Achieving Meaningful Efficiency in Coercion-Resistant, Verifiable Internet Voting

Oliver Spycher[1], Reto Koenig[2], Rolf Haenni[2], Michael Schläpfer[3]

[1]University of Fribourg
1700 Fribourg, Switzerland
oliver.spycher@bfh.ch

[2]Bern University of Applied Sciences
2501 Biel, Switzerland
{reto.koenig | rolf.haenni}@bfh.ch

[3]ETH Zurich
8092 Zurich, Switzerland
michschl@inf.ethz.ch

Abstract: In traditional voting schemes with paper, pens, and ballot-boxes, appropriate procedures are put in place to reassure voters that the result of the tally is correct. Considering that in Internet voting errors or fraud will generally scale over a much greater fraction of votes, the demand to get strong reassurances as well, seems more than justified. With the ambition of offering a maximum degree of transparency, so-called *verifiable* schemes have been proposed. By publishing the relevant information, each voter may verify that her vote is included in the final tally and that accepted votes have been cast using proper voting material. Remarkably, this can be done while guaranteeing the secrecy of the ballot at the same time. On the negative side, high transparency will generally make it easier for voters to reveal how they voted, e.g., to a coercer. In this paper we propose an Internet voting protocol that is verifiable and simultateously makes it practically impossible for vote buyers or coercers to elicit the voters' behaviour. We compare its efficiency with existing work under equal degrees of coercion-resistance using an appropriate measure (δ). The contribution of our scheme lies in its efficiency during the most critical phases of the voting procedure, i.e., vote casting and tallying. Moreover, during these phases, efficiency is insensitive to the desired degree of coercion-resistance.

1 Introduction

The secrecy of the ballot serves as a means to protect citizens from external influence that pressures them into casting a vote that does not reflect their personal preference. The key to protecting the secrecy of the ballot lies in preventing citizens from revealing to others how they voted. In traditional, paper-based schemes, precautions may require voters to fill out their ballots on-site, often in an isolated booth. Thus voters get the privacy it takes to render any information they take out of the polling station meaningless. Particularly, they cannot provide a coercer with a *receipt*, i.e., the information it takes to reveal the ballot they cast. In Internet voting, the quest for receipt-free, voter-verifiable systems is still ongoing. In a first phase, some propositions have been made that rely on strong assumptions, such as the existence of untappable channels [HS00] prior to the voting event. (In practice voters would need to register in person each time they are asked to vote using the Internet.) In 2005, Juels et al. achieved a breakthrough by proposing a receipt-free and yet verifiable protocol under strongly reduced trust assumptions [JCJ05] (henceforth referred to as *the* JCJ protocol). Remarkably their scheme is not only receipt-free but also highly resistant to coercers who want to push voters into handing out their credentials, voting at random, or abstaining from casting a ballot. Schemes that succeed at circumventing these coercion attacks are called coercion-resistant.[1] For putting these advances in security into practice, Juels et al. still need to make strong assumptions regarding the computational power of the tallying servers. Such assumptions make implementing JCJ infeasible for large-scale elections, as shown in [CCM08].

Since 2005 there have been a number of propositions that take the work of Juels et al. as a starting point and want to make coercion-resistant Internet voting practical while also preserving the security features of JCJ [Ar08, ABR10, CH11, SKH11, SHK11]. With one exception, the propositions are configured to achieve high degrees of coercion-resistance at the cost of efficiency.[2] The price is always paid by either the voter or the tallying servers, which still have to perform lots of computing. This paper also proposes a protocol that is parameterizable regarding coercion-resistance. However, the price for a high degree of coercion-resistance is only paid during the setup-phase, i.e. the phase which is the least time critical. Notably, the computations related to the set-up phase specific to a vote only (*post-registration*) needs to be completed only after the last vote has been cast. We may expect voting phases to be typically long enough for post-registration to be completed, thus allowing the first vote to be cast just after the last voter has registered. Casting votes is just as fast as in JCJ, and tallying becomes drastically faster. We hereby address the general notion that user-friendliness and the possibility to obtain the election results early are preconditions for the successful introduction of Internet voting.

In Section 2, we provide an explanation of how coercion-resistance can be measured and how the JCJ protocol is considered coercion-resistant. After presenting our protocol, in Section 3 we compare its efficiency with the known proposals from the literature in Section 4 . Finally we make concluding remarks in Section 5.

2 Quantifying Coercion-Resistance

There are a variety of definitions for coercion-resistance. [KTV10] gives a nice overview of the various approaches. In their 2005 protocol proposition, Juels et al. included their own particular notion. The paper proves the protocol to be coercion-resistant in terms of their definitions. Subsequent JCJ-related protocols that were introduced under a formal view on coercion-resistance, have essentially done so using this model or one with slight technical adaptations.

[1] As it is common in the technical literature, we do not distinguish between vote buyers (people who give) and coercers (people who take). As far as we are concerned, a coercer is an algorithm designed to obtain the information it takes to reveal whether a voter has adhered to some predefined instructions.

[2] The only exception is the protocol proposed in [ABRTY10]. However, the scheme does not provide the same degree of verifiability as JCJ. This special case will be revisited in the context of Section 3.4 and Section 4.

