronger guarantees that Prêt à Voter 2006, assuming the integrity of the paper audit trial. If the paper audit trial is vulnerable to manipulation, then arguably the HRPAT mechanism could undermine the assurance of accuracy of the original scheme.

Assuming the integrity of the paper audit trail for the moment, the additional auditing possibilities introduced by this HRPAT mechanism means that it will be significantly harder to violate accuracy in an undetectable way. For example, the fact that all actually voted ballot forms can be audited for correct construction means that is essentially impossible for votes to be incorrectly encoded in receipts undetected. In previous versions of Prêt à Voter, and indeed

similar schemes, these checks are probabilistic and require assumptions of lack of collusion between ballot creating processes and auditing processes.

5.1 Linking the Receipt Onions to the Candidate Lists

The fact that in this scheme, the ballot forms carry linked onions on both portions does create potential threats against ballot privacy. Thus, for example, if the adversary is able to link the L and R onions for a ballot form and is able to access the paper audit trail, then he will be able to compromise the secrecy of that voter's ballot. This could be achieved with the collusion of the P_1 and P_2 processes. It is of course difficult to gauge whether this is a good trade-off, and this judgement will probably vary according to circumstance, perceived threats etc.

The link between LH and RH onions is cryptographically protected and cannot be directly re-established without access to a threshold set of teller's keys. However, there is a danger that if booth keys are compromised, it may be possible to obtain the seeds for some ballots and link these to the decrypted values posted on the WBB. The coercer still has to link the HRPAT ballot to the voter who used it. He can do this if he can establish the link between the two onions. However, our construction ensures that it would require a collusion of the P_1 and P_2 processes to reveal these links.

We see that the HRPAT mechanism does introduce some threats against ballot privacy that are absent in conventional Prêt à Voter. However, we have striven to ensure that the threshold to exploit such vulnerabilities is quite high. It is a delicate trade-off to establish whether the introduction of such vulnerabilities is justified by the added assurance and confidence resulting from the HRPAT mechanism.

5.2 Voter Choices Differ between Pages

As the voter makes her marks on the form in the privacy of the booth, it is possible for a malicious or coerced voter to introduce different marks on the two pages in order to try to introduce inconsistencies between the paper and electronic records and so seek to discredit the election. To resolve this and to

Table 2. Another re-encryption mix of $onion_L$

$onion_L$	Re-encryption mix network	$onion_M$
O_{L_2}		O_{M_2}
O_{L_3}		O_{M_1}
O_{L_1}		O_{M_4}
O_{L_5}		O_{M_5}
O_{L_4}		O_{M_3}
All tellers		

prove that the marks were made differently on each sheet by the voter the tellers can take the list of $onion_L$s and run them through a re-encryption mix to form a list of $onion_M$s, as shown in Table 2. It is then possible to use the PET strategy to prove which $onion_M$ contains the same information as the $onion_L$, the extension of which is that the bottom page is valid but the voter's mark does not match. If the tellers, when prompted, find that $onion_L$ with the voter's choice V_{bottom} does not have the same plaintext as $onion_R$ with the choice V_{top} injected then they prove that $onion_L$ has the same plaintext as $onion_M$ to show that the marks are different on each of the pages.

6 Conclusions

We have presented a mechanism that can be incorporated in Prêt à Voter to generate a plaintext paper audit trail. This has a number of benefits: firstly there is the confidence building effect of having a paper audit trail as a safety-net. Secondly it provides a number of additional auditing possibilities: spot checks of correspondence between the paper ballots and decrypted ballots as well as checks on the correctness of the candidate order printed on the ballots by the booth devices. Note that these checks are applied directly to the candidate orders used by the voters, rather than on unused, audited forms as with the cut-and-choose audits.

A further benefit is to provide a mechanism to ensure that voters do submit the portion of the ballot that carries the candidate order, so countering dangers of voters attempting to smuggle these out to prove prove their vote to a coercer.

On the other hand, the HRPAT mechanism presented here does introduce some threats against ballot privacy that are not present in conventional Prêt à Voter. Evaluating this trade-off requires more systematic ways to evaluate voting systems than exist at present. Besides, it is likely that such trade-offs will be highly dependent on the context. For example, in the UK, it is required by law to maintain a link between voter id and ballots forms. Thus, in the UK, a mechanism along the lines proposed would not only be acceptable but would probably be required.

Another issue to be borne in mind, is that the paper audit trail may be vulnerable to manipulation. This is true of conventional pen and paper voting, but here it may be particularly problematic as such manipulation may serve to cast doubt on a completely valid electronic count. Again, this is a delicate trade-off against the comfort factor of having a paper audit trail fall-back.

Acknowledgements

The authors would like to thank Ron Rivest for suggesting enhancing Prêt à Voter with a human-readable paper audit trail. We would also like to thank Steve Schneider, Jacques Traore, Raphael Yahalom and Joson Xia.

