

Preventing Board Flooding Attacks in Coercion-Resistant Electronic Voting Schemes

Reto Koenig^{1,2}, Rolf Haenni¹, and Stephan Fischli¹

¹ Bern University of Applied Sciences, CH-2501 Biel, Switzerland
{rolf.haenni,stephan.fischli}@bfh.ch

² University of Fribourg, CH-1700 Fribourg, Switzerland
reto.koenig@unifr.ch

Abstract. This paper addresses the board flooding problem of Juels et al.’s coercion-resistant electronic voting scheme. A key property of this scheme is the possibility of casting invalid votes to the public board, which are indistinguishable from proper votes. Exactly this possibility is crucial for making the scheme coercion-resistant, but it also opens doors for flooding the board with an enormous amount of invalid votes, therefore spoiling the efficiency of the tallying process. To prevent such attacks, we present an adaption of the scheme in which each voter receives—in addition to the proper credential—some dummy credentials from the election registrars. Dummy credentials may be used to deceive possible coercers. The list of all dummy credentials is published along with the electoral register. Based on the electoral register and the list of dummy credentials, the system is now capable of making a distinction between invalid votes generated from dummy credentials and invalid votes generated from fake credentials. While the former are kept until the tallying phase, the latter are immediately rejected by the public board. If the public board additionally rejects all incoming duplicate votes, then its maximum size is bounded by the total number of issued credentials. This guarantees an efficient linear-time tallying phase even in case of a massive board flooding attack with a very large number of invalid votes. Although the solution presented in this paper does not yet entirely rule out vote selling or coercion, it makes it at least unbearable for the vast majority of voters.

1 Introduction

One of the most challenging problems in remote electronic voting is the design of a system that prevents voters from selling their votes or from being coerced. The first scheme that is resistant against both the selling of votes and the coercion of voters has been proposed by Juels, Catalano, and Jakobsson in [7]. To achieve *coercion-resistance* (which implies mere *receipt-freeness*), the so-called “JCJ-scheme” uses an anonymous authentication mechanism to guarantee that the identities of the voters remain hidden during the whole voting and tallying process. The anonymous authentication mechanism requires that during the registration phase each voter receives a *secret credential* over an untappable

channel. The knowledge of the secret credential allows the voter then to post an encrypted vote anonymously to the public board, such that its inclusion in the final tally is guaranteed. It is also possible to post invalid votes based on *fake credentials*, but those will be filtered out later during the tallying phase. Since board entries created from proper credentials are indistinguishable from those created from fake credentials, it is always possible to lie about the secret credential or to supply a coercer with a fake credential. The vote buyer or coercer will then see the posted invalid vote on the public board, but at this early stage of the protocol, there is no way to tell whether a particular board entry will be included in the final tally or not. This is the principal mechanism that renders the JCJ-scheme coercion-resistant.

The JCJ-scheme is the point of departure of most advanced protocols for remote electronic voting today dealing with coercion-resistance, but the protocol as presented in [7] has at least two major open problems.³ The first problem is the quadratic running time of the tallying process, where duplicate and invalid votes need to be eliminated. Detecting duplicate votes requires so-called *plaintext equivalence tests* (PET) [6] for every pair of votes cast, and detecting invalid votes requires each vote cast to be checked against the public electoral register, thus making the scheme quite inefficient for large scale elections. The *Civitas system* [3], an implementation of the JCJ-scheme, weakens this problem by breaking up the electoral register into various independent blocks of a given fixed size. Several other improvements based on hash tables were proposed by Smith, Weber, and others [8,13,14,16,17], but they have been shown to be vulnerable to Pfitzmann's attack against anonymous channels [12]. More recent developments in this direction are based on group signatures [1,2] or fake votes generated by the talliers [15].

The second major problem of the JCJ-scheme results from the aforementioned possibility of posting invalid votes based on fake credentials to the public board. Exactly this possibility is crucial for making the scheme coercion-resistant, but it also opens doors for flooding the public board with an enormous amount of invalid votes. Because invalid votes are indistinguishable from proper votes from the perspective of the public board, there are no direct countermeasures against such types of attack, i.e., as long as the incoming votes cast are well-formed and comply with the protocol, the public board needs to treat them all in the exactly same manner. A massive application-level flooding attack of that kind may therefore both jeopardize the availability of the public board and spoil the efficiency of the tallying process. To our best knowledge, no practical solution to this problem has yet been proposed in the literature. The problem itself seems to be intrinsic to the chosen approach.