All proposed protocols foresee the same defense strategy for the voter subjected to coercion: She hands out a fake credential to the adversary and casts the ballot of her choice through the anonymous channel using her real credential. In short, according to JCJ a protocol is coercion-resistant if an active, non-adaptive adversary cannot distinguish between dealing with the defense strategy and obtaining the real credential with a non-negligible probability of success. In order to prove the coercion-resistance of the JCJ protocol, the authors need to assume that along with the published result, the difference Γ between the number of cast votes n and the number of the ones that are actually counted (due to using a valid voting credential) gives the adversary no advantage in succeeding with coercion (*adversarial uncertainty*). As we will argue, adversarial uncertainty will always be low enough to allow coercion, even without any quantitative prior knowledge regarding Γ.

In [KTV10], Küsters et al. introduce their notion of a measure for quantifying coercion-resistance. They define the degree of coercion-resistance δ as the probability that the (reasonable) adversary will accept a run given that the voter submits to coercion minus the probability that the adversary will accept a run given that the voter applies the defense strategy.[3] They point out that there are opportunities of coercion already on the base of the expected and the effective tally, i.e., attacks that apply even in an ideal system. In that sense, JCJ seems justified in assuming adversarial uncertainty with regard to the expected tally. However Γ is a value specific to coercion-resistant Internet voting schemes. On one hand, since these schemes are not yet in practice, adversarial uncertainty with regard to Γ is to be expected in real life. On the other hand, since voters are also uncertain about Γ, the coercer can still launch an attack based on a wild guess $\Gamma = c$: he can offer money in case $\Gamma \leq c$ or scratch the car if $\Gamma > c$. The reasonable voter will then submit to coercion if she believes that the vote cast with the fake credential would cause Γ to exceed c by 1. Since in a scheme that is meant to be coercion-resistant there is no reason to actually take advantage of using fake credentials, c might initially be chosen relatively small, thus yielding a correspondingly high δ.

Given the exclusion of Γ from adversarial uncertainty, some parameterizable, JCJ-related protocols can be configured to achieve a degree of coercion-resistance that depends solely on the estimated Γ. However, in this case, the parameters have to be chosen such that no meaningful gains in efficiency as compared with JCJ remain. In any case, it seems that accelerating JCJ through parameterization inherently comes along with some loss in coercion-resistance. Nevertheless, this needs to be considered legitimate, knowing that JCJ would not have been considered coercion-resistant if adversarial uncertainty regarding Γ hadn't been assumed. Finally, it cannot be estimated whether coercion based on Γ promises less success than coercion based on the loss of coercion-resistance inherent to accelerating JCJ.

[3] If a vote buyer offers a voter 100 dollars for a vote when using a system that doesn't allow a defense strategy, the voter may expect to get the full reward when submitting to coercion and nothing otherwise. Intuitively speaking, δ signifies the fraction of the 100 dollars voters may on average expect to additionally get from a vote buyer when submitting to coercion as opposed to applying a defense strategy in a δ-coercion resistant system. Obviously, small δ values are what we are looking for.

The protocol we are about to introduce is δ-coercion resistant in a parameter β. We will compare its performance with others under parameters β that yield equal degrees of coercion-resistance δ, where δ signifies the reduction of coercion-resistance compared with the JCJ-protocol. Remarkably, unlike Γ, we are able to quantify δ for each of the protocols.

3 Protocol

Due to space constraints, we are not able to introduce JCJ beforehand. Instead we will indicate relevant divergencies from JCJ within our exposition. Due to the strong relation between both protocols, we find this approach to be justified. After showing the basic idea behind our protocol in Section 3.1 and presenting the applied cryptographic primitives in Section 3.2 , in Section 3.3 we start off by introducing a basic version of our protocol. It already holds strong security features. In Section 3.4 we will propose some slight enhancements to improve verifiability. We chose this step-by-step approach for the sake of readability. We will informally justify the δ-coercion resistance within the exposition of our protocol, i.e., assuming the ideality of the applied cryptographic primitives. The formal security proof is left for future work.

3.1 The Idea

Our scheme foresees the same defense strategy for voters under coercion as JCJ and the other well-known, verifiable, coercion-resistant protocols from the literature: they hand out an invalid credential and cast a vote to the public bulletin board ($\mathcal{PB}$) using their real credential. The protocol should not enable the coercer to decide whether an invalid or a real credential was obtained, despite verifiability. Evidently this requires that the voters' be able to cast votes to the $\mathcal{PB}$ an arbitrary number of times, regardless of whether using real or invalid credentials.[4] As a consequence, the $\mathcal{PB}$ may contain multiple votes cast using the same credential as well as votes cast with an invalid credential. Thus all coercion-resistant protocols need to include steps to *remove duplicates* and *authorize votes* prior to decryption.

