A Verifiable Voting Protocol based on Farnel

Roberto Araújo[1], Ricardo Felipe Custódio[2], and Jeroen van de Graaf[3]

[1] TU-Darmstadt, Hochschulstrasse 10, 64289 Darmstadt - Germany
rsa@cdc.informatik.tu-darmstadt.de
[2] UFSC, Campus Universitário, Trindade, 88040-900 Florianópolis (SC) - Brazil
custodio@inf.ufsc.br
[3] UFMG, Av. Antônio Carlos 6627, 31270-901 Belo Horizonte (MG) - Brazil
jvdg@ufmg.br

Abstract. *Farnel is a voting system proposed in 2001 in which each voter signs a ballot. It uses two ballot boxes to avoid the association between a voter and a vote. In this paper we first point out a flaw in the ThreeBallot system proposed by Rivest that seems to have gone unnoticed so far: it reveals statistical information about who is winning the election. Then, trying to resolve this and other flaws, we present a new, voter-verifiable version of the Farnel voting system in which voters retain copies of ballot IDs as receipts.*

1 Introduction

Secure voting systems are a cornerstone of modern democratic societies. They can prevent or detect frauds or faults, and so provide accurate results. To increase transparency in such systems, researchers have been designing voter-verifiable schemes. These schemes allow the voter to verify whether her vote was taken into account in the result, but without violating the vote privacy.

Different strategies have been used to design voter-verifiable schemes. Almost all solutions described in the literature uses cryptography as basis, but the resulting protocols are often hard to grasp by a common person. Recently, a new kind of scheme with verification property was proposed by Rivest [9] - the ThreeBallot voting system. His proposal attempts to satisfy the voter verifiability *without* employing cryptography. Many drawbacks, though, have been reported for this scheme and improvements were incorporated in its newer versions.

In 2001, Custódio et al. [3][4] proposed a protocol, called Farnel [4], in which uses two ballot boxes and the voters sign ballots. In fact, Rivest uses the concept of the Farnel to sidestep a flaw in his scheme.

This paper presents a new version of Farnel, which is voter-verifiable. Also, it points out another flaw in the ThreeBallot scheme which seems to have gone unnoticed so far; it leaks information. We do this as follows: Section 2 describes the original Farnel protocol. Section 3 shows how the ThreeBallot protocol leaks information. Section 4 describes the new Farnel protocol; it inherits some interesting characteristics that can be incorporated to obtain a verifiable voting

[4] Farnel means basket in Portuguese.

system. Section 5 presents an electronic version of our protocol. Finally, Section 6 presents our conclusions.

2 The original Farnel scheme

Farnel [3][4] was conceived to address the problems of a conventional ballot box. This paper-based scheme requires each voter to sign one ballot. However, in order to avoid an association between the voter and her ballot, the voter does not sign her *own* ballot, but another one chosen at random, as explained below. This way it is possible to know who the voters were, and any attempt to add, modify, or delete votes, after the voting period, can be detected.

Initialization phase Farnel uses *two* ballot boxes. Before voting starts, the first ballot box is publicly initialized with ballots filled out and signed by a ballot authority. This set of ballots must represent, with an equal probability, all possible votes. The second ballot box starts empty.

Voting phase In order to vote, the voter receives a blank valid ballot (signed by the ballot authority), makes her choice, and casts the ballot into the first (pre-initialized) ballot box. Then, through manual or mechanical shuffling, the first ballot box presents a ballot chosen randomly from its current set of votes to the voter. After receiving the ballot, the voter signs and drops it into the second box. This ends the voting process for the voter.

Tallying phase After the voting period has finished, the ballot authority opens and signs a second time all the votes of the first ballot box and adds them into the second box. Then the second box is opened and all ballots are counted. From this result the ballots from the initialization step are discounted.

Properties of Farnel Farnel gives warranties to the voter that her ballot will be counted, and that the exclusion or the addition of new votes is not possible after the voting phase. Anyone can, for example, verify that all ballots are signed, either by the voters or by the precinct. Moreover, everybody can check who voted without needing the list of voters. The scheme, however, is not voter verifiable.