In this paper, we propose an extension of the JCJ-scheme that addresses both aforementioned problems. The key idea is to equip the public board with a stronger filter on what is an acceptable vote cast. For this, the voters receive during the registration phase some *dummy credentials* (in addition to the secret credential), which may then be used to mislead potential vote buyers or coercers.

³ Further major and minor problems of the JCJ-scheme are discussed in [9,14].

Invalid votes generated from these dummy credentials will be accepted by the public board (and filtered out later), but invalid votes from fake credentials will be rejected immediately. As we will see, enhancing the JCJ-scheme in such a way has a number of potential pitfalls. These pitfalls will be discussed and possible solutions will be presented.

The major benefit of our method results from the public board’s ability to separate invalid votes created by fake credentials from those created by dummy credentials. Let n denote the number of voters, m the number of issued dummy credentials, and s the number of votes cast using fake credentials. Note that n and m are fixed during the registration phase, whereas s is unbounded (and possibly orders of magnitude larger than $n + m$). If we further assume that the public board is also capable of eliminating duplicate votes, we can introduce an upper limit $n + m$ for the size of the public board.

Another important benefit of our approach is the fact that the known attacks against the linear-time improvements proposed by Smith [14] and Weber [16,17] are no longer possible. The reintroduction of these improvements allows the elimination of all types of invalid votes (duplicate, fake, and dummy) in linear time, which reduces the total running time of the original JCJ tallying phase from $O(n^2 + s^2)$ to $O(n + m)$. If furthermore $m = d \cdot n$ for some constant $d > 0$ (the average number of dummy credentials issued per voter), then the tallying phase even runs in $O(n)$ time.

Unfortunately, this unprecedented leap in performance and robustness has some negative effect with respect to perfect coercion-resistance. It is possible to minimize this effect to a small subset of unfortunate voters, which receive the minimum amount of dummy credentials, but some residual affliction will remain. The same holds true for vote buying. We will discuss this topic and see how to further minimize this problem.

The paper is organized in the following way. In Section 2, we give a short overview of the original JCJ-scheme and discuss its properties and problems. The proposed solution for the board flooding problem and the corresponding extension of the JCJ-scheme is discussed in Section 3. We first exhibit the general idea, then give a semi-formal sketch of the adapted protocol, and finally discuss some of the above-mentioned pitfalls. Section 4 concludes the paper.

2 Coercion-Resistant E-Voting

The goal of the scheme proposed by Jules, Catalano, and Jakobsson in [7] is to make remote electronic voting resistant against all sorts of coercion. Coercion-resistance is defined as a stronger form of privacy. While privacy is defined in terms of an adversary that cannot interact with voters during the election process, it is assumed that a coercive adversary may interact with voters at any time. Thus an election scheme is called *private*, if the adversary cannot guess the vote of any voter better than an adversarial algorithm whose only input is the final tally, and the scheme is called *coercion-resistant*, if the adversary can be deceived into thinking that a coerced voter has behaved as instructed. Such

a scheme thus prevents voters from selling their votes or from being coerced in various ways, e.g. to vote in a particular way, to vote at random, to abstain from voting, or even to divulge the private keying material [7].

The JCJ-scheme is the first electronic voting protocol that offers full coercion-resistance under minimal assumptions. While many other protocols assume the existence of an untappable channel during the voting phase to offer mere receipt-freeness, an untappable channel is only required during the registration phase of the JCJ-scheme. Note that this assumption is realistic, because the registration process often requires the voters to visit the registration office in person. We will now briefly describe the JCJ-scheme in a semi-formal way. The main entities in the protocol beside the voters are the following:

Registrars They issue the secret credential to voter V_i and pronounce corresponding encryptions publicly to the system. A threshold encryption system guarantees that the secret credential is only known to the voter and that the protocol is safe even if a minority of the registrars is corrupted or under attack.

