

A New Approach Towards Coercion-Resistant Remote E-Voting in Linear Time

Oliver Spycher^{1,2}, Reto Koenig^{1,2}, Rolf Haenni², and Michael Schläpfer³

¹ University of Fribourg, CH-1700 Fribourg, Switzerland
{oliver.spycher,reto.koenig}@unifr.ch

² Bern University of Applied Sciences, CH-2501 Biel, Switzerland
{oliver.spycher,reto.koenig,rolf.haenni}@bfh.ch

³ ETH Zurich, CH-8092 Zurich, Switzerland
michschl@inf.ethz.ch

Abstract. Remote electronic voting has attracted increasing attention in cryptographic research. A promising protocol presented by Juels et al. is currently widely discussed. Although it offers a remarkably high degree of coercion-resistance under reasonable assumptions, it can not be employed in practice due to its poor efficiency. The improvements that have been proposed either require stronger trust assumptions or turned out to be insecure. In this paper, we present an enhancement of the protocol, which runs in linear time without changing the underlying trust assumptions.

1 Introduction

Many governments are aiming at introducing modern technology into their voting processes. Particularly, remote e-voting systems are meant to make voting easier, faster, and more attractive. As appealing as that may seem, introducing physical distance between the voter and the ballot-box comes with a price. Since voters can no longer witness their ballot reach its destination with their own eyes, they need to be provided with another means of assurance. At first sight, this seems to be a simple problem, easily solvable by publishing the set of collected ciphertext votes to let voters verify that their votes have been cast as intended. However, care needs to be taken. Generally, such an approach will allow voters to prove violent coercers or generous vote buyers how they voted. Since voter coercion and vote buying (short: coercion) are highly scalable in an electronic network environment, they need to be prevented. Unfortunately, it seems very difficult to prove voters that their vote is cast as intended (*individual verifiability*), without allowing them to prove others how they voted (*receipt-freeness*).

The protocol underlying this paper was published in 2005 by Juels, Catalano, and Jakobsson [8], often referred to as *the JCJ* protocol. Even today, it seems to be the only known protocol for remote e-voting that offers individual verifiability and receipt-freeness simultaneously under somewhat acceptable trust assumptions. Apart from disabling voters from proving how they voted, the protocol

even ensures immunity against coercers who try to force voters into handing out their credentials (simulation attack) or not casting their votes at all (forced abstention attack). Protocols that avoid all conceivable attacks of coercion are attributed *coercion-resistant*. The JCJ protocol offers a remarkably high degree of coercion-resistance.

Since JCJ imposes unrealistic computational requirements on the tallying authorities, it can not be employed in a real-world context. Nevertheless, the protocol is widely discussed and taken as a starting point for further improvements [2–4, 11, 12]. The ultimate goal of these proposals is to reduce the quadratic running time of the JCJ tallying procedure. We propose our modification of the JCJ protocol to allow tallying in linear time. Section 2, describes JCJ in more detail and points out its security properties and trust assumptions. Section 3 presents our modification of the protocol and shows why the security properties of JCJ are preserved without having to strengthen any trust assumptions. Section 4 concludes the paper and exposes some open questions.

2 The JCJ Protocol

To achieve receipt-freeness, other protocols need to assume an *untappable channel* [10] between authorities and voters at every voting event. Requiring voters to visit the authorities’ offices at each occasion clearly compromises the spirit of remote e-voting. JCJ is distinguished by assuming an untappable channel only during the distribution of the voters’ credentials. Since JCJ allows credentials to be re-used in many subsequent voting events, they can be distributed easily when citizens appear in person at the administration offices to register as new community members.

2.1 Description of the Protocol

In the following paragraphs, we present each phase of the JCJ protocol. Due to space constraints, we settle for a semi-formal style of exposition. In particular, we do not thoroughly explain well-known cryptographic techniques. Furthermore, we assume the application of publicly verifiable group threshold mechanisms whenever registering or tallying authorities perform joint computations, even if the text might suggest a single entity. All ciphertexts are ElGamal encryptions over a pre-established multiplicative cyclic group $(\mathcal{G}_q, \cdot, 1)$ of order q , for which the decisional Diffie–Hellman problem (DDHP) is assumed to be hard.