3 Information leakage in the ThreeBallot voting system

We give a brief description of Rivest's ThreeBallot voting scheme [9]. It gets its name from the fact that each ballot consists of three columns, each representing a full ballot. Each row of the ballot has a candidate name, and a ballot must have exactly one of the three cells following the candidate name marked. However, the candidate that gets chosen will have two cells marked. For instance, in the example ballot of Figure 1 the voter chose candidate 1.

candidate 1	X		X
candidate 2		X	
candidate 3			X

Fig. 1. A ballot for candidate 1.

Then the ballot is cut vertically, separating the three columns. One of the columns is copied; this is the voter's receipt. All three columns end up in the ballot box.

When the voting phase has completed all votes are tallied. Obviously each candidate gets one free vote per ballot, so these votes must be subtracted to obtain the final tally.

There is a flaw with this scheme which is not mentioned in the latest version dated October 1, 2006; information about the contents of the ballot box is leaking before the election has finished.

When reading the ThreeBallot paper superficially it may appear that the secrecy of the ballot is perfect, i.e., that no information leaks. However, each receipt in fact does reveal a tiny bit of information, so little that it cannot be used against the voter. But in a large set of receipts statistical pattern do emerge.

The issue is best explained using an extreme example: suppose that candidate 1 gets all the votes and the other two none (we are assuming 3 candidates). Furthermore, suppose that all voters behave uniformly random with regard to where they put the marks and which column they choose as a receipt. Finally, suppose that all voters are willing to show their receipt to some organization who are at the polling station awaiting people who have just voted.

Counting the number of marks for each candidate (row) on the receipts reveals information on who is winning the election at that particular polling place. In this example, the winning candidate can expect 2/3 mark per receipt, whereas all the others can expect only 1/3 mark per receipt. The information is of a statistical nature.

To show the effect we wrote a small simulation program. Table 1 shows ten simulations for an election with three candidates, where 100 receipts have been collected and candidate 1 gets all the votes. The lines show the number of marks for each candidate, leaving no doubt at all about who is winning already while voting is still going on.

Table 1. A simulation of ten elections where every voter votes for candidate 1 and 100 receipts are collected.

1	2	3	4	5	6	7	8	9	10
69	73	61	65	65	64	65	65	68	61
34	39	32	37	29	32	30	31	29	34
43	34	31	37	30	37	37	28	26	27

In fact we are dealing with two (p, n)-Bernoulli distributions: one with $p = 2/3$, and the others with $p = 1/3$. In both cases n=#receipts.

Observe that adding candidates (rows) to the ballot does not help. Adding columns does, because it flattens the distributions ($p = 1/4$ vs. $p = 2/4$; $p = 1/5$ vs. $p = 2/5$ etc.), but this is undesirable for practical reasons.

Observe also that a statistical analysis is more difficult if the voters do not behave randomly and the original scheme is used: the voter chooses which column to copy.

The flaw in the ThreeBallot system is debatable. It is true that the information obtained from the receipts has the same effect as exit polls. But there is a difference: not every country has or allows exit polls, and in addition voters can lie about how they voted, whereas in the threeballot system the receipts reveal actual information. In an election where the difference of votes among two candidates is small, for example, the information obtained from the receipts can certainly influence voters while the election is going on.

4 A variant of the Farnel scheme

As presented in Section 2, the Farnel scheme does not provide individual voter verification; it just ensures, through signatures, that after the voting phase votes cannot be excluded and that new votes cannot be added. In this Section we present a new paper-based scheme inspired on Farnel. It also uses two ballot boxes, but does not depend on signatures. It provides a receipt to the voter, but without leaking information during the voting phase.

4.1 Prerequisites

The ballot form The ballot form used is composed of two halves. The first half is not much different from the layout traditionally used in elections. It is composed of a list of voting options (including a blank vote option) where next to each option there is a space to select it. It also contains a unique identification number (ID) which identifies the ballot uniquely and associates it to the election. The second half contains only the same ID (see Section 4.4 for a discussion about the IDs). The halves are separated by a perforation to allow separation by the voter and the IDs are covered by scratch surfaces (see also Figure 2).

The ballot boxes Two ballot boxes are used. One of them is a conventional box; however, it must be initialized with filled out, fake ballots (i.e. just the part that contains the options) before the voting starts (otherwise the first voters would not have a set of random IDs to choose from).