Tallying Authorities They are responsible for processing the votes cast, jointly decrypting and counting the votes, and publishing the final tally. Again, a threshold encryption systems guarantees the safety of the protocol even if a minority of the tallying authorities is corrupted or under attack.

The votes cast are published on an append-only *public board*. Its task is to accept and publish every well-formed vote cast that complies with the protocol. To guarantee the integrity and availability of the board, it may be replicated in such a way that a minority of unavailable or corrupted board servers does not prevent its functioning properly as a whole [5,11].

The whole voting protocol is divided into three major phases, during which the voters are authenticated anonymously. The protocol uses numerous cryptographic primitives such as encryption, digital signatures, non-interactive zero-knowledge proofs of knowledge, plaintext equivalence tests, re-encryption mix-nets, anonymous channels, etc. A first overview of the protocol is given in Figure 1.

Registration Each voter V_i , $1 \leq i \leq n$, receives a secret credential σ_i jointly generated by the registrars. This credential constitutes a proof of eligibility, which is used in the voting phase to cast the vote. Additionally, an encryption $S_i = \text{Enc}_\varepsilon(\sigma_i, \gamma_i)$ of σ_i with randomness γ_i is appended to the (digitally signed) electoral register on the public board. ε denotes the tallying authorities' common public key. The protocol assumes the majority of the registrars to be trustworthy and the channel between the registrars and V_i to be untappable.

Voting The voters cast their candidate selection $c_i \in \mathcal{C}$ via an anonymous channel to the public board. The message posted to the board consists of $A_i = \text{Enc}_\varepsilon(\sigma_i, \alpha_i)$ and $B_i = \text{Enc}_\varepsilon(c_i, \beta_i)$ along with corresponding zero-knowledge proofs of knowledge of σ_i and c_i . It is important that the candidate set \mathcal{C} is finite and that an additional disjunctive proof that c_i represents




Fig. 1. Overview of the original JCY-scheme: the first filter eliminates votes with invalid proofs and duplicate votes from the public board, while the second filter checks the votes cast against the electoral register (and thus eliminates votes created from fake credentials).

a valid candidate choice is provided. This proof is needed to prevent the construction of receipts based on invalid candidate choices.

Tallying The tallying authorities check the proofs included in the votes cast and jointly perform pairwise PETs on the encrypted credentials to eliminate duplicates. The resulting adjusted list of votes cast is shuffled in a verifiable re-encryption mix-net to anonymize the votes and credentials included. Respective proofs of correct shuffling are published on the public board. Another verifiable re-encryption mix-net is applied to the electoral register, which finally allows the tallying authorities to jointly check the validity of the encrypted credentials involved in the votes cast (without decrypting them). Votes accompanied with fake credentials are discarded. The resulting adjusted list of proper votes is decrypted and tallied.

What makes this particular system coercion-resistant is the fact, that any posted entry to the public board is accepted if it is well-formed and complies with the protocol. It must thus consist of a valid candidate selection and some credential encrypted by the tallying authorities' common public key (together with corresponding proofs of knowledge). But the credential must not necessarily be a proper credential issued by the registrars and thus constituting a proof of eligibility, it simply needs to have the format of a proper secret credential. This enables the voter to deceive potential coercers with a *fake credential*, simply by choosing one at random. Votes accompanied with such fake credentials are discarded during the tallying phase. The two mix-nets involved in the tallying phase guarantee that no voter can prove to a third party whether a particular vote cast has been discarded before tallying or not. This feature makes the system resistant against selling votes or coercing voters.

Scheme by Smith and Weber [14,16,17] Instead of applying PETs on all pairs of distinct votes for removing duplicates, both Smith and Weber in essence suggest computing and decrypting $A_i^z = \text{Enc}_\varepsilon(\sigma_i^z, \alpha_i^z)$, where $z \in \mathbb{Z}_q$ is a random value shared among the talliers. The resulting *blinded credentials* σ_i^z are stored in a hash table for collision detection in linear time (clearly, $\sigma_i = \sigma_j$, iff $\sigma_i^z = \sigma_j^z$). Both authors propose using the same procedure for eliminating votes created from fake credentials, but since the same exponent z is used across all ciphertexts A_i , the coercer gets an attack strategy to identify whether a vote with known σ_i is counted, namely by posting two votes, one that includes an encryption of σ_i and one an encryption of σ_i^2 [1,3,12]. Note that this attack is not applicable to the mere removal of duplicates.

