Improving the Farnel, Threeballot, and Randell-Ryan Voting Schemes

Roberto Araújo[1] and Peter Y. A. Ryan[2]

[1] Department of Computer Science, TU-Darmstadt
Hochschulstrasse 10, D-64289 Darmstadt, Germany
rsa@cdc.informatik.tu-darmstadt.de
[2] Centre for Software Reliability, Newcastle University
Newcastle upon Tyne NE1 7RU UK
peter.ryan@ncl.ac.uk

Abstract

A number of recent voting schemes provide the property of voter verifiability: voters can confirm that their votes are accurately counted in the tally. The Farnel type voting schemes are based on the observation that to achieve voter-verifiability it is not necessary for the voter to carry away a receipt corresponding to their own vote. The Farnel approach then is to provide voters, when they cast their vote, with copies of receipts of one or more randomly selected, previous cast votes. This idea has a number of attractive features: ballot secrecy is achieved *up front* and does not have to be provided by anonymising mixes etc during tabulation. In fact, plaintext receipts can be used in contrast to the encrypted receipts of many other voter-verifiable schemes. Furthermore, any fears that voters might have that their vote is not truly concealed in an encrypted receipt are mitigated. The Farnel mechanism also mitigates randomization style attacks. In this paper we explore some enhancements to the original Farnel scheme and ways that the Farnel concept can be combined with some existing voter-verifiable schemes, namely Prêt-à-Voter, ThreeBallot, and Randell-Ryan.

1 Introduction

Voter verifiability is a novel security feature provided by several recent voting systems, such as Prêt-à-Voter [11, 3] and Punch Scan [7]. It allows voters to verify that their votes are accurately counted by means of *protected receipts* and so gives more confidence to the election process. The voters, though, cannot use their receipts to compromise their privacy, even if they are prepared to cooperate with the coercer.

High-assurance voting systems typically rely on cryptography to achieve security and to implement voter verifiability. Such technology makes the security of modern systems comparable or even better than traditional paper-based elections. However, systems that employ cryptography are not easily grasped by the average voter and so voters need to rely on the assurances of experts.

With the goal of making such schemes more understandable, Randell-Ryan [8], Rivest [9, 10], and Araújo et al. [2], introduced voter verifiable schemes that do not rely on cryptography. These schemes are simple and can be more easily understood by the voters. However, they do not achieve the same levels of assurance as the cryptographic systems. In the scheme proposed in [9], for example, the ballot secrecy is not perfect and it may reveal statistical indications of voting results before the voting end.

Although cryptographic voting systems may not be fully understandable by end users, it is difficult to design a secure system without cryptography. Thus, in this paper we introduce improvements to the scheme of Araújo et al. [2], and explore ways to incorporate the Farnel [1] concept into Rivest [9] as well as into Randell-Ryan [8] voting schemes. In addition, we present a new scheme based on combining Farnel [2, 4] with a Prêt-à-Voter style encoding of receipts. Our proposals make minimal use of cryptography to achieve security while keeping simplicity.

This paper is organized as follows: in the next Section we describe some basic elements required by our proposals. In Section 3 we introduce a new ballot form to the scheme of Araújo et al. Then, in Section 4, we show a new scheme based on Farnel that employs only one ballot box. In Section 5 we present and discuss improvements to Threeballot scheme. After that, we introduce improvements to the scheme of Ryan-Randell in Section 6. Finally, we present our conclusions in Section 7.

2 Preliminaries

We present here the basic elements of the Farnel approach. The Farnel type voting schemes, [2, 4] are based on the observation that to achieve voter-verifiable it is not necessary for the voter to carry away a receipt corresponding to their own vote. The Farnel approach then is to provide voters, when they cast their votes, with copies of receipts of one or more randomly selected, previous cast votes.

This idea has a number of attractive features: ballot secrecy is achieved *up front* and does not have to be provided by anonymising mixes etc during tabulation. In fact, plaintext receipts can be used in contrast to the encrypted receipts of many other voter-verifiable schemes, e.g. [11]. Furthermore, any fears that voters might have that their vote is not truly concealed in an encrypted receipt is mitigated. The Farnel mechanism also mitigates randomization style attacks.

2.1 The Farnel Ballot Box

The Farnel ballot box is a special kind of ballot box that was introduced by Custódio et al. [4, 5]. This ballot box performs differently from a conventional one. It has a shuffle mechanism and is initialized with elements (e.g. votes). After receiving elements from voters, it gives them elements that correspond to randomly selected previously cast. Recently, Araújo et al. [2] improved the Farnel box by adding it two more properties. That is, in addition to shuffle elements, the box also copies some of them and removes scratch surfaces.

[1] Farnel means basket in Portuguese.

Following these previous works, we describe the enhanced Farnel box as follows: it is a box that has mechanisms to remove scratch surfaces, and to shuffle and to copy elements (e.g. votes) in a memoryless way. The box has an initial set of elements cast before the voting. At the time of voting, it is able to receive an input, to add it to its initial set, to shuffle the new set, to copy one or more randomly selected items from its set, and to output the copies. For convenience, here we call Farnel the enhanced Farnel box and original Farnel the box introduced by Custódio et al.

Although the requisites of Farnel seems difficult to implement, a tombola (i.e. a raffle drum) normally used in lottery games to shuffle tickets could form the basis of an implementation of the Farnel box. A tombola has a slit to receive slips and a mechanism to spinning it, but it could be adapted to remove scratch surfaces. This way, after the tombola receives an input and shuffles its contents, the authorities could open it, take some elements from it and copy them, and deposit the original elements back into the tombola. These procedures would be observed by voters and by helper organizations. We believe, however, that the Farnel box could be also constructed by adapting a scanner in the slit and by adding a small printer in the tombola. This would avoid the contact of the authorities with the original elements in the tombola.