The other ballot box is able to receive a slip containing an ID, to add it to a set of already received slips, and to copy l randomly chosen IDs from this set. To this end we assume that the box has some mechanical device, and that copies are made in a memoryless way. The shuffling mechanism, for example, could be based on a bingo cage. We call this ballot box *Farnel*.

4.2 The protocol

Initialization phase In this phase the ballot authority establishes the following *voting parameters*: a number x of initial votes and a number l of IDs that should be printed on the receipts (see Section 4.4 for a discussion about these parameters). Moreover, he initializes publicly the ballot boxes. Let's say that there are v eligible voters in the election.

Before initializing the ballot boxes, the ballot authority performs a cut-and-choose process to prove the correct formation of the ballots. For a number $2x$ of blank ballots, he takes x ballots at random, detaches their protective layers, and publishes them; these ballots are no longer used. After that, the authority holds the other (entire) x blank ballots and tears each of them in two along the perforation. Next, he marks an option on each of the parts containing the options, detaches their layers, and casts them into the conventional ballot box. The options can be selected at random, but each of them should have at least one vote. The authority then scratches away the layers of the other parts (the slips that contain copies of the IDs) and casts them into the Farnel ballot box. The total number of fake votes for each option is published. Finally, the authority seals both boxes until the voting begins. **Note that neither the authority nor third parties should be able to record or remember the IDs of the initial votes**.

Voting phase After proving her eligibility to the voting authorities, the voter receives a blank ballot form. The following steps are performed to cast a vote and to obtain the receipt (see also Figure 2).

1. (Verifying and filling out the ballot form) The voter scratches away the layer covering the IDs and matches them (a). If they are equal, she makes her vote by marking one of the options available (b). We assume that the voter cannot record or remember the ballot's ID.
2. (Casting the vote) The voter separates the two parts (c) of the ballot form, casting the part containing the ballot ID and the options into the conventional ballot box (d). The other part, showing only the ID, is cast into the Farnel ballot box (e).
3. (Obtaining the receipt) The Farnel ballot box is shuffled (f) and l copies are produced of IDs which are printed as a receipt to the voter (g).

Tallying phase In a public session the talliers open the two ballot boxes and publish their contents on the bulletin board. To compute the results of the election, all votes are tallied. The fake ballots cast in the initialization phase are subtracted from the sums yielding the final result.

Ballot verification Anyone can check on the bulletin board whether each ballot from the conventional ballot box has a corresponding ID in the Farnel ballot

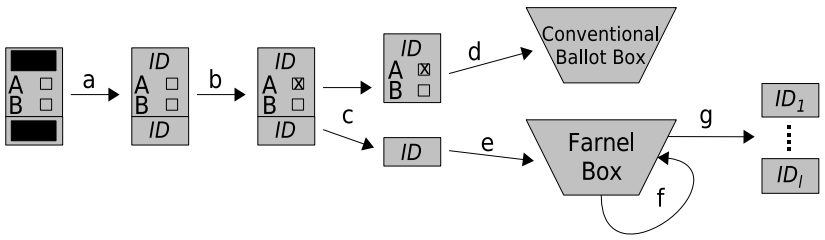

Fig. 2. Main voting steps of the new paper-based scheme.

box. In addition, the voters confirm whether their receipts (i.e. the IDs) match to ballots on the bulletin board. If one ballot and its ID were not published, the voter can complain by showing her receipt to a voting authority.

4.3 Security requirements and the new paper-based scheme

Here we sketch an analysis of our scheme based on the security requirements normally found in the literature. In this analysis, we supposed that the bulletin board cannot be compromised.

Accuracy In our scheme duplication, elimination, substitution, and addition of votes can be detected. The detection is accomplished by checking the information published on the bulletin board. Duplicates can be identified by checking if the IDs of the votes published are unique. Anyone can also detect elimination and substitution of votes. Every vote on the board should have a corresponding ID published. Moreover, the voters can independently match their receipts (i.e. the IDs) to the votes on the board. Note, though, that the detection is probabilistic since not all votes will have their IDs printed on the receipts. The addition of votes can be detected through the total number of votes published. The total should be the sum of the number of initial votes and of the number of voters that cast their votes.

Privacy The voter privacy in our scheme is assured even if she wants to violate it as follows. An adversary could try to violate the privacy by: obtaining an ID of a specific vote or extracting information about votes from the receipts.