3 Preventing Board Flooding Attacks

In this section, we describe a way of modifying the JCJ-scheme to become resistant against board flooding attacks and to allow a linear-time tallying phase. As this implies several major and minor changes to the JCJ-scheme throughout various parts of the scheme due to different reasons, we uncover them step by step. A discussion of some important related questions follows in the second part of this section.

3.1 The Modified JCJ-Scheme

To protect the public board against application-level flooding attacks, it needs to be equipped with a stronger filter on what to accept. The main idea of our approach is to accept only votes cast from legitimate voters. Since it is crucial for the original JCJ-scheme to accept any vote cast, even those accompanied with a fake credential, it seems to be impossible in the first place to make such a distinction between legitimate and non-legitimate voters. But by introducing a third category of credentials, so-called *dummy credentials*, which are distributed to the voters during the registration phase (together with the proper secret credential), it is possible to reject all votes accompanied by fake credentials right from the beginning. Thus the idea is that the dummy credentials take over the role of deceiving potential vote buyers or coercers. This means that during the vote casting phase, they need to be treated in the same way as the secret credentials, whereas fake credentials are immediately rejected. In other words, voters are equipped with several access keys for posting votes to the public board, but only one of them is a key to include votes in the final tally. A first overview of the extended protocol is given in Figure 2.

More formally, let $\{\tau_{ij} : 1 \leq j \leq d_i\}$ be the set of dummy credentials for voter V_i (note that that d_i might be different for every voter, see Subsection 3.2). They are generated jointly by the registrars during the registration phase, together with V_i 's secret credential σ_i . Corresponding encryptions $T_{ij} = \text{Enc}_\varepsilon(\tau_{ij}, \gamma_{ij})$ with randomness γ_{ij} are published on the public board together with $S_i = \text{Enc}_\varepsilon(\sigma_i, \gamma_i)$. The public board thus contains two separate lists of encrypted

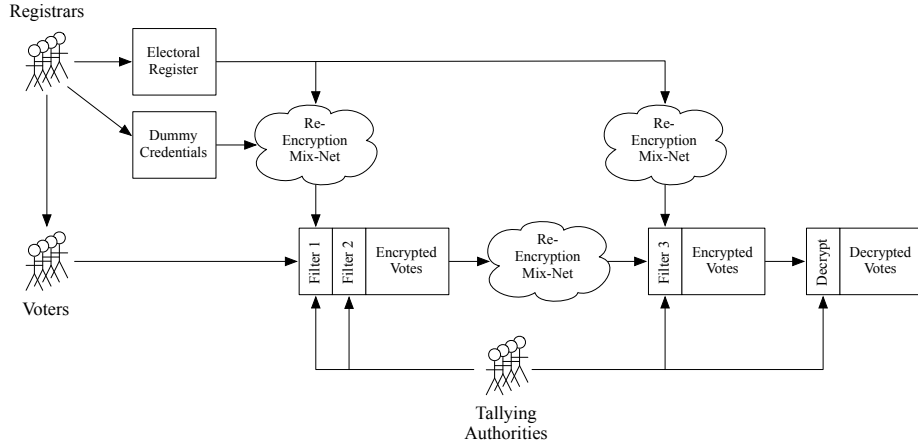


Fig. 2. Overview of the extended JCJ-scheme: the first filter discards votes created from fake credentials, the second filter removes duplicates, and the third filter checks the votes against the electoral register (and thus eliminates votes created from dummy credentials). The final list of proper votes is decrypted and counted.

credentials: the original electoral register $\mathcal{S} = \{S_i : 1 \leq i \leq n\}$ and the new set $\mathcal{T} = \{T_{ij} : 1 \leq i \leq n, 1 \leq j \leq d_i\}$ of dummy credentials. With $\mathcal{ST} = \mathcal{S} \cup \mathcal{T}$ we denote the complete set of encrypted credentials. Furthermore, we denote the number of all issued dummy credentials by $m = |\mathcal{T}| = \sum_{i=1}^n d_i$, which implies that a total of $|\mathcal{ST}| = n + m$ credentials have been issued in all. Those are the ones that will be accepted by the public board during the vote casting phase.