Registration. The *registrars* jointly establish the random credential $\sigma \in \mathcal{G}_q$ and pass it to voter V through an untappable channel. Additionally, they append a randomized encryption $S = \text{Enc}_\varepsilon(\sigma, \alpha_S)$ of σ to V ’s entry in the voter roll, which is modeled as a public bulletin board. Value α_S denotes the encryption’s randomness, and ε stands for the tallying authorities’ common public key. Assuming a majority of trustworthy registrars, in the end only V will know σ and no one will know α_S .

Vote Casting. Voter V identifies her choice c from the available set of valid choices (or candidates) \mathcal{C} . To cast the vote, she posts the encryptions $A = \text{Enc}_\varepsilon(\sigma, \alpha_A)$ and $B = \text{Enc}_\varepsilon(c, \alpha_B)$ to the public bulletin board, through an anonymous channel. The pair (A, B) must be accompanied by two non-interactive zero-knowledge proofs (NIZKP), one to prove knowledge of σ and one to prove $c \in \mathcal{C}$. Requiring the first proof prevents attackers from casting unauthorized votes by re-encrypting entries from the voter roll (recall that α_S is not known to anyone). Since each authorized vote on the voting board will be decrypted during the tallying phase, the second proof is needed to prevent coercers from forcing voters to select $c \notin \mathcal{C}$ according to some prescribed pattern, thus obtaining a receipt [6]. To circumvent coercion, the voter can deceive the coercer by posting a fake vote to the voting board. To do so, V simply claims some $\sigma' \in \mathcal{G}_q$ to be her real credential and uses it to compute A . She computes B according to the coercer’s preference and reveals the plaintexts of A and B to justify compliance. Alternatively, V can even let the coercer compute A and B and cast the vote using σ' .

Tallying. At the end of the vote casting phase, the voting board contains N posted votes, of which not all must be counted. First, the *talliers* verify all proofs that were cast along with the votes. If a proof does not hold for a vote (A, B) , it is marked accordingly and excluded from further processing. Then the talliers need to filter out votes that were cast multiple times with a proper credential and votes that were cast with a fake credential. For both tasks, the authors of JCJ propose the application of a *plaintext equivalence test* (PET) [7]. Given two ElGamal encryptions $X = \text{Enc}_\varepsilon(x, \alpha_X)$ and $Y = \text{Enc}_\varepsilon(y, \alpha_Y)$, the algorithm $\text{PET}(X, Y)$ returns *true* for $x = y$ and *false* for $x \neq y$, without revealing any information on x or y .⁴

Removing Duplicates. Exclude from further processing all (A_i, B_i) , for which the voting board contains (A_j, B_j) , $i \neq j$, such that $\text{PET}(A_i, A_j)$ returns *true*. Given that the voting board contains the votes in the order as cast, a “last-vote-counts” (“first-vote-counts”) policy is implemented by starting the search with big (small) values j . This exhaustive search over the entire voting board of size N runs in $\mathcal{O}(N^2)$ time.

Removing Invalid Votes. Invalid votes could easily be excluded from the tally by applying $\text{PET}(S_i, A_j)$ in an exhaustive search over all values S_i of the voter roll and all values A_j of the voting board, similarly to the previous step. However, that would allow the voters to prove the coercer how they voted. To prevent that, the voter roll and the voting board are mixed and re-encrypted using a verifiable re-encryption mixnet, resulting in values \widehat{S}_i and $(\widehat{A}_j, \widehat{B}_j)$, respectively. Now talliers compute $\text{PET}(\widehat{S}_i, \widehat{A}_j)$ for all pairs \widehat{S}_i and \widehat{A}_j . If the algorithm returns *true* for some index i , \widehat{B}_j is decrypted and counted in the tally. This procedure runs in $\mathcal{O}(N \cdot n)$ time, where n denotes the size of the voter roll.

⁴ A common way of performing PET in a homomorphic encryption scheme is to check whether the decryption of $(X/Y)^z$ equals 1 for some random value $z \in \mathbb{Z}_q$.

If the voting board is flooded with a large number of fake votes, N may be orders of magnitudes larger than n , which implies that the JCJ tallying procedure has an $\mathcal{O}(N^2)$ worst-case running time (quadratic with respect to the number of votes). This makes the scheme not only vulnerable to denial-of-service attacks, but also practically infeasible in large-scale settings. The authors of Civitas, a running prototype implementation based on JCJ, have shown this in [5].