As in JCJ, our protocol divides the authorities put in charge of the voting system among *registrars* and *talliers*. Regarding corruption by a coercive adversary, we advise the reader to assume all registrars and a majority of talliers are trustworthy. This could be weakened by requiring that all registrars be trustworthy only during the registration step and during the other phases by assuming that each voter knows a registrar who will not participate in a coercive attack against the voter. This weakening requires no change to the proposed protocol and the reasoning strictly follows [JCJ05]. Regarding *verifiability* (defined in [JCJ05] as *strong verifiability*) none of the authorites need to be trusted. The definition requires voters to be able to detect the exclusion of legitimate votes, changes to legitimate votes, and the inclusion of multiple votes cast with the same credential. In Section 3.4, we will change this definition as well as give more power to voters during verification under the notion of *improved verifiability* (the features of which are also mentioned in [JCJ05] though not formalized), e.g., voters can additionally verify that all credentials used to cast votes are assigned to eligible voters, whereas the basic protocol

[4] If the number of accepted votes were limited, the coercer could test the received credential for validity by counting the number of times he can use it to cast a vote.

would only allow voters to verify this given respective trustworthy majorities of registrars and talliers. In order to achieve *improved verifiability* in the full protocol, we will enhance the basic protocol in Section 3.4 accordingly. The conclusion will be that our scheme reaches δ-coercion resistance and a degree of verifiability equal to the JCJ scheme, notably under equal assumptions regarding the authorities and adversarial power. After showing the applied primitives, we are ready to introduce our protocol.

3.2 Cryptographic Primitives

The new scheme applies the following cryptographic primitives: the ones not employed by the JCJ protocol are identified accordingly. In justifying coercion-resistance and verifiability in the course of our exposition, we assume primitives to be ideal.

Multi-party ElGamal Cryptosystem with Threshold. We propose all ciphertexts to be ElGamal over a pre-established multiplicative cyclic group $(\mathcal{G}_q, \cdot, 1)$ of order q, for which the decisional Diffie-Hellman problem (DDHP) is considered to be hard.[5] Assuming no decryption, ElGamal ciphertexts are not meant to disclose any information in the encrypted plaintext, even in the event that the plaintext space is small and in the presence of other ciphertexts.

We also propose the application of a multi-party computation scheme derived from [Pe91, GJK99] to preserve the confidentiality of encrypted values throughout the protocol. Thus, malicious decryption is only possible in the event of a conspiring majority (the number depends on the chosen threshold) of group members, i.e., registrars or talliers.

Verifiable Mix-Nets. Trustworthy mix-nets take an ordered set of ciphertexts and output re-randomized encryptions in a random order such that the link is not able to be retrieved. They are implemented as a sequence of shuffles, each performed by a distinct mix-node. The link between elements from input and output is only retrieved in the event of all nodes conspiring. Correctness of execution is proven using NIZKP.

[5] We thus follow Civitas [5], which basically instantiates the JCJ protocol. However they do deviate in the choice of the underlying cryptosystem. The reason behind JCJ choosing a modified version of ElGamal (M-ElGamal) lies in the reasoning of their security proof. Although we could allow our protocol to adopt M-ElGamal as well, we adhere to ElGamal, thus making its performance more easily comparable to most of the other known proposals for coercion-resistant Internet voting. Furthermore, the question whether to choose ElGamal or M-ElGamal does not seem sensitive to the design of a particular verifiable voting protocol but rather to the desired security reassurances of the cryptosystem itself. Notably, ElGamal has recently been proven to have the beneficial IND-CCA1 property (resistance against non-adaptive chosen ciphertext attacks) just as much as M-ElGamal [Li11]. Underlying our informal security argumentation within the protocol description, we assume that the plaintexts of all ciphertexts are unconditionally hidden, even when the plaintext space is restricted, and given the ideality of the remaining primitives.

Plaintext Equality Test PET. Given two ElGamal encryptions E_1 and E_2, the algorithm returns $true$ if the plaintexts are equal and $false$ otherwise. This is done by checking whether the decryption of $(E_1/E_2)^z$ equals 1 for a random value $z \in \mathbb{Z}_q$. [JJ00] PET is verifiable and reveals no non-negligible information on the plaintexts.

Additional Primitive M-PET. Unlike JCJ, the new scheme relies on an additional method for efficiently testing the equality among the elements encrypted by a set of ciphertexts as described in [We08]. Clearly, applying PET pair-wise on all elements of the set would result in quadratic runtime. This is exactly the approach chosen in the JCJ protocol and the reason for its inefficiency during the tallying stage.

Given ciphertexts $X_1, \ldots X_n$, the modified PET (M-PET) raises all values to a random value $z \in \mathbb{Z}_q$, and decrypts them to obtain the blinded plaintexts $x_1^z = DEC(X_1^z), \ldots, x_n^z = DEC(X_n^z)$. The blinded plaintexts can be efficiently compared for equality, for instance, by sequentially saving them in a hash table. If a hit is made, the algorithm returns as $true$ and as $false$ otherwise. M-PET doesn't reveal any non-negligible information on the plaintexts, given that the discrete logarithm of any plaintext x_i is unknown in the base of any plaintext x_j, $1 \leq i < j \leq n$.

Communication Channels. There is a public board PB which is used as a *public broadcast channel*. Voters post their votes to PB and the authorities post all output of the tallying phase to PB. For the sake of simplicity we also assume that all public information, including public values from the employed PKI, is accessible on the PB. Further there is an *untappable, authenticated channel* from the registrars to the voters to hand the voters their credentials. Finally an anonymous channel is in place to allow one cast votes anonymously to the PB.

Non-Interactive, Zero-Knowledge Proofs NIZKP. To provide verifiability, many computations throughout the protocol need to be paired with with non-interactive zero-knowledge proofs. These proofs allow voters to prove knowledge of a plaintext by proving plaintext membership of a given sub-domain of $\mathcal{G}_q$, authorities can also prove the correct execution of PET, M-PET, correct mixing, encryption and decryption. We rely on the Fiat-Shamir heuristic for secure non-interactivity, i.e., negligible knowledge-errors and overwhelming witness-hiding.