To detect fake credentials for filtering out corresponding invalid votes, the public board needs to check for each incoming vote cast whether the included encrypted credential matches with one of the entries in the list \mathcal{ST} . We can safely apply Smith’s and Weber’s linear-time scheme here, because all votes based on arbitrarily fake credentials are already dropped at this early stage. Note that to install this first filter, we require the help of the talliers already during the voting phase. Compared to the original JCJ-scheme, this is a true disadvantage of our approach, but it is the key for restricting the size of the public board to an upper limit. To do so, the detection and removal of duplicate votes needs to be performed simultaneously, again by applying safely Smith’s and Weber’s scheme and with the help of the tallying authorities. In Figure 2, these tasks of the public board are called “Filter 1” and “Filter 2”, respectively. Note that \mathcal{ST} needs to be shuffled in a verifiable re-encryption mix-net, similar to the shuffling of \mathcal{S} in the original JCJ-scheme. This is important for disguising the links between the voters and their entries in \mathcal{ST} .

The rest of the tallying phase is similar to the original JCJ-scheme, except that the elimination of duplicate votes has already been conducted. Therefore both, the list of encrypted votes registered on the public board and the list \mathcal{S} of encrypted secret credentials, are shuffled in corresponding re-encryption

mix-nets. Respective proofs of correct shuffling are published. The output of the two mix-nets are then used to separate the valid votes from those generated by dummy credentials, again by applying safely Smith’s and Weber’s linear-time scheme. In Figure 2, this task is called “Filter 3”. At the end, the adjusted list of encrypted votes is jointly decrypted and tallied by the tallying authorities.

3.2 Discussion

The above description of the adapted JCJ-scheme outlines the general ideas of our approach. The modifications raise several important questions. Some of them will be discussed below.

How many dummy credentials are needed? To answer this question, suppose first that each voter receives exactly $d \geq 1$ dummy credentials from the registrars, i.e., let $d_i = d$ for all $1 \leq i \leq n$. Each voter would then have d extra credentials to deceive potential vote buyers or coercers. The problem of such a scenario is that the secret credential could only be withheld as long as not all d dummy credentials are “expended”. A coercer could thus force the voter to release all $d + 1$ credentials and use them to cast $d + 1$ identical votes. If all $d + 1$ votes cast appear on the public board, it follows that one of them (the one that includes the proper credential σ_i) will be included in the final tally. Otherwise, if some of the votes cast do not appear on the public board, then the coercer knows that the voter was lying about some of the credentials. Using a similar line of reasoning, votes could be sold by passing all $d + 1$ credential to a vote buyer. Therefore, a constant number of dummy credentials clearly ruins the coercion-resistance property of the scheme.

The above argument leads to the conclusion that the registrars have to generate a varying number of dummy credentials for each voter. Suppose that V_i receives $d_i \in_R \{1, \dots, d\}$ dummy credentials, i.e., d_i is chosen at random between 1 and a fixed upper limit d . If the scheme guarantees that d_i is unknown to potential coercers (which includes the registrars and the tallying authorities), then V_i may lie about d_i , for example by passing all d_i (or less) dummy credentials to the coercer and by claiming that the secret credential is included in that list. This argument works for every V_i with $d_i > 1$, but unfortunately not for those with $d_i = 1$. Under coercion, such (unfortunate) voters could only give away a single dummy credential, but they could not claim it to be the secret credential. Note that this problem does not disappear by increasing the lower limit of the interval $\{1, \dots, d\}$ to some value $c < d$ or by decreasing it to 0. Even worse, a similar problem exists for the upper bound d , because voters with $d_i = d$ dummy credentials could sell their votes by simply handing over all $d + 1$ credentials to the vote buyer (as in the case of a constant number of dummy credentials). This problem could be solved by not imposing an upper limit to the interval, but this brings up new problems of practicability in cases where d_i becomes very large.