We can specify the Farnel ballot box in the process algebra CSP formally as follows:

Let *Init* denote the initial set of dummy ballots with which the box is initialized. Let l denote the number of number of receipts to be output to each voter when they cast their votes and *ballots* the set of all possible ballots (or receipts). Then the Farnel box will start in state $Farnel(Init)$ and its subsequent behavior is defined recursively as:

$$Farnel_l(X) := cast?b : Ballots \rightarrow \Box receipt!r : \wp_l(X) \rightarrow Farnel_l(X \cup \{b\})$$

We have used the notation $\wp_l(X)$ to denote set of subsets of X of cardinality l.

Thus, the Farnel ballot box is parametrised by the integer l and it's initialization *Init*. At any point, the box can accept a ballot b, after which it outputs a set ballots of size l chosen at random from it's current set X. After this, the new ballot is added to X and the box is ready to receive the next ballot.

2.2 The Initialization Process

Most of schemes described in the next Sections employs a Farnel box. As this ballot box requires an initialization and this process is almost the same in the schemes, we present a generalization of it here and refer to this Section when necessary.

The initialization process takes place before the election and is performed by the authorities in a public session. The main objective is to cast a predefined number of votes (or receipts) into the Farnel ballot box and to publish the number of elements cast per option on the bulletin board.

For some of the schemes that we describe here, it is necessary to ensure that ballots cast during the initialization are well-formed in some way. This will typically involve some form of random auditing. Thus, for example, we might require that $2x$ blank ballots be created beforehand. The authorities perform the following steps to initialize the ballot box:

1. Select x blank ballots at random and audit them as necessary. Ballots audited are discarded;
2. Mark the other x unaudited blank ballots according to the number of votes per option specified in advance;
3. Cast the x marked ballots (or receipts) into the Farnel box and publish the number of elements cast on the bulletin board.

Notice that in schemes which employ a conventional and a Farnel box, the conventional box is initialized with votes and the Farnel is initialized with the corresponding receipts. Also, for schemes using plaintext ballots, the auditing for well-formedness is not necessary and would be omitted.

In order to prevent manipulation, the initialization process should be scrutinized by help organizations. They should check the ballot box is empty before it is initialized as well as verify all procedures above are performed correctly. Further, the ballot box should be sealed and continually supervised by third parties after the initialization. The seal is removed when voting starts.

The initialization of the Farnel box is necessary mainly for ensuring the anonymity of the early voters. As the Farnel receives an input from each voter and outputs copies of random elements, it must have an initial set of elements to choose from. Otherwise, after receiving the early inputs, the Farnel would not have enough elements to select at random and make the copies.

2.2.1 Initialization of the Farnel box with void ballots

Where we are using encrypted receipts we have an alternative way to initialize the Farnel box: we include a *void* option on the ballots and initialize the box with ballots representing votes for the void option. This has the advantage that we do not have to keep a log of the actually votes cast for each candidate during initialization. We do need a robust mechanism to ensure that all initializing votes are cast for void, but it seems likely that this is easier to enforce than maintaining a record of an initial tally. We can use this approach for the Prêt-à-Voter and ThreeBallot style ballots, but not where plaintext receipts are used.

2.3 The Parameters of the Farnel box

The Farnel box is initialized with a number of elements (votes or receipts) before voting starts and outputs copies of its elements during the voting, as described. The initial elements ensure the voter's anonymity while the copies are handed to the voter as her receipt. The number of initial elements as well as the number of copies given to each voter compose the parameters of the box.

Because the Farnel box outputs copies of its contents, it may reveal information that affects the voter's anonymity. The information remains concealed until the tally if encrypted receipts are employed. However, it would be revealed after the receipts be decoded (unless an anonymising mix tabulation is employed). The quantity of information revealed increases according to the number of copies output by the box. The more copies the Farnel box outputs, the more information about its elements it reveals. Consequently, although more elements can be verified by the voters, the information revealed may be sufficient to compromise the anonymity.

In order to preserve the anonymity of the voters, the initial elements and the voters elements cannot be distinguished through the copies output by the Farnel box. The number of initial elements is fundamental for guaranteeing this. As the Farnel box output elements for each voter, the elements of the early voters have more chance to be output. Hence, these elements may be distinguished from other elements. Depending on the number of initial elements, however, the chance of distinction may be negligible as the initial elements may also be output.

In order to achieve verifiability while maintaining anonymity, the number of initial elements and the number of copies should be defined such that:

1. The voter anonymity is preserved even if the Farnel box is able to output a copy of her element;

2. An individual receipt or a set of them do not provide enough information to distinguish initial elements from voters elements;

3. The number of copies of elements in all receipts is sufficient to detect accuracy problems with an acceptable probability (i.e. the probability that the corruption of any given ballot is detected be at least 50%).

We require that the voter should not be able to obtain any information other than the option she marked when casting her element.

Taking into account these requisites, we have a number of possible strategies for the initialization of the box: ballots marked at random (with the totals carefully recorded), a predetermined number of votes per option, votes for a void option, or a combination of these methods. If we adopt an initialization with votes for void, and do not include an anonymising mix tabulation, we must include a minimal number of votes for the other options. Otherwise, the first voter may vote and receive a copy of her own vote as receipt.

In principle, an initialization with at least one element for each option may be sufficient to preserve the voter's anonymity. However, depending on the number of voters and on the number of options, the voter anonymity may not be preserved. For example, supposing a voting with two options, one initial element for each option, and two copies per voter as receipt, after the first voter casts her vote, she may receive her element and an initial element for the same option from the Farnel box.

Note that in the specification of the Farnel box presented before, the box is not able to output the element it receives.

3 A New Ballot Design for the Farnel Variant

The Farnel scheme was originally proposed by Custódio et al. [4, 5] (see [2] for a description). The scheme employs an original Farnel ballot box and relies on physical signatures. However, it is not voter-verifiable and requires trustworthy authorities. Recently, Araújo et al. [2] introduced a variant of the Farnel scheme. In contrast to the original version, the scheme is voter-verifiable and does not employ signatures. It relies, though, on trustworthy talliers to tabulate the votes.

In this Section we introduce a new ballot design for the Farnel variant proposed by Araújo et al. [2]. Our proposal aims at reducing the need to trust the talliers.