In the first case, a voter or a voting authority could attempt to remember or record the ID of a ballot. In order to prevent this, the ballots in our scheme have their IDs covered by scratch surfaces. We suppose, though, that the voter cannot remember or record the ID of her own ballot (see Section 4.4 for a discussion about the IDs).

In the second case, an adversary could ask a voter to point the ID of her vote on her receipt. As the receipt is composed of a set of IDs chosen at random, the voter can only try to guess an ID related to her option. Again, we consider that the voter cannot remember the ID of her own ballot.

Alternatively, the adversary could collect the receipts of most of the voters and try to determine the votes of the first voters; we call this the *attack of collecting receipts*. He could explore the fact that the IDs of the first voters are more probable to appear on the receipts than the IDs of the last ones. To attempt to determine the votes, the adversary would check the most repeated IDs on the receipts and match them later to the votes on the board. Note, though, that the IDs cast before the election are as probable as the IDs of the first voters and that the adversary cannot distinguish among them at least as long as there is one initial vote for each option.

Verifiability Our scheme can be verified by the voters and by third parties. The IDs on Farnel ballot box aim at verifying the votes on the conventional ballot box. The publication of the IDs and of the votes allows anyone to verify the exactness of the voting results.

The voter verifiability in our proposal is different from the normally found in the literature. Instead of verifying if her own vote is in the final tally, the voter verifies a small subset of all votes. This is accomplished by matching the IDs on the receipt to the votes on the board. Note that here the verifiability depends on the *voting parameters*.

Voter-verifiable election schemes usually take into account that the voters will use their receipts to verify their ballots. Karlof et al. [7], though, pointed out that some voters can discard their receipts. Consequently, an adversary could take advantage of this and replace votes without being detected. Our scheme employs two approaches to mitigate this problem. First, each receipt is composed of a set of IDs and these IDs (or some of them) can be printed in others receipts. This way, even if a voter discards her IDs, others voters can possibly verify the votes related to these IDs. Second, the Farnel ballot box maintains the IDs of all votes and they are also published. This way, it adds redundancy to the verification process.

4.4 Discussion

We discuss now some aspects inherent to our scheme.

The IDs The IDs on the ballot form should be easy to compare and difficult to remember. The voter should compare the IDs to detect a possible malformation of her ballot (i.e. different IDs) and should not be able to remember it afterwards. Although these properties seem to be contradictory and difficult to implement, barcodes could be used to encode IDs and prevent the voters to recall them; in addition, the voter could compare barcodes easily as long as they are thick enough. However, some voters may not perform the comparison or ignore the malformation of their ballots.

A better solution is to avoid the voter comparing the IDs of her own ballot. This way, the ID could be just an alphanumeric string. The drawback is that

malformed ballots would not be detected. To mitigate this, we employ a cut-and-choose process for auditing the ballots: before receiving a blank ballot, the voter chooses some random ballots and detaches their scratches to verify their IDs; these ballots are discarded. As the voter does not remove the scratches of her own ballot, the Farnel ballot box now needs a special mechanism to remove the scratch of the slip; the other scratch can be removed in the tallying phase.

Chain voting This is a real threat to our scheme. An adversary can smuggle a valid blank ballot, mark an option on it, and corrupt a voter to use this ballot. After voting, the voter returns to the adversary the blank ballot received from the voting authorities. The adversary now can use the new blank ballot to corrupt another voter in the same way. Besides the usual chain voting attack, the adversary here can perform differently. He can obtain a slip containing an ID and corrupt a voter to use it; the voter gives back the adversary the ID of her own ballot and the adversary uses this ID to confirm the voter vote in the tallying phase. In both cases, therefore, the security of the scheme would be compromised.

In order to prevent these attacks, we modify our ballot form. We add a serial number to the ballot such as Jones [6]. The number is printed over the scratch surface that covers the ID on the slip. We also change the position of the other ID to allow the ballot to be folded showing just the scratch surfaces; now it is printed on the back of the ballot (see Figure 3 for an example of this ballot form). In addition to the ballot form, the following voting steps of the scheme must be modified. The authorities should record the serial number of the ballot before giving it to the voter and should confirm the number before the voter casts her vote. Note that the scratch containing the serial number should be removed after the authorities confirm it and that the voter does not compare the IDs on her ballot. As above, we assume that the scratch is removed by the Farnel ballot box.