As an answer to the above question, we suggest here that d_i , the number of dummy credentials for voter V_i , is determined according to some non-uniform probability distribution over sets $\mathbb{N}^d = \{1, \dots, d\}$ or $\mathbb{N} = \{1, \dots, \infty\}$ of natural

numbers.⁴ The most natural candidate distribution with an upper limit d is a *binomial distribution* $\mathcal{B}(d, p)$ with shape parameters d and p . The idea is to choose d and p such that only a very small fraction of voters get the minimum number ($d_i = 1$) or the maximum number ($d_i = d$) of dummy credentials. This is the case if the variance of the distribution is relatively small compared to d . In this way, we cannot entirely rule out vote selling or coercion, but we can at least make it unbearable for the vast majority of voters.

The most natural candidate distribution with no upper limit is a *normal distribution* $\mathcal{N}(\mu, \sigma^2)$ with some reasonable values for the mean and the variance. Since normal distributions are density functions defined over \mathbb{R} , they need to be applied in some discretized manner over \mathbb{N} . Many other distributions are possible, but a more exhaustive discussion of these questions is beyond the scope of this paper.

How do the registrars generate a random number of dummy credentials? The naïve approach for the registrars to generate a random number of dummy credentials for voter V_i is to jointly apply the chosen probability distribution to determine d_i and to generate each of the d_i dummy credentials using the same distributed procedure as for the secret credential σ_i . The problem of this simple approach is that d_i is not a secret of V_i alone, i.e., V_i cannot lie about it towards potential voter buyers or coercers if they collude with one of the registrars.

As a solution to this problem, we suggest to split up the group of registrars into r sub-groups. Each of these sub-groups is then responsible for generating roughly d_i/r dummy credentials, but without informing the other groups about the exact number. To do so, we need to decompose the chosen probability function into a sum of r probability functions with parameters adapted accordingly. In the case of a binomial distribution $\mathcal{B}(d, p)$, for example, each sub-group may simply use the distribution $\mathcal{B}(d/r, p)$ to determine their numbers independently, because $\mathcal{B}(d, p) = r \cdot \mathcal{B}(d/r, p)$. Similarly, a normal distribution $\mathcal{N}(\mu, \sigma^2)$ can be split up into a sum of r normal distributions $\mathcal{N}(\mu/r, \sigma^2/r^2)$, because normal distributions are closed under linear combination. Note that we need to assume a majority of each sub-group to be trustworthy.

How should the public board store the encrypted dummy credentials?

In the original JCJ-scheme, the encrypted credentials S_i are published on the public board together with the plaintext identities of the voters. The list \mathcal{S} plays thus the role of a *electoral register*, which can be inspected and verified by everybody. By doing the same with the encrypted dummy credentials, i.e., by linking each $T_{ij} \in \mathcal{T}$ publicly with V_i 's identity, we would allow potential coercers or vote buyers to derive the secret number $d_i = |\{T_{ij} \in \mathcal{T}\}|$ from \mathcal{T} . But as already discussed above, coercion-resistance can only be guaranteed as long as d_i is V_i 's secret, since otherwise V_i loses the ability to lie about it.

⁴ We explicitly exclude the borderline case $d_i = 0$, because it would completely disallow V_i to deceive a passive coercer who does nothing but directly observing V_i 's vote casting process (*shoulder surfing* attack).

As a simple solution to this problem, we suggest that the set \mathcal{T} of encrypted dummy credentials is published anonymously without any links to the voters. Since \mathcal{T} does not serve as an electoral register, it does not necessarily need to be treated in exactly the same way as \mathcal{S} . However, this solution is only applicable if deleting entries from the electoral register is prohibited over multiple voting events. Otherwise, additional mechanisms need to be introduced to delete the entries in \mathcal{T} that belong to the deleted entry in \mathcal{S} .

What is the benefit of the modified scheme? The original JCJ-scheme has three critical time-consuming components: the elimination of duplicates, the mixing of the votes in the re-encryption mix-net (as well as the verification of the proofs produced by the mix-net), and the elimination of invalid votes (see Figure 1). The input size of each of these components depends directly on the number of *votes*, not the number of *voters*. Since costly cryptographic computations such as zero-knowledge proofs and multi-party computations are needed to perform these tasks, processing a single additional vote is expensive. Techniques that avoid the processing of votes that will not appear in the final tally are therefore inherently appealing.