References

1. Chaum, D., Ryan, P.Y.A., Schneider, S.: A practical voter-verifiable election scheme. In: de Capitani di Vimercati, S., Syverson, P.F., Gollmann, D. (eds.) ESORICS 2005. LNCS, vol. 3679, pp. 118–139. Springer, Heidelberg (2005)
2. Chaum, D., Ryan, P.Y.A., Schneider, S.: A practical voter-verifiable election scheme. Technical Report, University of Newcastle, CS-TR:880 (2005)
3. Cohen, S.: Auditing technology for electronic voting machines. Ph.D, MIT Cambridge (July 2005), `http://www.vote.caltech.edu/theses/cohen-thesis_5-05.pdf`
4. ElGamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. IEEE Transactions on IT 31(4), 467–472 (1985)
5. Fisher, K., Carback, R., Sherman, T.: Punchscan: Introduction and system definition of a high-integrity election system. In: Pre-Proceedings IAVoSS Workshop on Trustworthy Elections, pp. 19–29 (2006)
6. Heather, J., Lundin, D.: The append-only web bulletin board. Technical Report at the University of Surrey CS-08-02 (2008)
7. Jakobsson, M., Juels, A., Rivest, R.: Making mix nets robust for electronic voting by randomized partial checking. In: USENIX Security Symposium, pp. 339–353 (2002)
8. Jones, D.W.: A brief illustrated history of voting (2003), `http://www.cs.uiowa.edu/~jones/voting/pictures`
9. Juels, A., Catalano, D., Jakobsson, M.: Coercion-resistant electronic elections. In: Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society, pp. 61–70 (2005)
10. Kohno, T., Stubblefield, A., Rubin, A.D., Wallach, D.S.: Analysis of an electronic voting system. In: Symposium on Security and Privacy. IEEE, Los Alamitos (2004)
11. Lundin, D., Treharne, H., Ryan, P.Y.A., Schneider, S., Heather, J.: Distributed creation of the ballot form in prêt à voter using an element of visual encryption. In: Proceedings of Workshop On Trustworthy Elections (WOTE 2006), pp. 119–125 (2006)
12. Lundin, D., Treharne, H., Ryan, P.Y.A., Schneider, S., Heather, J., Xia, Z.: Tear and destroy: chain voting and destruction problems shared by prêt à voter and punchscan and a solution using visual encryption. In: Proceedings of Workshop on Frontiers in Electronic Elections (FEE 2006) (2006)
13. Mercuri, R.: A better ballot box? IEEE Spectrum Online (October 2002)
14. Neff, A.: Practical high certainty intent verification for encrypted votes (2004), `http://www.votehere.net/documentation/vhti`
15. Neff, C.A.: A verifiable secret shuffle and its application to e-voting. In: Proceedings of the eighth ACM conference on Computer and Communications Security (CSS 2001), pp. 116–125 (2001)
16. Pedersen, T.: A threshold cryptosystem without a trusted party. In: Davies, D.W. (ed.) EUROCRYPT 1991. LNCS, vol. 547, pp. 522–526. Springer, Heidelberg (1991)
17. Randell, B., Ryan, P.Y.A.: Voting technologies and trust. IEEE Security & Privacy (November 2006)
18. Rivest, R.L.: The three ballot voting system. MIT Press, Cambridge (2006), `theory.lcs.mit.edu/~rivest/Rivest-TheThreeBallotVotingSystem.pdf`
19. Rivest, R.L., Wack, J.P.: On the notion of software independence in voting systems. Philosophical Transactions of the Royal Society (to appear, 2008)

20. Ryan, P.Y.A.: Prêt à voter with human readable paper audit trail. Technical Report of University of Newcastle, CS-TR:1038 (2007)
21. Ryan, P.Y.A.: Prêt à voter with paillier encryption. Technical Report of University of Newcastle, CS-TR:1014 (2007)
22. Ryan, P.Y.A.: Prêt à voter with paillier encryption. In: Mathematical and Computer Modelling, Mathematical Modeling of Voting Systems and Elections: Theory and Application (2008)
23. Ryan, P.Y.A., Peacock, T.: Prêt à voter: a system perspective. Technical Report of University of Newcastle, CS-TR:929 (2005)
24. Ryan, P.Y.A., Peacock, T.: Putting the human back in voting protocols. Technical Report of University of Newcastle, CS-TR:972 (2006)
25. Ryan, P.Y.A., Peacock, T.: Threat analysis of cryptographic election schemes. Technical Report of University of Newcastle, CS-TR:971 (2006)
26. Ryan, P.Y.A., Schneider, S.: Prêt à voter with re-encryption mixes. In: Gollmann, D., Meier, J., Sabelfeld, A. (eds.) ESORICS 2006. LNCS, vol. 4189, pp. 313–326. Springer, Heidelberg (2006)
27. Ryan, P.Y.A.: A variant of the chaum voting scheme. Technical Report CS-TR-864, University of Newcastle upon Tyne (2004)
28. Ryan, P.Y.A.: A variant of the chaum voting scheme. In: Proceedings of the Workshop on Issues in the Theory of Security, pp. 81–88. ACM, New York (2005)
29. Ryan, P.Y.A.: Verified encrypted paper audit trails. Technical Report Newcastle Tech. Report 966, June 2006, University of Newcastle upon Tyne (2006)
30. Xia, Z., Schneider, S., Heather, J., Ryan, P.Y.A., Lundin, D., Peel, R., Howard, P.: Prêt à voter: all in one. In: Proceedings of Workshop On Trustworthy Elections (WOTE 2007) (2007)