3.1 An overview of the Farnel variant

The scheme employs a ballot form composed of two halves that are linked by a unique ID and that are separated by perforations. More specifically, the ballot has an options half composed of voting options as well as an ID and an ID half that has the same ID of the options half (see Figure 1). These IDs are covered by scratch surfaces.

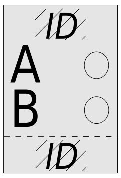

Figure 1: The ballot form of the Farnel variant.

Besides the unusual ballot form, the scheme depends on two ballot boxes. One of them is conventional. The other is a Farnel box, such as described in Section 2.1. These boxes are initialized before the voting. That is, the conventional box receives votes (i.e. marked option halves) and the Farnel box receives the ID halves (i.e. receipts) corresponding to the votes. The scratch surfaces in the halves are detached during the initialization and at the end the number of votes cast are published on a bulletin board.

At time of voting, the voter receives a blank ballot and detaches its scratch surfaces. She then compares the IDs on the halves and if they match, she marks her option. After that, she separates the two halves of her ballot, casts the option half into the conventional box, and casts the other half into the Farnel box. Upon receiving the half, the Farnel box shuffles its ID halves and copies a set of them as receipt to the voter. As alternative to avoid comparison of IDs, the scheme may have an auditing to check ballots before the voter receives her blank ballot and require the voter to cast her vote without removing the scratches. The Farnel box now removes the scratch of the half it receives.

After the voting, the authorities publish the content of both ballot boxes on the bulletin board and count all votes from the conventional box. The initial votes are then subtracted from the total of votes to obtain the results.

In order to verify the votes, voters and third parties compare the ID halves with the IDs in the options halves. The voters can also match the IDs on their receipts to the options halves on the board.

3.1.1 Drawback

Due the receipt style employed, the proposal requires trustworthy talliers. These authorities should supervise the votes strictly after opening the ballot boxes. On the contrary, an adversary can compromise the voting results as follows.

According to the scheme, the two halves of all ballots are published after the voting. This way, they can be compared to verify the exactness of the voting results. Before publishing the options halves, though, an adversary could replace votes (i.e. a marked option half) by a new one marked to a different option, but that contains the same ID of the replaced vote. This substitution would not be detected by voters and third parties as they only compare IDs.

3.2 Combining the Farnel variant and Prêt-à-Voter

The main problem of the receipt used in Farnel variant is that it does not depend on the vote option. This way, an adversary can easily replace votes without being detected. In order to detect such problems, a receipt should contain some information related to the option selected. However, this information should not reveal the option itself before voting closes and should still be able to detect replacement of votes. Otherwise, the receipt can leak statistical information about the voting results as the Threeballot scheme [9, 10] (see Section 5.1). We introduce now a new ballot design for Farnel variant that satisfies these requirements.

Our ballot form is based on Prêt-à-Voter [11, 3] ballot and on some ideas of Randell-Ryan scheme [8] and of Scratch-and-vote [1]. It is formed of two pages that are overlaid initially. The top page has a list of voting options in a random order and each option is associated to a bubble to select it. The top page also contains a commitment to the list of options and its respective decommitment value. The bottom page contains the *same* bubbles and the *same* commitment of the top page. The commitment printed on both pages as well as the value to open it in the top page are covered by scratch surfaces. A carbon mechanism transfers the selections made by the voter on the top page to the bottom page (see Figure 2 for an example of this ballot form).

Formally, the new ballot form is described as follows. Let C be a set of options available, π_C the permutation of C, H a secure hash function used here as commitment, and r a random number from a large (key) space. π_C, $H(\pi_C, r)$, r, and bubbles to select an option compose the top page. The bottom page contains only $H(\pi_C, r)$ and the bubbles in same position of the top page.

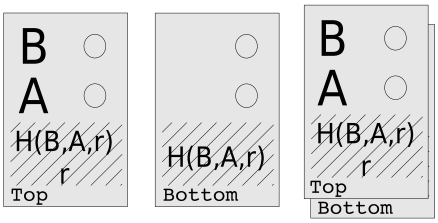

Figure 2: The proposed ballot form for the Farnel variant. The two pages of the ballot and the pages overlaid.

The votes are tabulated from the top pages and the receipts are made from the bottom pages (without the scratch surfaces). Because each bottom page contains the *same* selections of its corresponding top page and also has the commitment to the options on the top page, an adversary cannot replace a top page by another with a different permutation or with a selection for a different option, without being detected. Moreover, since the bottom page does not include the option selected, an adversary cannot use receipts to obtain indication of the results before voting closes. Thus, the new ballot form satisfies our requirements.

3.3 New steps for voting and tallying phases

Due the modification of the ballot form, the conventional box and the Farnel box are now initialized with marked top pages and with bottom pages respectively. In addition, the voting as well as the tallying steps in the original scheme need to be adapted. Now, the voter performs the following steps to vote:

1. (Selecting the option) The voter marks her vote on the top page and the mark is transfered to the bottom page;
2. (Verifying the ballot) She then puts the ballot in a special envelope, which has transparent borders and has a window to show just the scratch, and hands the envelope to the authorities. They verify the scratch on the top page is intact and the voter did not separate the two pages;
3. (Casting the top page) The voter separates both pages and casts the top page into the conventional ballot box;
4. (Obtaining the receipt) The bottom page is cast into the Farnel box that outputs random bottom pages as receipt.

As the Farnel variant scheme, the contents of the two ballot boxes are published on a bulletin board in the tallying phase. Now, the scratch surface on the top pages should be removed before publishing the ballots and the commitments should be decommited to verify the ballots: the random number and the options on the top page are hashed together and the resulting hash is compared with the hash on the ballot.

Even using the special envelope with transparent boards, it is difficult to ensure the ballot pages were not separated before the authorities verify it. This way, a malicious voter could attempt to discredit a voting by marking different options on the two pages. In order to counter this problem, a physical mechanism to prevent the voter to separate the pages could be used, such as the folder with key used in Punch Scan [7]. Another solution requires the modification of the ballot form. Instead of using a carbon mechanism to transfer the marks, the top page would have holes à la Punch Scan and the voter would use a bingo dauber to select her option. Each role, however, would have a different pattern. Thus, if the voter marks a pattern on the top page, she could not mark a different pattern on the bottom page as the patterns would not match on both pages.