Fig. 3. Ballot form to prevent chain voting.

The voting parameters As describe before, the voting parameters are composed of the number of initial votes x and the number of IDs l printed on the receipts. These parameters as well as the number of voters v affect the voter verifiability. Also, they are related to the privacy, that is, they can facilitate or not the attack of collecting receipts (see Section 4.3). From these remarks, we deduce the following:

Considering x much bigger than the number of voters v (e.g. 10 times bigger), the IDs cast by the voters will be almost statistically indistinguishable from the initial IDs. As result, if l is small (e.g. $l = 1$), an adversary cannot violate the voters privacy (particularly the privacy of the first voters) by distinguishing among the voters IDs and the initial IDs from the receipts. A small l, though, affects the voter verifiability as the chance of detecting problems in the tally decreases.

A v bigger than x, on the other hand, results in more IDs of voters on the receipts even if l is small. Consequently, the adversary has more chance to distinguish among the voters IDs and the initial IDs. For example, if $v = 500$, $x = 2$, and $l = 4$, the voters IDs will appear more on the receipts than the initial ones and this facilitates the distinction.

Certainly, the voter verifiability and the privacy in our scheme are related. To measure the relation of these properties, we performed experiments considering a fix number of voters $v = 500$ for different l and x. The experiments show the probability of the ID cast by the first voter to appear on at least one of the 500 receipts. Also, they show the chance of detecting the elimination of a vote in the tally through the receipts. Table 2 presents some results of these experiments.

Table 2. x - number of initial votes; l - number of IDs on each receipt; *1st ID* - prob. of the 1st ID to appear on at least one of the receipts; *Detection* - prob. of detecting the elimination of a vote in the tally by means of the receipts.

x	100		300		1000			1500		
l	1	2	1	2	3	4	5	4	5	6
1st ID %	83	97	62	86	70	80	87	68	76	82
Detection %	49	66	43	63	60	69	76	61	68	74

The last votes and the receipts As described in Section 4.3, the initial IDs as well as the IDs cast by the first voters have more chance to appear on the receipts. This way, the votes corresponding to these IDs are more probable to be verified by the voters. Conversely, the IDs cast by the last voters have less chance to be printed on the receipts. That is, the probability of an ID_2 cast after ID_1 to appear on at least one receipt is less than the probability of ID_1 to appear on the receipts. Thus, considering a number of voters v, and the first and the last voter, we observe that the ID of the first voter can appear in any of the v receipts, while the ID of the last voter can appear only in the last receipt.

As the chance of the last IDs to appear on the receipts is less than the chance of first IDs, an adversary could substitute a vote of one of the last voters without being detected. Though, the adversary would need to identify these votes before substituting them. In principle, because the adversary cannot distinguish the first votes from the last ones, he cannot identify the last votes to replace them. Hence, the best that the adversary can do is to try to guess the last votes.

On the other hand, by observing most of the receipts, the adversary could identify the votes that would be possibly verified and substitute only the votes in which the IDs do not appear on the receipts. We call this the *attack of substitution of votes without receipts*. As the Farnel box keeps all the IDs that compose the receipts and these IDs are published on the bulletin board immediately after the voting, the attack would be detected. However, the adversary can still succeed if he is able to substitute votes (without receipts) and their corresponding IDs before they are published.

In order to mitigate the attack, all IDs in the Farnel box should appear on the receipts. Although this could be achieved by increasing the number of IDs on the receipts, the IDs of the last voters still would have less chance to appear on them. As solution, we propose to print a set of receipts at end of the voting. In other words, after the last voter cast her ID and obtains her receipt, the Farnel box shuffles its IDs (without receiving an input), prints a predefined number of receipts, and outputs them. These receipts would be handed to help organizations and could be also published on the bulletin board. Alternatively, the receipts could also be printed during the voting. However, the period defined for printing should consider the chance of the last votes to appear on the receipts.

Supervising the votes The ballot design is fundamental for the verification of votes in our scheme. As presented, it is composed of two halves that have the same ID. The halves are separated (i.e. by casting the two halves into different boxes) during the voting phase and in the tallying phase they are associated (i.e. by publishing the halves) to allow anyone to verify the vote. Still, the ID on the halves can appear on the receipts, so the voters can verify the corresponding vote independently.