3.3 Basic protocol

Pre-Registration. The talliers jointly establish a multi-party ElGamal threshold PKI, publish their public key ε on the PB, and keep their shares of the corresponding private key to themselves. The registrars jointly establish a number of $\beta \cdot N_+$ random credentials, where β denotes the security parameter underlying the degree of coercion-resistance δ, and N_+ denotes the maximum expected number of individual voters ever to participate at elections hosted by the voting system. The credentials are tuples of the form (σ, i), whereas we use the terms σ-credential and i-credential to refer to the respective components. Each component is random from $\mathcal{G}_q$ and only computable if the registrars maliciously co-operate. They jointly encrypt and post each of the two components $(\mathrm{E}_\varepsilon(\sigma, \alpha_\sigma), \mathrm{E}_\varepsilon(i, \alpha_i))$ on the PB and memorize their share of the randomnesses α_σ and α_i, both random from Z_q. We call the resulting list of encrypted credential components the *credential pool*. Finally, they pass all $E_\varepsilon(i, \alpha_i)$ through a mix-net and the talliers decrypt the output to form the list $\mathcal{UNL} < i >$, i.e., the list of i-credentials, the elements of which are unlinkable to the *credential pool* by the coercer. The pre-registration step is needed only prior to the first election hosted by the voting system. Since valid i-credentials need to be made public later in the protocol, the list $\mathcal{UNL} < i >$ is meant to enable voters, as in JCJ, to lie about their credentials directly after registering. The *credential pool* however will be processed at a later stage to allow the exclusion of votes cast with an invalid credential.

Registration. The voter roll is initialized as an empty list on the PB. After successful authentication for registration, the registrars choose an unassigned ciphertext tuple from the *credential pool* and post it to the voter roll along with an identifier of the voter. They hand voters their credential (σ, i), along with a proof that the credential corresponds with the ciphertext tuple. As with all computations by registars and talliers, this procedure is conducted by the means of multi-party computation, such that only a malicious collusion can compute the secret, i.e., the plaintexts. The proof is implied by one proof from each registrar computed by the respective partial knowledge of the randomness of α_σ and α_i. Finally, the voter secretly chooses the random elements $\hat{\sigma} \in \mathcal{G}_q$ and $\hat{i} \in \mathcal{UNL} < i >$. Whenever the coercer asks the voter to hand out her credentials, she can lie and hand out $(\hat{\sigma}, \hat{i})$. In the basic version of the protocol, the *voter roll* only serves as a reference for locating the unassigned credentials from the *credential pool* and for identifying the credentials to be retained in case voters lose eligibility.

Post-Registration. The registrars pass all the ciphertext tuples $(\mathrm{E}_\varepsilon(\sigma, \alpha_\sigma), \mathrm{E}_\varepsilon(i, \alpha_i))$ of the *credential pool* to a mix-net. From the output, the talliers decrypt the second component, the ciphertexts containing i-credentials. We call the resulting list $\mathcal{UNL} < E_\varepsilon(\sigma), i >$, as the coercer cannot link its elements to the credential-pool or to the non-anonymous voter roll. The post-registration step needs to be completed only prior to tallying, i.e., the phase in which voters cast their votes can be used for this step. Thereby the negative impact of the time-consuming mix-nets is mitigated, or even fully compensated, given that the voting phase is sufficiently long.

Vote Casting. The voter selects the representation c of her prefered candidate(s) from a set $\mathcal{C} \subset \mathcal{G}_q$, which we assume to be available on the PB. To cast the vote, she uses the anonymous channel and posts the two ciphertexts $A = \mathrm{Enc}_\varepsilon(\sigma, \alpha_A)$ and $B = \mathrm{Enc}_\varepsilon(c, \alpha_B)$ to the voting board on the PB, along with her i-credential in plaintext. The voter aditionally needs to post one non-interactive, zero-knowledge proof (NIZKP) per ciphertext. The first one requires voters to prove their knowledge of σ. This is done indirectly by proving knowledge of α_A. We thereby exclude the attempt to cast an illegitimate vote by undetectably copying and re-randomizing σ-ciphertexts from the PB.[6] The other proof shows that $c \in \mathcal{C}$. Since each authorized vote on the voting board will be decrypted during the tallying phase, requiring the second proof prevents coercers from forcing voters to select $c \notin \mathcal{C}$ according to some prescribed pattern, thus obtaining a receipt (*Italian attack*) [Di07] or from using the talliers as a decryption oracle to obtain σ-credentials for subsequent votes.