The procedure of verifying the marks on the ballots could expose the voter option if the authorities are able to identify the list of options. In the next Section we introduce a scheme in which the authorities just need to verify the marks in one page.

4 Single Box Farnel Scheme

The design presented above is awkward in several respects: it requires two ballots boxes and the vote casting procedure is rather complicated and vulnerable to certain threats. We present here an improved version of the Farnel variant that requires just one ballot box and uses a simpler vote casting procedure.

4.1 Requisites

4.1.1 The ballot form

As the form presented in Section 3.2, the ballot here has two pages that are initially overlaid. The top page, though, contains only the candidates in random order along with bubbles to select them. The bottom page contains a commitment that is printed on the left of the page and the same bubbles of the top page. It also has a decommitment value that should be printed in the middle of the page and an index value that indicates the permutation used in the top page; this index helps the authorities to identify the permutation in the tallying process. The commitment is covered by a scratch apart from the index and the decommitment.

More formally, let C be a set of options available, I a set of positive integers used as index, π_C the permutation of C, H a secure hash function used as commitment, i a unique number in I, and r a random number from a large (key) space. The top page has π_C and the bubbles to select the options. The bottom page has $H(\pi_C, r)$, r, i, and the same bubbles of the top page. Figure 3 illustrates the ballot form required by the scheme.

The list of possible permutations for all ballots as well as the index of each permutation are published on the bulletin board before the voting.

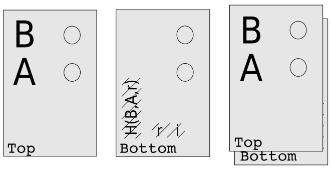

Figure 3: A ballot form in the single box Farnel scheme.

4.1.2 The ballot box

The scheme employs just a Farnel ballot box that is initialized (see Section 2) with marked bottom pages before the voting starts; the corresponding top pages are destroyed.

4.2 The Scheme

Before the Voting As the Farnel variant, we define a number of copies l that each voter receives as receipts and initialize the Farnel ballot box with a number of dummy votes (Section 2 details this process).

For this initialization phase, as well as for the voting phase, we require an auditing process. The audit is necessary to detect malformed ballots and is performed as follows: the authorities selects a set of ballots at random, separates the two pages of each ballot, and detaches their scratch surfaces. In order to verify a ballot, the authorities hash the options on the top page (π_C) along with the random number (r) on the bottom page and compares the result with the hash ($H(\pi_C, r)$) also on the bottom page. Moreover, they verify the randomization on the top page and the randomization indicated by i match. In the voting phase, external organizations can help the voter to audit votes in the same way, that is, the voter selects some blank ballots at random and hands them to the organizations that verify the commitments on the ballots.

Voting The voting authorities hand a blank ballot to the voter in a sealed envelope after verifying her eligibility. The voter can either use the blank ballot to vote or ask the authorities to audit it. In the latter case, the authorities detach the scratch surfaces on the ballot and check the commitment (as before) through a computer. This procedure is watched by the voter and can be also performed by help organizations that would employ their own computers. Assuming that the ballot is verified as well-formed, it is discarded and the authorities hand a new blank ballot to the voter. In principle, we could allow the voter to opt to audit a number of ballots before accepting one to use to cast her vote. If any ballot fails the audit checks then recovery mechanisms will need to be invoked. Discussion of these is beyond the scope of this paper though.

To cast her vote, the voter now performs the following steps to vote (see also Figure 4):

1. (Selecting the option) The voter chooses her option on the ballot form and marks the corresponding bubble;

2. (Casting the ballot) She separates the two pages of her ballot and destroys the top page by means of a paper shredder. She then shows the bottom page to the officials who verify the two scratch surfaces are intact. If the scratch surfaces are entire, the voter casts the page into the Farnel box in the presence of the officials and observers;

3. (Obtaining the receipt) After receiving the bottom page, the Farnel box removes the scratch surface that covers *only* the commitment value on the left side, shuffles its set of bottom pages, and copies l of them. The copies are held by the voter as her receipt.

Recovering and Tallying the votes In order to tally the votes, the authorities open the Farnel box, detach the scratch surfaces on all ballots, and publish the ballots on the bulletin board. Voters can, as usual, visit the bulletin board and confirm that their receipts appear accurately, and complain if not. Then, the authorities start the process to recover the votes. In this process, they compare the index on the ballot with the index on the bulletin board to identify the permutation of the options; remember that the permutations as well as their indexes were previously published. From the permutation identified and the mark on the ballot, the authorities determine the option chosen by the voter. After recovering the votes, the authorities open all commitments using

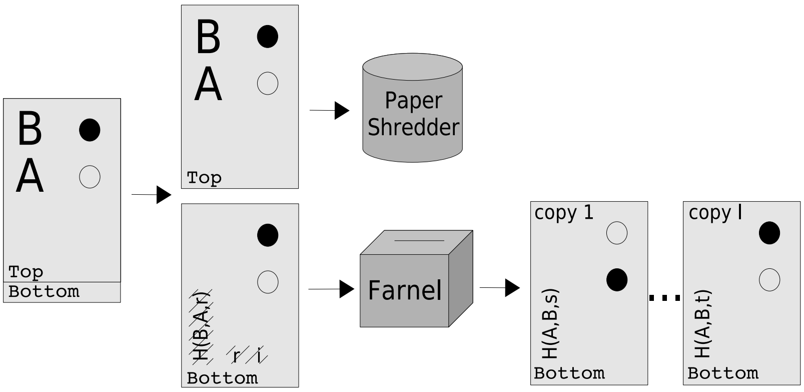

Figure 4: The voting steps of the single box Farnel scheme.

the random numbers. In this step, they hash the random number along with the permutation identified before and compare the resulting hash with the hash on the ballot. Now, the authorities count the ballots in the same way of Farnel, that is, all votes are counted and the votes cast during the initialization phase are subtracted from this sum. This last step is unnecessary if all initializing votes are void votes.