However, due to the verification through the IDs, the talliers must be trustworthy. Particularly, from the opening of the conventional ballot box until the publication of all votes on the bulletin board, the talliers should supervise strictly the votes. Otherwise, an adversary (e.g. a malicious tallier) can compromise the security of the scheme.

Suppose this adversary has access to the set of votes of the conventional ballot box before the votes are published on the bulletin board. In order to compromise the scheme, the adversary could smuggle a vote from the set, makes a fake vote for a different option but using the same ID of the vote smuggled, and includes the fake vote in the set. This way, after publishing all ballots, the fake votes would appear on the board instead of the smuggled one.

This attack would undermine the security of the scheme. Because the votes are verified through the IDs, voters and third parties would not detect the substitution of the original vote by the fake one.

5 An electronic version of the paper-based scheme

We now introduce an electronic voting scheme of the scheme presented in the previous Section. It uses commitments as the IDs, which are constructed by a

voting machine (DRE). Also, it uses a special ballot box which accepts a ballot ID and hands out copies of other ballot IDs.

5.1 Building blocks

Threshold ElGamal Cryptosystem As a basis for the scheme we employ the ElGamal public key cryptosystem [5] under a subgroup of order q of Z_p^*, where p and q are large primes and $q|p-1$. More specifically, we utilize a threshold variant, as described by Cramer et al. [2].

We utilize the following notation: T is an ElGamal public key corresponding to a secret key $\widehat{T}$, while $E_T(i,s)$ is the ElGamal encryption of a message i constructed with T and a random number $s \in_R Z_q$, and $D_{\widehat{T}}(i)$ is the ElGamal decryption of i.

Mix Net In order to make encryptions anonymous during the tallying, we employ a mixnet. This primitive was introduced by Chaum [1] and further improved by many others authors. Specifically, we require a verifiable reencryption mixnet such as the proposal of Neff [8].

Commitment scheme Another cryptographic primitive is a commitment scheme , which must be homomorphic and will be used to commit to the voting options. We use the ElGamal cryptosystem for this purpose.

Cut-and-choose We employ a cut-and-choose process to prove the voter that her vote was correctly formed by the voting machine. This is accomplished in a similar way to Lee et al. [10]. Especially, the voting machine makes some encryptions with ElGamal and presents them to the voter; the voter selects some of them for verification and the machine opens them by revealing the random numbers employed.

5.2 Prerequisites

The ballot The ballot is constructed by the voting machine and is presented to the voter. It contains each possible option and some encrypted stuff next to it: each option i is associated to two commitments, as follows: $\langle i, commit(i, r_{i1}), commit(i, r_{i2})\rangle$ for $r_{i2}, r_{i1} \in_R Z_q$. In particular, each commitment is represented by: $commit(i,r) = \langle E_T(i,s_1), E_T(r,s_2),\ E_M(ir,s_3)\rangle$ for $r, s_1, s_2, s_3 \in_R Z_q$. Here r is chosen uniformly at random from Z_q, and ir is the product of i and r, while $(T,\widehat{T})$ stands for the ElGamal keys of the talliers, and $(M,\widehat{M})$ stands for the keys of the special ballot box which uses the *same* ElGamal modulus.

The special ballot box The paper-based scheme from the previous Section required a ballot box that received a ballot ID, shuffled its contents, and output copies of randomly chosen IDs. Here our special ballot box is initialized with a set of encrypted IDs which it keeps in a private list, L. It receives an enciphered ID, adds it to L, decrypts some random elements from L, and prints the result on the receipt. Elements selected are not deleted from the list, though.

There are four parties involved in our scheme:

Voters The voters cast votes and receive receipts for checking data later. Each receipt is composed of three parts: the ballot (with commitments and hidden commitments) and some decommitments, the commitment of the option chosen, and some plaintexts from the list L.

Voting machine The voting machine generates ballots, makes the first two parts of the receipts, and publishes commitments on the bulletin board.

Special ballot box It holds a private list of encryptions L and acts as described before. It has a barcodes reader and receives new encryptions through this reader. It also prints the last part of the receipt.