2.2 Security Properties and Assumptions

We briefly want to point out, to which degree JCJ satisfies the key requirements *privacy* and *accuracy*, and why JCJ provides a high level of coercion-resistance. Privacy is motivated by the notion of the secrecy of the vote. It is satisfied if no vote can be linked to the voter from whom it originates. Accuracy captures the notion that all (and only) legitimate votes are tallied as cast.

Privacy. With respect to privacy, JCJ relies on the security of the anonymous channel and the trustworthiness of the tallying and mixing authorities. Since a majority of tallying authorities could collude to jointly decrypt entries of the voter roll and the voting board, they could easily break privacy. Similarly, the mixing authorities could violate privacy by jointly establishing a link from the decrypted votes back to the voter roll. In both cases, the violation of privacy could be hidden by the conspiring parties.

Accuracy. By observing the voting board, voters verify that their vote has been cast as intended. Changing or removing votes from the tally would be detected by the public. Adding illegitimate votes requires the knowledge of a credential σ that complies with a value S in the voter roll. Since all values S are related to a voter enlisted in the voter roll, adding an illegitimate value could be noticed by voters that are about to register. Attacks of that kind are thus not scalable. As pointed out in the previous paragraph, a colluding majority of authorities could secretly decrypt S to obtain V 's valid credential σ . However, if they use σ for casting votes, they could be exposed by V when the corresponding PET algorithm returns *true* at removing duplicates during the tallying procedure.

Coercion-Resistance. Assuming that the coercer cannot communicate with the registrars, voters can always lie about their credentials σ . They are thus protected against coercers that want to push them into voting in a prescribed way, voting at random, or handing out their credentials. If the coercer wants V to abstain from voting, V can still cast a vote, given at least one moment of privacy. As pointed out before, we allow a minority of authorities to be untrusted. Disallowing communication between the coercer and all registrars would strengthen that assumption. However, allowing communication would enable the coercer to force the voter into handing out the proper credential: The coercer could claim knowledge about the secret share that a colluding registrar has provided to V , without saying which one. To be safe, V needs to hand out all secrets truthfully. We therefore need to assume that the voter knows at least one registrar not colluding with the coercer. Note that this is not implied by assuming *any* majority of trustworthy registrars. Thus, V can lie to the coercer about that secret.

3 Coercion-Resistance in Linear Time

We pointed out that the steps to remove duplicate and illegitimate votes are inefficient in the original JCJ protocol. This issue is widely discussed and has been addressed in the literature. Before presenting our enhancement of the scheme, we introduce two highly promising known approaches, that also aim at improving efficiency at tallying.

Scheme by Smith and Weber [11–13]. Instead of applying $\text{PET}(A_i, A_j)$ on all pairs of distinct ciphertexts for removing duplicates, $1 \leq i, j \leq N$, both Smith and Weber in essence suggest computing and decrypting $A_1^z = \text{Enc}_\varepsilon(\sigma_1^z), \dots, A_N^z = \text{Enc}_\varepsilon(\sigma_N^z)$, where $z \in \mathbb{Z}_q$ is a random value shared among the talliers. The resulting *blinded* values 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 illegitimate votes. In that case, however, based on the fact that the same exponent z is used across all ciphertexts, the coercer gets an attack strategy to identify whether a vote with known σ is counted [2, 5, 9]. Note that this attack does not apply at removing duplicates.

Scheme by Araujo et al. [1–3]. To solve the efficiency problem of the JCJ scheme, the authors suggest an approach based on group signatures. At registration, voters obtain their credential. Unlike JCJ, no public values are related to voter roll entries. Their credentials enable the voters to deduce invalid credentials and mislead coercers. If the provided proofs hold, duplicates on the voting board are publicly identifiable by the equality of two values that are cast along with the vote. After mixing the relevant values on the voting board, the tallying authorities use their private keys to identify the legitimate votes. Notably, all information on their legitimacy is sent along with the vote itself, but can only be assessed by a sufficiently large group of talliers. Fully avoiding matches between cast values and voter roll entries summarizes the essence of this elegant approach to avoid the inefficient comparison procedure.

An inherent weakness of this approach is the fact that a majority of colluding registrars could compute valid (but illegitimate) credentials unnoticed. As described earlier, adding illegitimate votes to the tally in JCJ requires the knowledge of a credential σ that complies with an entry S in the voter roll, i.e., such attacks could easily be detected. This is not the case in Araujo et al.’s scheme. Nevertheless, we believe that the approach holds much potential.