4.2.1 Human Readable Paper Audit Trail

In the manner of Ryan [12], the scheme could be adapted to provide a HRPAT by employing a conventional ballot box as alternative to the paper shredder. This way, instead of destroying the top page in a paper shredder, this page may be cast into the conventional ballot box. The box would store the top pages as an audit trail so that the votes can be counted without depending on the votes from the Farnel box.

5 Improving Threeballot Voting System

The ThreeBallot voting system was proposed by Rivest [9, 10]. This system attempts to satisfy end-to-end voter-verifiability without relying on cryptography. Several drawbacks, though, have been reported for Threeballot and improvements were incorporated in its newer versions. In this Section we introduce a variant of Threeballot that aims mainly at solving the information leakage problem pointed out by Araújo et al. [2] and independently by Clark et al. [6].

5.1 An overview of ThreeBallot voting system

The scheme employs a ballot design that consists of three single ballots. The ballots are identical except for the random IDs printed on the bottom of them. That is, they have the same list of options and bubbles to select them, but each ballot has a unique random ID. The IDs are encoded in a way that the voter cannot remember them. Figure 5 shows an example of the three single ballots.

In Threeballot, the voter should follow some rules to mark her ballot. That is, in order to vote for an option, the voter should mark the same option in two of the three ballots. The other options, though, should receive one mark each one in one of the three ballots. The choice of which of the three ballots she places these marks should otherwise be random.

After marking her three ballots, the voter inserts them into a machine that verifies the voter has marked the three ballots according to Threeballot rules. If the ballots were marked correctly, the voter chooses one of the three ballots as her receipt and the machine copies the ballot selected. Ideally, this should be done in a way that prevents the system learning which of the three ballots the voter chose to retain as her receipt. Now, the machine casts the three original ballots into a conventional ballot box to finish the process.

At end of the voting, the ballots cast are published on a bulletin board and all votes are counted. As each option received one extra vote, these votes are subtracted from the count to obtain the final results.

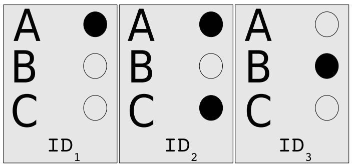

Figure 5: An example of a vote for option A in Threeballot voting system.

5.1.1 Drawbacks

The version of Threeballot presented in [9] has several drawbacks as also discussed by Rivest. However, most of them were mitigated or even solved in the last version presented in [10]. We briefly describe here two known problems of Threeballot: the reconstruction attack and the information leakage problem. Although Rivest proposed mitigations to the former problem, the latter is still unsolved so far.

The reconstruction attack As described above, all three single ballots are published on the bulletin board after the voting. Strauss [14] showed through simulations that from the ballots published, it might be possible to reconstruct the triples and so violate the voters privacy. In order to accomplish the reconstruction, the adversary chooses a targeted single ballot (possibly the receipt of a voter) and matches it to every possible pair of ballots on the board. As result, the attacker might find either the two other ballots that compose uniquely the valid triple or a set of possible pairs that form valid triples with the targeted ballot.

In order to mitigate this attack, Rivest proposed to replace receipts by means of the (original) Farnel idea in [9], that is, instead of keeping her own receipt, the voter casts it into a Farnel like box and receives another one. In [10], Rivest et al. consider a short ballot form (i.e. a ballot with few races and few candidates per race) to increase the possibility ballots being cast with the same pattern.

The leakage of information problem The voter receipt in Threeballot is a copy of a single ballot and so it contains part of the marks of the three ballots. The receipt, as observed by Araújo et al. [2], exposes some statistical information about the vote, but it cannot be used to violate the voter privacy. However, early statistical information about the voting results can be obtained from a large set of receipts. By means of simulations, Araújo et al. pointed out that by observing a large number of receipts and by counting the marks per option on these receipts, an adversary can acquire an indication of who is winning the voting before voting closes.

In the next Section, we present a version of Threeballot that mitigates the reconstruction attack and the leakage of information problem.

5.2 Our Proposal - Combining Farnel, Threeballot, and Prêt-à-Voter schemes

As described before, the votes in the Threeballot system can be reconstructed if no countermeasure is used and its receipt leaks statistical information about the voting results. We now introduce a version of Threeballot that overcomes these problems. Our solution can be seen as a combination of Threeballot scheme with Prêt-à-Voter [11, 3] based ballot forms and with Farnel. It employs cryptography and does not relies on a machine to check the validity of the voter selections.

The main component of Threeballot is its special ballot form. Thus, we firstly introduce a new ballot form for Threeballot scheme. Our ballot is based on the ballot form presented in Section 3.2 and is described as follows.

5.2.1 An initial ballot design

As the proposal of Rivest, the ballot form is composed of three single ballots. Here, though, the options on the three ballots are permuted and every ballot has a top and a bottom pages that are initially overlaid. More specifically, the top page of each single ballot is composed of the *same* permuted list of options, bubbles to select the options, and a distinct commitment and its decommitment to the options; the corresponding bottom page has a copy of the bubbles and of the commitment of the top page. Formally, let C be a list of options, π_C the permutation of C, H a secure hash function used as commitment, and a random number r_j $(1 \leq j \leq 3)$. Each single top page is composed of π_C, r_j, $H(\pi_C, r_j)$, and the bubbles. The bottom page contains only $H(\pi_C, r_j)$ and the bubbles. As before, the commitments and decommitments are covered by scratch surfaces and a carbon mechanism copies the mark from the top page to the bottom page. Figure 6 shows the new ballot design.

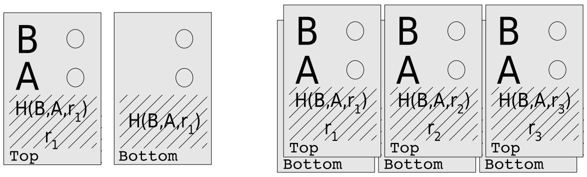

Figure 6: On left the two pages of the first single ballot. On right the three single ballots that compose the new ballot design for Threeballot scheme.