Tallying authorities These authorities are responsible for running a mixnet, and for decrypting and counting the votes. They also define the number of initial votes and generates them; in addition, they define the number of votes that each voter verifies. They hold the keys T,$\widehat{T}$. We suppose that a subset of the talliers are trustworthy.

5.3 The protocol

Initialization phase In this phase, the following parameters of the voting are established and published on the bulletin board: the voting options (or candidates), the number of initial votes, the number of IDs that should be printed on the receipt. Let's say that there are m options i ($i = 1, \cdots, m$), a number x of initial votes, and a number l of IDs printed on the receipt.

The talliers generate the initial votes according to x and to the commitment scheme explained before. Then the talliers publish commitments of the form: $commit(i, r) = \langle E_T(i, s_1), E_T(r, s_2), E_M(ir, s_3)\rangle$ for different $r, s_1, s_2, s_3 \in_R Z_q$ in each commitment. The values $E_M(ir, s_3)$ are handled by the special ballot box as its private list L of encryptions.

Voting phase After proving her eligibility to the voting authorities, the voter is allowed to interact with the voting machine in the voting booth. The following steps are executed to cast a vote and to obtain the receipt. Figures 4 and 5 exemplify the scheme for three voting options.

1. (Generating the ballot) For each option i ($i = 1, \cdots, n$), the machine generates the triple: $\langle i, commit(i, r_{i1}), commit(i, r_{i2})\rangle$ for $r_{i2}, r_{i1} \in_R Z_q$. As described before, each commitment is composed of: $\langle E_T(i, s_1), E_T(r, s_2), E_M(ir, s_3)\rangle$ for $r, s_1, s_2, s_3 \in_R Z_q$. After that, the machine prints the ballot on the receipt (a). Here as well as in the next two steps the ballot is not shown to the voter.

2. (Opening some commitments) The voter informs the machine to open the first or the second commitment for each option i. After this, the machine opens the corresponding commitments (b). In other words, for a commitment $\langle E_T(i, s_1), E_T(r, s_2), E_M(ir, s_3)\rangle$ already printed on the receipt, the machine prints r, s_1, s_2, s_3 on the receipt as decommitment.
3. (Voting) In order to vote, the voter informs her option i and the machine prints the corresponding, not opened, commitment on the receipt (c). In particular, the machine prints $\langle E_T(i, s_1), E_T(r, s_2), E_M(ir, s_3)\rangle$.
4. (Verifying the ballot) Now, the machine shows the receipt (A) to the voter. The voter should verify if the commitments selected were opened and if the commitment corresponding to her vote was printed. If the receipt is correct, the voter confirms her vote. The machine then prints a stripe on the not-open commitments of the ballot to erase them. She also prints the barcode of $E_M(ir, s_3)$ of the voter's vote and adds a digital signature to the receipt; the voter holds this receipt (B). The other elements of the vote, $\langle E_T(i, s_1), E_T(r, s_2)\rangle$, are sent to the bulletin board.

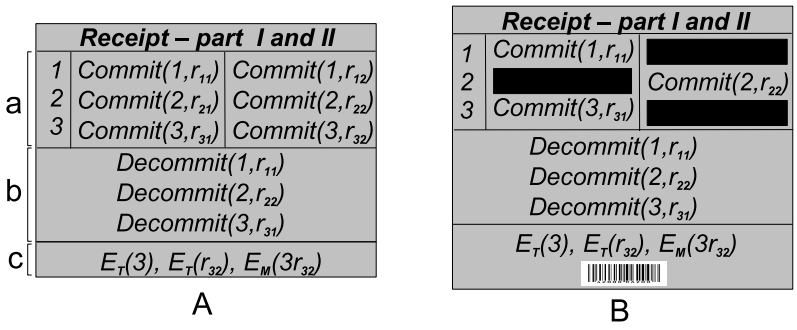

Fig. 4. Parts I and II of the receipt. A - the receipt that the voter verifies; B - the receipt that the voter holds.

5. (Obtaining the last part of the receipt) Using the barcodes reader, the voter adds $E_M(ir, s_3)$ to the special ballot box (d). The box writes $E_M(ir, s_3)$ to itsprivate list L. Then, it chooses y elements at random from L, decrypts them, and prints the results on the receipt (e). In doing so the elements are not deleted from L. Figure 5 illustrates this step.