3.1 Description of the Enhanced Protocol

Our enhancement strongly relates to the original JCJ, so the modifications are easily summarized. For removing duplicates, we propose using the linear-time scheme proposed by Smith and Weber. For identifying the legitimate votes, we suggest preserving the use of the voter roll. The key to efficiency lies in requiring voters to indicate which voter roll entry their vote (A, B) relates to. Talliers then apply PET only on respective re-encryptions of A and S , where S is the public value copied from the indicated voter roll entry. Authorizing legitimate

votes thus becomes linear over the total number of cast votes. In the following paragraphs, we present the protocol in further detail. Later we justify why the security properties of JCJ are preserved under unchanged trust assumptions.

Registration. The registration step is conducted according to JCJ. Additionally, we assume that a distinct public number i is assigned to each voter. For simplicity, we take i to be the index of V 's entry in the voter roll.

Vote Casting. To cast a vote, V performs the same steps as in JCJ. Additionally to posting values A and B along with corresponding proofs, V posts the value $C = \text{Enc}_\varepsilon(i, \alpha_C)$, accompanied by a non-interactive zero-knowledge proof to prove knowledge of i . The tallying authorities will later use i to locate S_i on the voter roll and efficiently detect legitimate votes. Note that the voting board must also accept wrong values $C \neq \text{Enc}_\varepsilon(i, \alpha_C)$.

Tallying. After excluding votes with invalid proofs, the talliers add a random number X_i of additional fake votes for each voter (see discussion below). After removing duplicates by applying Smith's and Weber's scheme on values A_i , the resulting adjusted list is passed as input to a first re-encryption mixnet, which outputs tuples $(\widehat{A}_j, \widehat{B}_j, \widehat{C}_j)$. Next, the talliers jointly decrypt \widehat{C}_j into i and establish a list of tuples $(\widehat{A}_j, \widehat{B}_j, S_i)$. Votes for which the decryption renders an invalid index $i \notin \{1, \dots, n\}$ are excluded from further processing. The remaining tuples $(\widehat{A}_j, \widehat{B}_j, S_i)$ are then passed to a second re-encryption mixnet, which outputs tuples $(\widetilde{A}_j, \widetilde{B}_j, \widetilde{S}_i)$. Now the talliers perform $\text{PET}(\widetilde{A}_j, \widetilde{S}_i)$ for each tuple. If the algorithm returns *true*, \widetilde{B}_j is decrypted and counted.

The generation of additional fake votes is important to conceal the existence of a real vote after employing the first mixnet, where the encrypted voter roll indices i are decrypted. The presence of fake votes at that point enables voters to repudiate the fact of having posted a valid vote. The fake votes must be generated and published by trustworthy authorities, such that the exact number of fake votes for voter V is not revealed. We will later argue that it is sufficient if the registrar, who enjoys V 's trust, posts a random number $X \geq 1$ of fake votes designated to V . Clearly, if X is independent of N for all voters, then our tallying procedure runs in $\mathcal{O}(N)$ time (provided that each of the two mixnets runs in linear time).

3.2 Security Properties and Assumptions

In JCJ, coercion-resistance is based on the voter's ability to lie about σ and secretly cast the real vote in a private moment. Note that by witnessing the voter casting a vote, the coercer will need to assume that the voter did not reveal the proper credential. Thus, that one moment of privacy is required in any mode of coercion, not only in the event of a forced abstention attack. The voter may then claim not having cast *any* vote, except possibly the one instructed by the coercer, posted with a fake credential. We will now argue why this argument yields coercion-resistance in our scheme as well.

During the vote casting phase, the voting board reveals no more information to the coercer than in JCJ. During the tallying phase, however, the coercer learns how many votes are related to V 's entry in the voter roll. Let x denote that number. The coercer's strategy is to decide, whether x is distributed according to the same random distribution as for the other voters or if it is greater by one.

Taking the Risk. To decide whether x originates from X or $X + 1$ based on one sample (or even a few in case of repeated coercion in subsequent voting events) seems hardly feasible. This conjecture is additionally supported by the fact that other voters may also attribute fake votes to V . We believe that V is likely to take the small risk of being caught, in case V is not exposed to the risk of being punished. If the coercer needs to assume that voters will generally not fear getting caught, any coercion attack seems obsolete. Therefore, we are confident that this notion suffices for solving the vote buying problem, which is the only concern in many countries regarding the notion of coercion.