5.2.2 The ballot boxes

In addition to the special ballot design, the scheme employs two ballot boxes, that is, a Farnel box, such as described in Section 2, and a conventional box.

5.2.3 The Scheme

Following the Farnel idea, our proposal has a setup phase where the ballot boxes are initialized. Considering the ballot design described above and the initialization process presented in Section 2, the authorities initialize the conventional box with marked top pages and the Farnel box with the related marked bottom pages. Note that, with ThreeBallot, it is easy to encode void votes without needing to introduce an explicit void option: the ballot form is simply marked with exactly one bubble against each candidate. We still need to keep a protected record of how many initial, void votes are cast as this information is needed for the final check-sums.

As described in Section 2, the initialization process includes an audit to detect malformation of ballots (e.g. ballots with invalid commitments). In order to audit a ballot here, the scratch surfaces on both pages of the three single ballots are detached. Then, the commitment on both pages of each ballot are compared and opened using decommitment value from the top page. Specially, the options (π_C) printed on the form are hashed along with the random number (r_j) and the result is compared with the commitment ($H(\pi_C, r_j)$) on the top page. This auditing can also be performed in the voting phase. That is, before receiving her blank ballot, the voter assisted by help organizations can audit some blank ballots.

Voting Upon proving her eligibity to the voting authorities, the voter receives a sealed envelope with a blank ballot and performs the following steps to cast her vote:

1. (Marking the ballot) As the original Threeballot scheme, the voter marks her option in two of the three ballots and each other option once in one of the three ballots;

2. (Verifying the scratches on the top pages) The voter separates the top and the bottom pages of the three single ballots, and puts the top pages into an envelope. This envelope has windows to show just the scratch surfaces. She then hands the envelope to the authorities that verify the scratches are intact;

3. (Verifying the marks on the bottom pages) The voter now hands the three bottom pages to the authorities. They verify the pages were correctly marked and their scratches are intact.

4. (Casting the vote) The voter casts the three ballots from the top page of her ballot apart into the conventional ballot box;

5. (Obtaining a receipt) In order to obtain her receipt, the voter casts the three bottom pages of her vote into the Farnel box. The Farnel box removes the scratch surfaces on the pages, shuffles its set of bottom pages, and returns one or more copies of random bottom pages to the voter.

Tallying the votes When the voting has finished, the talliers open the conventional ballot box, detach the scratch strips on all top pages from the box, and publish the ballots on a bulletin board. Additionally, they also publish the bottom pages from the Farnel box. In order to compute the voting results, the authorities count all votes on the top pages, and subtract from them the votes cast before the voting and the extra votes cast à la Threeballot scheme.

Verifying the votes From the top pages posted on the board, anyone verifies if the decommitment values open their respective commitments. As before, this is performed by hashing the options along with the random number, and then comparing the result with the original hash. In addition, the voters can search top pages on the board that correspond to their receipts (i.e. copies of bottom pages).

Notice that the scheme could use the original Farnel ballot box (see Section 2.1) adapted to remove scratches surfaces in replacement of the enhanced Farnel box. In this case, the box would only exchange receipts as in [9]. That is, instead of casting the three bottom pages of her ballot (see step 5 above), the voter chooses one of the three pages, casts it into the adapted box, and destroys the other two pages; the box gives the voter a random selected receipt from its set. Alternatively, the voter could cast her three bottom pages and receives three random pages from the Farnel box. These alternatives have the advantage of employing a more simple Farnel box while not requiring a check machine.

5.3 Achieving ballot secrecy without the Farnel box

The version of Threeballot scheme proposed above does not leak statistical information. On the other hand, it actually makes the reconstruction attack more virulent, but this is countered by the Farnel mechanism. It, however, relies on Farnel ballot box that adds complexity to the scheme. We discuss now some ideas to simplify the scheme by making it independent of the Farnel box. Notice that the ideas are only of theoretical interest.

At first glance, one might think that the new ballot design could be enough to accomplish a scheme without the Farnel box. Indeed, the new design would overcome the leakage of information problem. However, it would make the reconstruction attack easier than in the original Theeballot scheme. Because the ballot forms have different permutations and each form employs the same permutation in the three single ballots, the bulletin board can be segmented in groups [2] according the permutations used. This way, an attacker only needs to compare ballots in the same group to reconstruct the votes.

In order to render the reconstruction attack harder to perform, we change our ballot design to allow distinct permutations in the same ballot. That is, instead of using the same permutation in the three ballots, a random permutation is selected for each of the three ballots. However, now some voters can have problems to select options in different positions in the three ballots. In addition, the authorities cannot verify the marks on the three bottom pages as the options

[2] For a number of voters N and a number of candidates C, the number of groups is $3n/C!$ or $3n/C$ for cyclic permutations.

can be in different positions. The first problem could be reduced by training the voters, so we concentrate in a solution for the second one.

A possible solution for verifying the marks on a ballot is to introduce a mapping between the two pages. That is, the options on the top page are mapped to elements on the bottom page and each option can be associated to different elements in other ballots (see Figure 7 for an example). Thus, the authorities can verify the bottom page without recognizing the options. This mapping could be, for example, a bijection between the set of options and a set of positive integers. Although the mapping hides the options from the authorities and still makes possible the check, it would reveal the relationship between the three permutations to the authorities. This way, a malicious authority could perform the reconstruction attack by grouping the permutations as before.

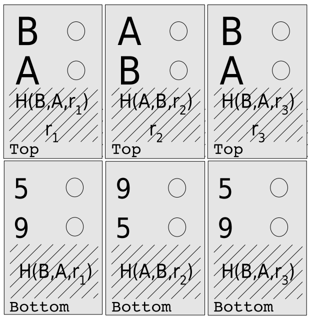

Figure 7: An example of ballot with mapping.