Tallying phase After the election, the talliers send all pairs of encryptions published on the bulletin board, $\langle E_T(i, s_1), E_T(r, s_2)\rangle$, to a mixnet. The mixnet shuffles the pairs and publishes them on the bulletin board. After this, the talliers cooperate to decrypt the pairs to obtain the options i and the random numbers r. The talliers then multiply i and r, and publish the triples $\langle i, r, ir\rangle$ on the bulletin board. To compute the elections results, all votes from the voting phase (i.e. the

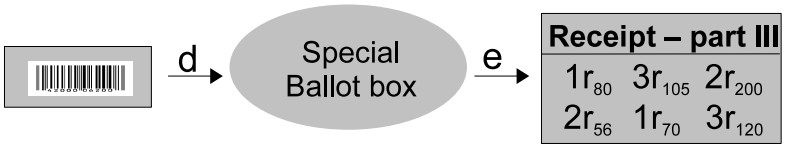

Fig. 5. The last part of the receipt.

i's) are counted and from this result are subtracted the fake votes generated in the setup phase.

Ballot verification The voter receives a receipt composed of three parts. The first part, which contains the ballot (with commitments and hidden commitments) and the decommitments, is used to verify the construction of the ballot. This may be accomplished by a helper organization through a computer. It constructs the commitments from the decommitment values of the receipt and then checks if the resulting commitments match the commitments printed on the receipt.

The second part of the receipt, which contains the commitment of the option chosen, is used to check if the commitment of the receipt appears on the bulletin board.

The third part of the receipt contains a list of ir to verify if the values ir match the values published by the talliers on the board.

6 Conclusion

We presented a new version of the Farnel voting system; a paper-based scheme and an electronic one. In addition, we showed a flaw in the ThreeBallot voting system.

The schemes introduce a new way to verify votes: the voter does not verify her own vote, but copies of a subset of votes cast so far. More precisely, the voter receives copies of some ballot IDs. These are used later to compare with the IDs of the ballots published on the bulletin board.

The paper based version uses a simple ballot form. It just requires the voter to compare IDs and to mark her option. However, the scheme relies on a ballot box that can shuffle and copy IDs. Also, the security of the scheme depends on the voting parameters.

We used the paper version to model the electronic version. The scheme works as expected, but is not very practical and has several drawbacks. It requires a verifiable mix net in the tally phase and the special ballot box must perform correctly. Moreover, the voter must compare a lot of information. We believe, though, that this scheme can be improved and are working in this direction.

References

1. David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, 1981.
2. Ronald Cramer, Rosario Gennaro, and Berry Schoenmakers. A secure and optimally efficient multi-authority election scheme. In *Advances in Cryptology - EUROCRYPT 97*, volume 1233 of *LNCS*, pages 103–118. Springer-Verlag, 1997.
3. Ricardo Custódio. Farnel: um protocolo de votação papel com verificabilidade parcial. Invited Talk to Simpósio Segurança em Informática (SSI), November 2001.
4. Ricardo Custódio, Augusto Devegili, and Roberto Araújo. Farnel: um protocolo de votação papel com verificabilidade parcial. Unpublished notes, 2001.
5. Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4):469–472, 1985.
6. Douglas W. Jones. Chain voting, August 2005. http://vote.nist.gov/threats.
7. Chris Karlof, Naveen Sastry, and David Wagner. Cryptographic voting protocols: A systems perspective. In *Proceedings of the Fourteenth USENIX Security Symposium (USENIX Security 2005)*, pages 33–50, August 2005.
8. C. Andrew Neff. A verifiable secret shuffle and its application to e-voting. In *SIGSAC: 8th ACM Conference on Computer and Communications Security*. ACM SIGSAC, 2001.
9. Ronald L. Rivest. The threeballot voting system. http://theory.lcs.mit.edu/∼rivest/Rivest-TheThreeBallotVotingSystem.pdf, October 2006.
10. Seungjoo Kim Yunho Lee, Kwangwoo Lee and Dongho Won. Efficient voter verifiable e-voting schemes with cryptographic receipts, June 2006. IAVoSS Workshop On Trustworthy Elections (WOTE2006), Cambridge.