[6] Due to this measure, votes cannot be cast by stealing the credentials of other voters, given a trustworthy majority of registrars (a majority could still compute σ and i) and talliers (a majority could compute the private decryption key and decrypt $sigma$-credentials from list $\mathcal{UNL} - (E_\varepsilon(\sigma), i)$

Apart from casting the ι-credential, this step is exactly the same as in JCJ. Although the coercer has no means of deciding to whom, among the uncontrolled voters, the ι-credentials refer to, he still gains a quantifiable advantage at coercion. Recall that the voter under coercion had to choose an arbitrary value $\tilde{\iota}$ from $\mathcal{UNL} < \iota >$ and pretend that this was his ι-credential. The reasonable coercer will therefore observe the voting board to find out whether someone has cast a vote using $\tilde{\iota}$. If this is the case, the coercer could conclude that $\tilde{\iota}$ is in fact an ι-credential that belongs to another voter and that the voter under coercion has revealed a false credential.[7] The probability that a voter is unfortunate enough to choose $\tilde{\iota}$ is less than $\frac{1}{\beta}$. The further exhibition of our protocol shows that the coercer doesn't gain any additional useful information for distinguishing the behaviour of the voter under coercion. This will lead to the conclusion that our scheme is indeed δ-coercion resistant, when $\delta = \frac{1}{\beta}$. [8]

Tallying. At the beginning of the tallying step, the voting board contains tuples of votes (A, B, ι) that might have been cast with wrong proofs, that were cast with the same credential as other votes (we call these votes *duplicates*), or that hold A- or ι-components that do not correspond with a valid credential (σ, ι) from $\mathcal{UNL} < E_\varepsilon(\sigma), \iota >$. Prior to decryption and counting, these invalid votes need to be excluded.

First, votes with wrong proofs as well as votes with ι-credentials that are not contained as the second component of an element enlisted by $\mathcal{UNL} < E_\varepsilon(\sigma), \iota >$ are marked and excluded from further processing. In order to efficiently remove duplicates, the talliers only consider votes not cast with a distinct ι-credential and apply $\mathrm{M} - \mathrm{PET}$ on the A-components of votes cast with the same ι-component.[9] At this stage a last-vote-counts or a first-vote-counts policy is enforced. Note that the steps described so far could also be performed each time a vote is posted, i.e., prior to the tallying stage.

To authorize votes, the ι-credentials are used to link the A- and B-components of the votes with the encrypted σ-credentials from $\mathcal{UNL} < E_\varepsilon(\sigma), \iota >$ to form tuples $(E_\varepsilon(\sigma), A, B)$. These tuples are passed to a mix-net. We call the output $\mathcal{UNL} < E_\varepsilon(\sigma), A, B >$, since its elements are unlinkable to both $\mathcal{UNL} < E_\varepsilon(\sigma), \iota >$ and the voter roll and the votes on the voting board. For each element, the talliers apply PET to the first two components. If the algorithm comes back as *true*, A is an encryption of a valid σ-credential. In that case, the corresponding ciphertext B is decrypted and counted in the tally, otherwise the vote is excluded from further processing. Note that since votes are being assessed for the validity of σ-credentials encrypted by the A-component, we should not apply $\mathrm{M} - \mathrm{PET}$ at this stage as such an approach would allow the coercer to

[7] Note, that this conclusion can only be drawn in the strict model proposed by JCJ, where it is assumed that exactly one voter is under coercion and that invalid credentials are only used to the degree of achieving adversarial uncertainty regarding Γ. If we now allow the coercer to believe that the vote cast with $\tilde{\iota}$ as the ι-credential is a fake vote (one with an invalid σ-credential), coercion will become even more difficult. However, we adhere to the strict model proposed in the JCJ paper.

[8] The precise value of δ is $\frac{N_+ - 1}{\beta N_\perp - 1}$. Firstly, this is always smaller than $\frac{1}{\beta}$ and secondly, the difference is very small and irrelevant for a reasonable N_+. We thus justify the facilitation of saying $\delta = \frac{1}{\beta}$.

[9] We hereby adhere to the approach proposed by Smith and Weber. However unlike Smith / Weber, we apply $\mathrm{M} - \mathrm{PET}$ only when removing duplicates, not when authorizing votes as proposed by them. Since we do not check the validity of the values encrypted by A at the current stage, and since the coercer does not know the discrete logarithm of any valid σ-credential in the base of any other, the coercer learns nothing useful for his attack.

check the validity of $\hat{\sigma}$ by the means of another vote cast by him with an A-component encrypting, e.g., $\hat{\sigma}^2$, or in other words, a value the logarithm of which is known in base $\hat{\sigma}$. The basic protocol is illustrated in figure 1.