The mapping described makes the scheme simple as the authorities can verify the marks, but it is not secure. Thus, we rely on a check machine to solve the problem such as the original Threeballot scheme and avoid the mapping. The drawback is that here the machine needs to identify the options to verifies the marks. We assume, however, that the machine can perform comparisons without storing the options and describe the ballot form as follows: let C be a set of options, π_C^j $(1 \leq j \leq 3)$ be the permutation of C, H be a secure hash function, and a random number r_j $(1 \leq j \leq 3)$ from a large (key) space. Each single top page is composed of π_C^j, r_j, $H(\pi_C, r_j)$ along with the bubbles to mark the options. The corresponding bottom page contains the same $H(\pi_C^j, r_j)$ and bubbles of the top pages. Figure 8 illustrates this ballot design.

Based on this modified ballot form and considering the check machine to verify the votes, the voter now performs the following steps to vote:

1. (Selecting the option) She marks her vote on her ballot and verifies the marks by means of the check machine as the Threeballot scheme;

2. (Verifying the ballot) The voter inserts her ballot into an envelope that hides the options and hands it to the authorities. The authorities verify the ballot was not separated and check the scratch on the top page;

3. (Casting the vote) She separates her ballot and casts the three single top pages apart into the conventional ballot box;

4. (Obtaining the receipt) In the presence of the authorities, the voter keeps one of the three bottom pages as receipts and destroys the other two in a paper shredder.

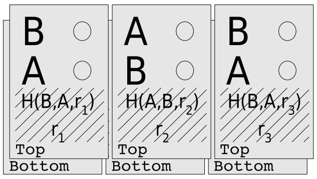

Figure 8: A ballot design for three ballots with different permutations.

As we changed just the structure of the ballot by employing different permutations, the tallying of votes remains the same described in the last Section.

6 Randell-Ryan Scheme with Farnel

Randell and Ryan [8] proposed a simple scheme that intends to improve the voter secrecy in the existing manual systems and to provide voter-verifiability. The scheme is based on Prêt-à-Voter [11, 3] ballot forms and does not rely on cryptography.

6.1 An overview of Randell and Ryan voting scheme

The scheme employs a ballot form that has a left (LC) and a right (RC) column. These columns are separated vertically by a perforation. The LC contains a randomized options list and the RC has the respective spaces to select the options. A serial number, CIN (code identification number), identifies the ballot uniquely and is printed at the foot of the LC. The RC has the same CIN at its foot, but it is covered by a scratch surface; this scratch has a receipt identification number (RIN) printed over it. The proposal also requires two conventional ballot boxes. One for receiving left columns and the other for receiving right columns.

During the voting, the voter receives a blank ballot and marks her vote on the right column. She then drops the LC into the ballot box for left columns. After that, in the presence of officials and observers, a photocopy of the right column, with scratch strip intact, is made as the receipt and the voter drops the original RC into the other ballot box.

In the tabulation process, the authorities open the ballot boxes and publish the RCs on the bulletin board so that the voters can verify their votes. The authorities then scratch off the strip on the RCs to reveal the CIN. After that, they match the CIN number on the LC and on the RC columns to identify the votes.

6.1.1 Drawback

The scheme has a drawback similar to that of the Farnel variant [2], that is, it requires trustworthy talliers. Votes could be altered before being counted and the voters could not detect this. According to the scheme, the voter receipt is a copy of the right column with its RIN. The voter compares the RIN and the mark on her receipt with the right column published on the bulletin board. However, after the authorities scratch off the RIN strips to show the CIN, she cannot verify her vote. This way, after exposing the CIN to count the votes, a malicious authority could potentially alter votes without being detected. Of course, various mechanisms can be proposed to counter this threat, but the fact remains that the resulting levels of assurance will inevitably be lower that those achievable using cryptographic mixing/tabulating mechanisms. Put differently: we need to make stronger trust assumptions with Randell/Ryan than are need in Prêt-à-Voter.

6.2 An improved scheme

Now we present an improved scheme that intends to overcome the drawback of Randell-Ryan proposal.

6.2.1 The ballot design

We use almost the same design of Randell-Ryan scheme, that is, the ballot form has a left and a right column that are separated vertically by perforations. However, the left column has a commitment to the options and its decommitment. In addition, it includes an index number to help the talliers aligning the columns in the tallying phase. The right column has spaces to select the options as well as a copy of the commitment and of the index. Scratch surfaces cover the commitments, the decommitment, and the index on both sides. Formally, let C be a set of options, I be a set of positive integers, π_C the permutation of C, H a secure hash function, i a unique number in I, and r a random number from a large (key) space. π_C, $H(\pi_C, r)$, r, i, compose the left column. $H(\pi_C, r_i)$, i, and spaces to select the options compose the right column. Figure 9 illustrates this ballot form.

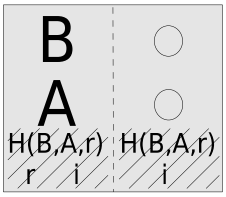

Figure 9: The new ballot form based on Randell-Ryan proposal.

6.2.2 Ballot boxes

As the Randell-Ryan scheme, we employ two ballot boxes. However, one of them is a Farnel ballot box (see Section 2 for details). The other is a conventional box.

6.2.3 The scheme

Due to the Farnel ballot box, the scheme has a pre-voting phase where the Farnel box as well as the conventional box are initialized. This process is the same described in Section 2. Here, though, the conventional ballot box is initialized with left columns and the Farnel box receives the corresponding marked right columns.

In order to audit blank votes, the scratch surfaces on the both sides of the ballot are detached. By doing this, the commitments ($H(\pi_C, r_i)$) as well as the index on the sides can be compared. After this comparison, the commitment on the left column is decommited using the random number on the bottom of this page. This is performed by hashing the options on the left column along with the random number and then comparing the result with the commitment on the page.

Voting The voter votes as in Randell-Ryan scheme. She selects her option on the left column and marks the corresponding space on the right column. Then, she performs the following steps to vote and to receive her receipt:

1. (Casting the ballot) In the booth, she separates the left and right columns of her ballot form and casts the left column into the conventional box;
2. (Obtaining the receipt) She drops the right column into the Farnel box. The box removes the scratch surface on the ballot and returns her a copy of a random column in already its set as receipt.