Understanding the Risk. The more voter V fears consequences in the event of getting caught, the more important it becomes to quantify the risk. V will agree to co-operate with the coercer unless V is actually convinced that the risk of getting caught is vanishingly low. That it is infeasible to decide whether x originates from X or $X + 1$ given a distribution function F_X with a high standard deviation, even over reasonably small values, is a hypothesis we will quantify in our future work. For the time being, we relate the problem to an analogy in JCJ.

We thus give an idea of how the distribution function F_X for determining the random value X needs to be defined, in order to make our scheme as secure as JCJ. Recall that in JCJ, V needs a time-slice Δt of privacy for casting the real vote. Note that the coercer may monitor the voting board at the beginning and at the end of that time-slice. Let x' denote the number of votes posted during that time. The coercer's strategy is to decide whether x' is distributed according to the same random distribution as the other values x'_i of the remaining $m - 1$ time-slices of equal length (if T denotes the total length of the voting period, we have $m = \frac{T}{\Delta t}$). If JCJ is coercion-resistant for defined n and m , an average of $x' = \frac{N}{m}$ votes is sufficient to disguise V 's additional vote (or $x' = \frac{n}{m}$ assuming all voters participate and post one vote only). V obviously enjoys the same protection in our scheme, if the registrar's random function F_X produces an average of $x = \frac{n}{m}$ fake votes. The same distribution function F_X can then be employed in the case of a greater number of voters n .

4 Conclusion and Future Work

We have shown a new protocol that solves the efficiency problem in coercion-resistant remote e-voting, without changing the trust assumptions and the security features of JCJ. We pointed out that coercion-resistance of both JCJ and our scheme assumes a private moment for voters to secretly cast their vote. This is only sufficient if the coercer can not deduce from the voting board whether voters took advantage of their privacy. We have related this problem to distin-

guishing whether x is distributed as X or $X + 1$ for reasonably small values and a high standard deviation of F_X , where there is one sample per voting event. In our future work, we will show formally that this problem is infeasible to solve and thus justify JCJ and our scheme to be sufficiently secure against coercion attacks under the known assumptions.

Acknowledgments. Research supported by the *Hasler Foundation* (project No. 09037), the *Mittelbauförderung* of the Bern University of Applied Sciences, and the Swiss Federal Chancellery.

References

1. Araujo, R.: On Remote and Voter-Verifiable Voting. Ph.D. thesis, Department of Computer Science, Darmstadt University of Technology, Germany (2008)
2. Araújo, R., Foulle, S., Traoré, J.: A practical and secure coercion-resistant scheme for remote elections. In: Chaum, D., Kutylowski, M., Rivest, R., Ryan, P. (eds.) FEE'07, Frontiers of Electronic Voting. pp. 330–342. Dagstuhl, Germany (2007)
3. 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)
4. Clark, J., Hengartner, U.: Selections: Internet voting with over-the-shoulder coercion-resistance. In: FC'11, 15th International Conference on Financial Cryptography. St. Lucia (2011)
5. 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 (2008)
6. Di Cosmo, R.: On privacy and anonymity in electronic and non electronic voting: the ballot-as-signature attack. Hyper Articles en Ligne hal-00142440(2) (2007)
7. Jakobsson, M., Juels, A.: Mix and match: Secure function evaluation via ciphers. ASIACRYPT'00, 6th International Conference on the Theory and Application of Cryptographic Techniques. pp. 162–177. LNCS 1976, Kyoto, Japan (2000)
8. 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)
9. Pfizmann, 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)
10. Sako, K., Kilian, J.: Receipt-free mix-type voting scheme: A practical solution to the implementation of a voting booth. In: Guillou, L.C., Quisquater, J.J. (eds.) EUROCRYPT'95, 15th International Conference on the Theory and Applications of Cryptographic Techniques. pp. 393–403. LNCS 921, Saint-Malo, France (1995)
11. Smith, W.D.: New cryptographic voting scheme with best-known theoretical properties. In: FEE'05, Workshop on Frontiers in Electronic Elections. Milan (2005)
12. 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)
13. Weber, S.: Coercion-Resistant Cryptographic Voting: Implementing Free and Secret Electronic Elections. VDM Verlag, Saarbrücken, Germany (2008)