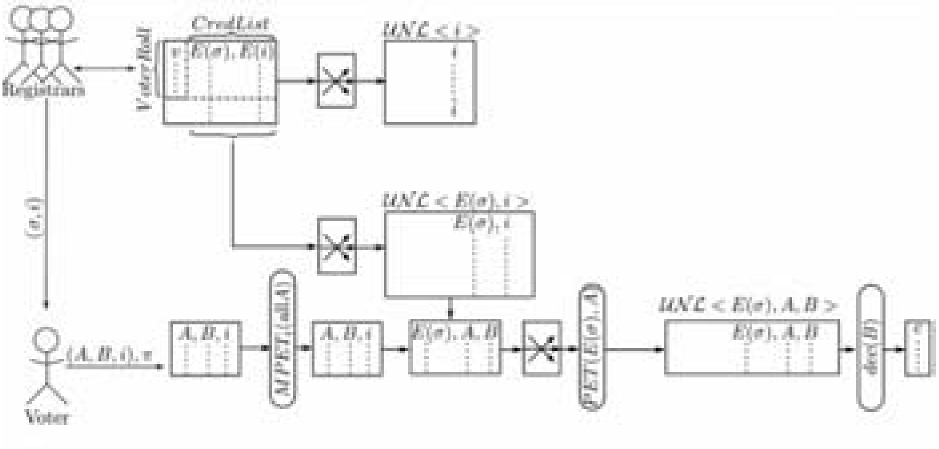

Fig. 1: Basic protocol

Credential Retention. As implied above, our scheme allows voters to re-use the same credential (σ, i) at numerous voting events. We therefore need to provide a mechanism that disallows voters to cast votes after losing eligibility, for instance when they leave the voting district. Removing their credential from the *credential pool* at post-registration is clearly not an option, since the coercer could verify the validity of the previously received i-credential by observing whether the value still appears on $\mathcal{UNC} < E_\varepsilon(\sigma), i >$ after the post-registration step of the following election. The protocol therefore defines credential retention by having the registrars compute a new σ-credential and replace $(E_\varepsilon(\sigma, \alpha_\sigma))$ in the *credential pool* with an encryption of this new value. However, the encryption of the i-credential remains the same. Finally, the voter's ID on the *voter roll* is marked as non-eligible. The new credential in the *credential pool* is marked and may not be assigned to new voters, since the coercer would know the true value of the i-credential, in case it previously belonged to a voter controlled by him. Clearly, voters who have moved will not be able to use their retained credential for voting since such votes would be discarded upon *vote authorization*. Just as all unassigned credentials in the *credential pool*, the new credential can only be used for voting unnoticed in the event of colluding registrars or talliers (a case to be ruled out in the full protocol).

Now we observe whether credential retention gives the adversary an advantage at judging if the voter, who previously lost eligibility, lied to him. We consider two cases: 1) where the voter has submitted to coercion and 2) where the voter has applied the defense strategy. In the first case, the coercer would expect the distribution of Γ, i.e., votes not to be counted, to remain the same and the number of counted votes to decrease by one. In the second case, the coercer would also expect Γ to decrease by one. This is exactly the distinguishing factor we need to assume irrelevant by means of *adversarial uncertainty* when proving the coercion-resistance of the JCJ-protocol, i.e., independent of credential retention.

3.4 Full Protocol and Improved Verifiability

Evidently, the basic protocol complies with the definition of *verifiability* in the JCJ paper: it allows one to detect the exclusion of legitimate votes, changes to legitimate votes, and the inclusion of multiple votes cast with the same credential. Notably the definition already captures the commonly quoted requirement imposed on verifiable systems, i.e., that voters need to be able to verify that their vote has indeed been cast as intended, recorded as cast, and tallied as recorded. Regarding verifiability, our basic scheme is no less powerful than the well-known coercion-resistant scheme by Araújo et al. [ABR10, AFT07, Ar08]. However, the JCJ paper mentions that it may be desirable for any election observer to verify, that credentials have only been assigned to voters whose names are on a published roll. The JCJ-protocol does indeed provide this kind of verifiability. However our basic protocol only does so when assuming trustworthy majorities among registrars and talliers. In order to ensure that one can detect the event where registrars or talliers collude to cast votes with a credential enlisted by the *credential pool* but not by the *voter roll*, we propose an enhancement to the tallying step.

In the tallying step prior to decryption, the *voter roll* is passed to a mix-net which outputs the list $\mathcal{UNL} < E_\varepsilon(\sigma) >$. The coercer cannot link the entries of this list to the entries of the voter roll. After votes from $\mathcal{UNL} < E_\varepsilon(\sigma), A, B >$ with A-components that encrypt an invalid σ-credential have been excluded from further processing (at vote authorization as described above), the talliers apply $\mathrm{M-PET}$ on all A-components of $\mathcal{UNL} < E_\varepsilon(\sigma), A, B >$ and all entries in $\mathcal{UNL} < E_\varepsilon(\sigma) >$. If no collision is detected for any of the entries of the $\mathcal{UNL} < E_\varepsilon(\sigma) >$ for an A-component of $\mathcal{UNL} < E_\varepsilon(\sigma), A, B >$, the corresponding vote has obviously been cast with a credential that corresponds to an entry in the credential pool that has not been assigned to any voter. These votes are excluded from further processing, i.e., their B-components are not decrypted. The full protocol is illustrated in figure 2. Note, that since all input values to $M - PET$ are encryptions of valid σ-credentials, no discrete logarithm of any value in the base of any other is known. Therefore the coercer does not have any advantage, and it is justified to apply $M - PET$.