Tallying After the voting, the authorities open the two ballot boxes, and detach the scratch surfaces on all left and right columns. Then, they open all commitments on the left columns to verify the votes. As before, this is performed by hashing the options along with the random number and then comparing the commitments. After that, the authorities use the hash values to match the right and left sides. This allows the marks on the RH columns to be interpreted and counted. Then, the results are computed as the Farnel scheme. That is, all votes are counted and initial votes are subtracted from the count.

Notice that the scheme can be implemented using a ballot style without the hashes. However, in this case it may be possible for the tabulating authorities to manipulate the candidate lists unless appropriate anti-counterfeiting measures are in place. A nice feature of dropping the hashes is that we do not have to worry about checking the correctness of the hashes: a ballot are well formed as long as the CIN #s match. Thus, if we have on-demand printing of the candidates lists we simply need to ensure that the ordering are randomized.

If the threat of manipulation of the candidate orders during tabulation is regarded as sufficiently serious, then the hash function would be used. Here, though, we do need to check that the ballots are well-formed w.r.t. the hashes.

7 Conclusion

We have described a number of ways to enhance the Farnel box mechanism and combine it with existing schemes. We have shown that the Farnel concept is a fruitful one and worthy of further exploration.

We have also introduced a novel way to initialize the Farnel box that employs purely void ballots. This initialization, however, only works with the ballot forms that give rise to protected receipts with a void option, e.g., Prêt-à-Voter or ThreeBallot style ballots. The new process would be easier to monitor and verify than having to maintain a record the totals of the various votes cast in the initialization phase. Even so, ensuring only void votes are cast during the initialization phase is still challenging and will require carefully designed monitoring procedures.

Implementing the concept of the Farnel box in a way that requires minimal trust in the mechanism or procedures remains challenging. We have suggested a number of possible implementations, but none seem entirely satisfactory. In order to overcome the reconstruction attack in Threeballot by means of Farnel, Rivest [9] suggests separating the two functions of the Farnel box into two stages: one to make the receipt and one to exchange the receipt held by the voter. This may prove easier to implement with less trust assumptions, but will still need to ensure that the two steps performed in close proximity to prevent any possibility of the voter wandering off with her original receipt.

An interesting feature of the Farnel mechanism is that it may help counter certain psychological style attacks on voter-verifiable schemes in which voters are persuaded that the secrecy of their vote is not guaranteed. Using Farnel, the voters do not retain their own receipts, so any fear that the vote can be extracted should be mitigated. The down-side is that voters may be less motivated to check receipts if the receipt they hold is not their own. This may be offset by ensuring that voter helper organizations are on hand to perform the checks on behalf of the voters. If voters are given more than one receipt each this should also help as long as a reasonable proportion of voters are diligent enough to check all or many of their receipts.

Besides helping counter psychological attacks, the Farnel idea can also mitigate the randomization attack. This attack was introduced by Schoenmakers [13] and intends to nullify votes. To perform the attack, an adversary instructs the voter to generate a receipt that has a certain property. The attacker will not know what vote will be encoded, this is effectively random. The effect then is to force voters to vote for a random candidate, so nullifying their right to vote freely. The attack can be applied to Prêt-à-Voter and to Punch Scan schemes as the voter receipt in these schemes contains the position chosen by the voter. This way, an adversary may ask the voter to place her X in a specific position and to show him afterwards the receipt marked in this position. By means of Farnel, however, the voter exchanges her receipt before leaving the voting place. Thus, the adversary cannot verify that the voter followed his instructions.

References

[1] Ben Adida and Ronald L. Rivest. Scratch & vote: self-contained paper-based cryptographic voting. In *WPES '06: Proceedings of the 5th ACM workshop on Privacy in electronic society*, pages 29–40, New York, NY, USA, 2006. ACM.

[2] Roberto Araújo, Ricardo Felipe Custódio, and Jeroen van de Graaf. A verifiable voting protocol based on farnel. IAVoSS Workshop On Trustworthy Elections (WOTE'07), June 2007.

[3] David Chaum, Peter Y. A. Ryan, and Steve A. Schneider. A practical voter-verifiable election scheme. In Sabrina De Capitani di Vimercati, Paul F. Syverson, and Dieter Gollmann, editors, *ESORICS*, volume 3679 of *Lecture Notes in Computer Science*, pages 118–139. Springer, 2005.

[4] Ricardo Custódio. Farnel: um protocolo de votação papel com verificabilidade parcial. Invited Talk to Simpósio Segurança em Informática (SSI), November 2001.

[5] Ricardo Custódio, Augusto Devegili, and Roberto Araújo. Farnel: um protocolo de votação papel com verificabilidade parcial. Unpublished notes, 2001.

[6] Aleks Essex Jeremy Clark and Carlisle Adams. On the security of ballot receipts in e2e voting systems. IAVoSS Workshop On Trustworthy Elections (WOTE2007), Ottawa, Canada, June 2007.

[7] Stefan Popoveniuc and Ben Hosp. An introduction to punchscan. IAVoSS Workshop On Trustworthy Elections (WOTE'06), June 2006.

[8] Brian Randell and Peter Y.A. Ryan. Voting technologies and trust. *IEEE Security and Privacy*, 04(5):50–56, 2006.

[9] Ronald Rivest. The threeballot voting system. October 2006.

[10] Ronald Rivest and Warren Smith. Three voting protocols: Threeballot, vav, and twin. Electronic Voting Technology Workshop (EVT'07), August 2007.

[11] P.Y.A. Ryan. A variant of the chaum voting scheme. Technical Report CS-TR-864, University of Newcastle upon Tyne, 2004.

[12] P.Y.A. Ryan. Pret a voter with a human-readable, paper audit trail. Technical Report CS-TR-1038, University of Newcastle upon Tyne, 2007.

[13] Berry Schoenmakers. Personal communication, 2000.

[14] Charlie E. M. Strauss. A critical review of the triple ballot voting system. part 2: Cracking the triple ballot encryption, October 2006.