Let $n = |\mathcal{S}|$ denote the number of voters (or the number of proper votes if all voters participate in the election) and s the number of duplicate or invalid votes. In the original JCJ-scheme, eliminating duplicates by performing pairwise PETs over all $n + s$ votes requires $O(n^2 + s^2)$ relatively expensive steps. If no duplicates are removed (worst case), $n + s$ is the input size for both the re-encryption mix-net and the final procedure for eliminating fake votes. In the literature of verifiable mix-nets, we find techniques with proofs of linear size [4,10,18], but all of them involve relatively high constant factors. The final elimination of fake votes again requires $O(n^2 + n \cdot s)$ expensive PETs. In total, the JCJ-scheme runs in $O(n^2 + s^2)$ time and $O(n + s)$ space.

Our modified approach improves the overall performance of the scheme in respect of both, computation time and memory space. Note that we have the same three critical components, only arranged in a different order. If $m = \sum_{i=1}^n d_i = |\mathcal{T}|$ denotes the total amount of dummy credentials issued during the registration phase, we get an $O(n + m)$ upper limit for both the size of the public board and the input of the mix-net. As s may become orders of magnitudes larger than m in case of a large scale board flooding attack, this states a major improvement over the original scheme. It prevents situations where the system becomes unavailable due to a memory overflow of the public board.

The modified scheme also eliminates the need for the quadratic number of PETs for eliminating invalid votes. Here we benefit from the methods proposed by Smith [14] and Weber [16,17], which allow duplicate and fake votes to be detected in linear time. Therefore, all components involved in the tallying phase run in $O(n + m)$ time and space, which implies an overall $O(n + m)$ running time for the whole modified scheme. This is a considerable improvement over the original scheme under all possible circumstances.

What is the downside of the modified scheme? In the presented form, our scheme allows statistical attacks which may possibly influence the outcome of a

voting event. For example, a vote buyer may offer a certain amount of money for each additional credential (dummy or proper) handed over by the voter. In this way, a potential vote seller gets a personal interest in handing over as many credentials as possible—including the proper one. An analogous strategy may be applied by the coercer. However, depending on the total number of dummy credentials and the parameters of the chosen distribution function, these attacks may become very cost-intensive for both, the vote buyer and the coercer. This raises the question of finding the optimal distribution function to maximize the cost of such statistical attacks. Answering this question is beyond the scope of this paper and is left for future work.

4 Conclusion

In this paper, we have discussed the board flooding problem in the JCJ-scheme for remote electronic elections. As a solution, we propose that each voter receives a set of dummy credentials along with the proper secret credential during the registration phase. For this enhancement to work, it is important that the number of dummy credentials varies from one voter to another, and that only few voters will get the minimum or maximum number of dummy credentials. The votes posted to the public board can then be filtered such that only votes created from proper or dummy credentials are retained. Duplicate votes are also immediately eliminated in linear time. In this way, there will never be more entries on the public board than the total number of issued (proper and dummy) credentials. The lack of such a filter leads to the board flooding problem in the JCJ-scheme.

This paper is a first step in transforming the impractical JCJ-scheme—applicable under the assumption of unrealistic computing power only—into a practical scheme. Our proposal allows the voting authorities to trade-off the efficiency of the tallying procedure against the obtained level of coercion-resistance. Future work will focus on the residual statistical vulnerability for coercion and vote buying. Currently, we are studying the possibility of obtaining additional dummy credentials on demand during the voting phase, for example by exchanging dummy credentials between voters.

Acknowledgments. Research supported by the *Hasler Foundation* (project No. 09037).