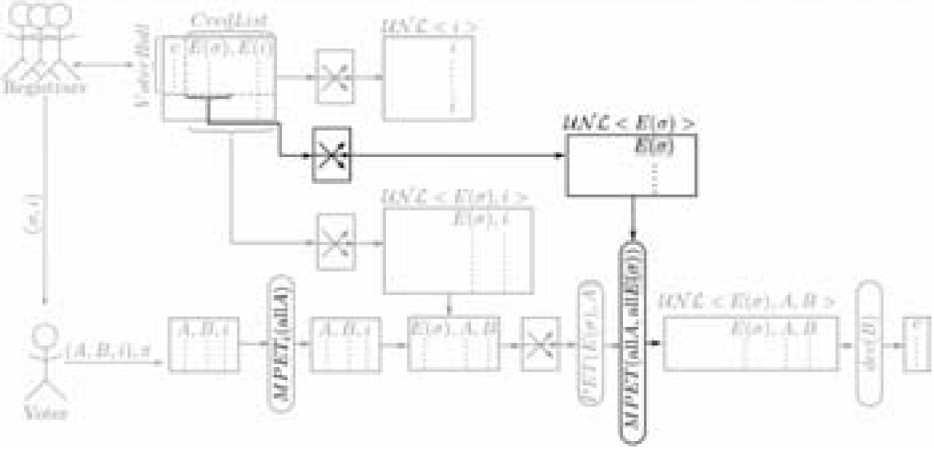

Fig. 2: Enhancement to the basic protocol to achieve full protocol

4 Efficiency

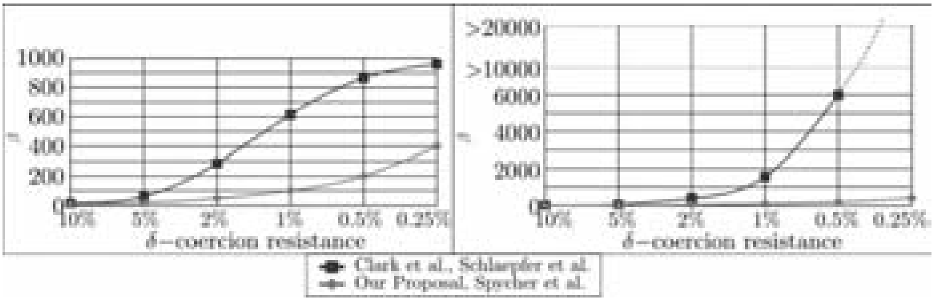

Fig. 3: The two drawings show the parameter β dependent on the degree of coercion-resistance δ. The diagram on the left shows the case for 1000 voters and 1000 votes on the voting board, the one on the right 100000 voters and 100000 votes on the voting board.

We now present the efficiency properties of our protocol through comparison with the schemes known from the literature. In the schemes by Clark et al. [CH11] and Schläpfer et al. [SHK11], voters associate their vote with non-anonymous information on the PB that refers to themselves. In order to mislead coercers, they randomly choose a set of other voters, who they can associate their vote with, thus forming an anonymity set of size β.[10] In the case of Clark et al., the *computation time on the voter's platform* scales in the parameter β. Particularly the number of modular exponentiations is $4 \cdot \beta + 10$, assuming a set $\mathbb{C}$ of two candidates to choose from. However, the tallying stage remains unaffected by the parameter and efficient, i.e., it is equally efficient as our basic protocol. The tallying time of our full protocol takes slightly longer, depending on the size of the mix-net but not more than twice as long. In Schläpfer et al. the *tallying time* scales in β, i.e., a mix-net during the tallying stage will need to perform $48 \cdot \beta \cdot N$ modular exponentiations, where N denotes the number of cast votes when assuming four mix-nodes.

The scheme by Spycher et al. [SKH11] does not rely on anonymity sets. Instead the registrar, who enjoys the voter's trust even after registration, assigns the voter an average number of β votes, under uniform distribution, cast with a false credential. Clearly this will also scale the time of tallying. $156 \cdot \beta \cdot n + 156 \cdot N$ is the number of modular exponentiation due to the most expensive steps, where n denotes the number of voters.

[10] In both cases coercion-resistance of degree $\delta = 0$ can be achieved by selecting $\beta = n$, where n is the number of voters. Moreover, it is sufficient for coerced voters to hide their votes in the anonymity set of size n, assuming adversarial uncertainty regarding the number of such votes. However this is a strong requirement, given large n.

Figure 3 shows the choice of β depending on the desired degree of coercion-resistance for the schemes with a corresponding parameter.[11] The scheme by Araújo et al. [ABR10] is by nature efficient at all stages and coercion-resistant with $\delta = 0$. However, as shown in Section 3.4 , it gives no means to verify whether authorities have created illegitimate credentials and cast extra votes.

We conclude that our protocol is efficient at both vote-casting and tallying. It does scale over β, but only during the non-critical pre-registration and post-registration steps. We therefore omit exact quantification. Furthermore, our protocol allows high levels of coercion-resistance, even under relatively small parameters. Since the pre-registration step may be conducted independent of the voting procedures, it will not have a negative impact on the elections. Also, the post-registration step can begin right after last voter has registered and only needs to end prior to tallying. The phase when citizens cast their votes should give enough time for completion.

5 Conclusion

It is true that the verifiable JCJ protocol offers coercion resistance but only under conditions that dot now allow such a protocol to be implemented for large-scale elections. Other proposed solutions either compromise verifiability or require a trade-off between coercion-resistance and efficiency during the critical phases of tallying vote-casting. Our proposal also requires more computation than conservative verifiable schemes; however, we have shown that when compared with other schemes, the factor that scales the computation time is small for relatively high degrees of coercion-resistance. Moreover, the expensive computations specific to coercion-resistance can be performed while the polls are open, i.e., while nobody is waiting.

Bibliography

[AFT07] R. Araújo and S. Foulle and J. Traoré. A Practical and Secure Coercion-Resistant Scheme for Remote Elections. In D. Chaum and M. Kutylowski and R. L. Rivest and P. Y. A. Ryan, editors, FEE'07, Frontiers of Electronic Voting, pages 330--342, Schloss Dagstuhl, Germany, 2007.