References

1. Araújo, R., Foulle, S., Traoré, J.: A practical and secure coercion-resistant scheme for remote elections. In: Chaum, D., Kutyłowski, M., Rivest, R.L., Ryan, P.Y.A. (eds.) FEE'07, *Frontiers of Electronic Voting*. pp. 330–342. Schloss Dagstuhl, Germany (2007)

2. Araújo, R., N. Ben Rajeb, R.R., Traoré, J., Youssfi, S.: Towards practical and secure coercion-resistant electronic elections. In: Heng, S.H., Wright, R.N., Goi, B.M. (eds.) CANS'10, 9th International Conference on Cryptology And Network Security. pp. 278–297. LNCS 6467, Kuala Lumpur, Malaysia (2010)
3. Clarkson, M.R., Chong, S., Myers, A.C.: Civitas: Toward a secure voting system. In: SP'08, 29th IEEE Symposium on Security and Privacy. pp. 354–368. Oakland, USA (2008)
4. Groth, J.: A verifiable secret shuffle of homomorphic encryptions. *Journal of Cryptology* 23(4), 546–579 (2010)
5. Heather, J., Lundin, D.: The append-only web bulletin board. In: Degano, P., Guttman, J., Martinelli, F. (eds.) FAST'08, 5th International Workshop on Formal Aspects in Security and Trust. pp. 242–256. LNCS 5491, Malaga, Spain (2008)
6. Jakobsson, M., Juels, A.: Mix and match: Secure function evaluation via cipher-texts. In: Okamoto, T. (ed.) ASIACRYPT'00, 6th International Conference on the Theory and Application of Cryptographic Techniques. pp. 162–177. LNCS 1976, Kyoto, Japan (2000)
7. Juels, A., Catalano, D., Jakobsson, M.: Coercion-resistant electronic elections. In: Atluri, V., De Capitani di Vimercati, S., Dingledine, R. (eds.) WPES'05, 4th ACM Workshop on Privacy in the Electronic Society. pp. 61–70. Alexandria, USA (2005)
8. Meister, G., Hühnlein, D., Eichholz, J., Araujo, R.: eVoting with the European citizen card. In: Brömme, A., Busch, C., Hühnlein, D. (eds.) BIOSIG'08, Special Interest Group on Biometrics and Electronic Signatures. pp. 67–78. Darmstadt, Germany (2008)
9. Meng, B.: A critical review of receipt-freeness and coercion-resistance. *Information Technology Journal* 8(7), 934–964 (2009)
10. Neff, C.A.: A verifiable secret shuffle and its application to e-voting. In: Samarati, P. (ed.) CCS'01, 8th ACM Conference on Computer and Communications Security. pp. 116–125. Philadelphia, USA (2001)
11. Peters, R.A.: A Secure Bulletin Board. Master's thesis, Department of Mathematics and Computing Science, Technische Universiteit Eindhoven, The Netherlands (2005)
12. Pfitzmann, B.: Breaking an efficient anonymous channel. In: De Santis, A. (ed.) EUROCRYPT'94, International Conference on the Theory and Applications of Cryptographic Techniques. LNCS 950, vol. 950, pp. 332–340. Perugia, Italy (1995)
13. Schweisgut, J.: Coercion-resistant electronic elections with observer. In: Krimmer, R. (ed.) EVOTE'06, 2nd International Workshop on Electronic Voting. pp. 171–177. Bregenz, Austria (2006)
14. Smith, W.D.: New cryptographic voting scheme with best-known theoretical properties. In: FEE'05, Workshop on Frontiers in Electronic Elections. Milan, Italy (2005)
15. Spycher, O., Koenig, R., Haenni, R., Schläpfer, M.: A new approach towards coercion-resistant remote e-voting in linear time. In: FC'11, 15th International Conference on Financial Cryptography. St. Lucia (2011)
16. Weber, G., Araujo, R., Buchmann, J.: On coercion-resistant electronic elections with linear work. In: ARES'07, 2nd International Conference on Availability, Reliability and Security. pp. 908–916. Vienna, Austria (2007)
17. Weber, S.: Coercion-Resistant Cryptographic Voting: Implementing Free and Secret Electronic Elections. VDM Verlag, Saarbrücken, Germany (2008)
18. Wikström, D.: A commitment-consistent proof of a shuffle. In: Boyd, C., González Nieto, J. (eds.) ACISP'09, 14th Australasian Conference on Information Security and Privacy. pp. 407–421. LNCS 5594, Brisbane, Australia (2009)