[ABR10] R. Araújo and N. Ben Rajeb, R. Robbana and J. Traoré and S. Youssfi. Towards Practical and Secure Coercion-Resistant Electronic Elections. In S. H. Heng and R. N. Wright and B. M. Goi, editors, CANS'10, 9th International Conference on Cryptology And Network Security in LNCS 6467, pages 278--297, Kuala Lumpur, Malaysia, 2010.

[11] In Section 3.3 we have shown that the coercion-resistance of our scheme follows $\delta = \frac{1}{\beta}$. It is easy to see that the same relation applies to the scheme by Spycher et al. as well. In the case of the protocols that rely on anonymity sets we have followed the definition from [KTV10]. To obtain δ, we need to compute $\sum_{r \in R} Prob(r|(\sigma,i)) - Prob(r|(\hat{\sigma},i))]$, where the condition in the first term signifies submission to coercion, the condition in the second one signifies applying the defense strategy. R denotes the set of results (i.e. the number of votes assigned to the voter under coercion) that the coercer would accept. Note, that inherent to assuming a reasonable coercer, the difference within the sum is inherently never negative. $\mathrm{Prob}(r|(\sigma,i)$ we compute as $F_1(r)$, where F_1 is the distribution function of a binomial distribution with N trials and a success probability of $\frac{\beta-1}{n-1}$, where N denotes the number of cast votes and n the number of voters. $\mathrm{Prob}(r|(\hat{\sigma},i))]$ we compute as $F_2(r-1)$, where F_2 again is the distribution function of a binomial distribution, this time with $N-1$ trials.

[Ar08] R. Araujo. On Remote and Voter-Verifiable Voting. PhD thesis, Department of Computer Science, Darmstadt University of Technology, Darmstadt, Germany, 2008.

[CH11] J. Clark and U. Hengartner. Selections: Internet Voting with Over-the-Shoulder Coercion-Resistance. FC'11, 15th International Conference on Financial Cryptography, St. Lucia, 2011.

[CCM08] M. R. Clarkson and S. Chong and A. C. Myers. Civitas: Toward a Secure Voting System. SP'08, 29th IEEE Symposium on Security and Privacy, pages 354--368, Oakland, USA, 2008.

[Di07] R. Di Cosmo. On Privacy and Anonymity in Electronic and Non Electronic Voting: the Ballot-as-Signature Attack. Hyper Articles en Ligne, hal-00142440(2), 2007.

[GJK99] R. Gennaro and S. Jarecki and H. Krawczyk and T. Rabin. Secure Distributed Key Generation for Discrete-Log Based Cryptosystems. In J. Stern, editors, EUROCRYPT'99, International Conference on the Theory and Application of Cryptographic Techniques in LNCS 1592, pages 295--310, Prague, Czech Republic, 1999.

[HS00] M. Hirt and K. Sako. Efficient Receipt-Free Voting based on Homomorphic Encryption. In G. Goos and J. Hartmanis and J. van Leeuwen, editors, EUROCRYPT'00, International Conference on the Theory and Applications of Cryptographic Techniques in LNCS 1807, pages 539--556, Bruges, Belgium, 2000.

[JJ00] M. Jakobsson and A. Juels. Mix and Match: Secure Function Evaluation via Ciphertexts. In T. Okamoto, editors, ASIACRYPT'00, 6th International Conference on the Theory and Application of Cryptographic Techniques in LNCS 1976, pages 162--177, Kyoto, Japan, 2000.

[JCJ05] A. Juels and D. Catalano and M. Jakobsson. Coercion-Resistant Electronic Elections. In V. Atluri and S. De Capitani di Vimercati and R. Dingledine, editors, WPES'05, 4th ACM Workshop on Privacy in the Electronic Society, pages 61--70, Alexandria, USA, 2005.

[Li11] Lipmaa, Helger. On the CCA1-security of Elgamal and Damgard's Elgamal. Proceedings of the 6th international conference on Information security and cryptology in Inscrypt'10, pages 18--35, Berlin, Heidelberg, 2011. Springer-Verlag.

[Pe91] T. P. Pedersen. A Threshold Cryptosystem without a Trusted Party. In D. W. Davies, editors, EUROCRYPT'91, Workshop on the Theory and Application of Cryptographic Techniques in LNCS 547, pages 522--526, Brigthon, U.K., 1991.

[KTV10] R. Küsters and T. Truderung and A. Vogt. A Game-Based Definition of Coercion-Resistance and its Applications. Proceedings of the 23nd IEEE Computer Security Foundations Symposium (CSF 2010), pages 122-136, 2010. IEEE Computer Society.

[SHK11] Michael Schläpfer and Rolf Haenni and Reto Koenig and Oliver Spycher. Efficient Vote Authorization in Coercion-Resistant Internet Voting. 3rd International Conference on E-Voting and Identity (VoteID 2011), 2011. Springer-Verlag.

[SKH11] O. Spycher and R. Koenig and R. Haenni and M. Schläpfer. A New Approach Towards Coercion-Resistant Remote E-Voting in Linear Time. FC'11, 15th International Conference on Financial Cryptography, St. Lucia, 2011.

[We08] S. Weber. Coercion-Resistant Cryptographic Voting: Implementing Free and Secret Electronic Elections. VDM Verlag, Saarbrücken, Germany, 2008.