Using Selene to Verify Your Vote in JCJ

Vincenzo Iovino, Alfredo Rial, Peter B. Rønne$^{(\boxtimes)}$, and Peter Y.A. Ryan

University of Luxembourg, Esch-sur-Alzette, Luxembourg
vinciovino@gmail.com, {alfredo.rial,peter.roenne,peter.ryan}@uni.lu

Abstract. We show how to combine the individual verification mechanism of Selene with the coercion-resistant e-voting scheme from Juels, Catalano and Jakobsson (JCJ). This results in an e-voting scheme which allows the voter to check directly that her vote is counted as intended, but still allows her to mitigate coercion.

We also construct variants of the protocol which provide everlasting privacy or better verifiability. Further, both improvements of JCJ and Selene are discussed.

1 Introduction

Remote e-voting gives voters the opportunity to conveniently vote from home, work or even abroad. However, it also presents cryptographers with the difficult task of integrating both verifiability and privacy properties in a secure, efficient and usable e-voting protocol. One of the hardest problems of leaving the reassuring frame of a voting booth is to protect voters against coercion attempts. Juels, Catalano and Jakobsson (JCJ) [JCJ05] found a way to provide coercion-resistance across multiple elections, assuming only a single coercion-free registration. The registration provides the voters with credentials which they use for voting. Coerced voters can provide the coercer with a fake credential, and a vote cast using this will not be counted. The system was later implemented as Civitas [CCM08].

The JCJ-mechanism might be worrisome to the normal user. Was the credential entered correctly? Did someone else manage to override my vote? In the end, it would be reassuring for the voters to be able to directly check that their votes were counted correctly. However, providing voters with such a service endangers the receipt-freeness and coercion-resistance if not done carefully. Fortunately, Selene [RRI16] provides us with a mechanism for individual tallied-as-intended verifiability while being able to mitigate the coercion threat. This is done by giving each vote a unique tracking number, but first revealing this to the voter after the tally has been published. Unfortunately, Selene was developed for Helios style protocols, but in this paper we will show that the construction can also be applied to the coercion-resistant vote casting system from JCJ/Civitas. Indeed we will consider different variants of JCJ and show how Selene can be added to JCJ even in the case when we want to provide everlasting privacy via pseudonyms, or when we offer better verifiability properties. We will also see

M. Brenner et al. (Eds.): FC 2017 Workshops 2017, LNCS 10323, pp. 385–403, 2017.
https://doi.org/10.1007/978-3-319-70278-0_24

how to address the secure platform problem with the extra verifiability gained from Selene. Along the way we will discuss some problems and solutions of the JCJ construction with cross-election and dynamic coercion. Further, we will give a more efficient construction of the zero-knowledge proofs needed in Selene.

1.1 Related Work

Since the seminal paper defining coercion-resistance [JCJ05], there have been numerous paper analyzing the JCJ protocol and providing alternatives, see e.g. [NFVK13a] and references therein.

Selene [RRI16] is based on the idea of having trackers for the votes, an idea already suggested in Schneier's book [Sch94], which later independently also appeared in a scheme used for ANR (Agence National de la Recherche) funding committee meetings. Recently, sElect [KMST16] uses trackers to achieve good accountability. However, in all of these cases the tracker directly represents a receipt, whereas Selene mends this by delaying when the voter can obtain the tracker.

The idea of everlasting privacy goes back to Moran and Naor [MN06] and have been studied in several works, see e.g. [CPP13] for how to make perfectly private audit trails in general election schemes, or [ACKR13] for how to do automated verification of everlasting protocols. Here we focus on pseudonymity rather than anonymity. However, if we follow JCJ closely, this is the best we can do since the credentials themselves will be like pseudonyms to a future adversary.

The secure platform problem is one of the main problems in e-voting. One solution is to use out-of-band channels and code-voting, see [Cha01,RT09]. In e.g. Helios [Adi08] Benaloh challenges [Ben06] should help to detect malware, but are unfortunately not often used [KOKV11]. Relying on hardware tokens is yet another possibility, see [HK14,GRCC15], but is not always unproblematic, see [KR16].

2 Building Blocks

Our construction uses the following building blocks: a non-interactive zero-knowledge proof system (NIZK) [BFM88] in the random oracle model [BR93], the ElGamal public key encryption scheme [Gam85], threshold encryption with a plaintext equivalence test [JJ00], a verifiable re-encryption mixnet [SK95], the Pedersen commitment scheme [Ped91], a web bulletin board [HL09], untappable channels [HH07] and anonymous channels [Fre00].

3 System Model and Setup

We first describe the parties involved in an e-voting scheme.

Voters. The voters V_i $(i = 1, \ldots, n)$ register for voting, cast ballots, obtain trackers and verify the voting results.

Tally Tellers. The Tally Tellers T_j tally the cast ballots and publish the results.

Registration Tellers. The Registration Tellers RT_k register voters.

Tracker Tellers. The Tracker Tellers TT_l process trackers. They could be the same parties as the Tally Tellers or the Registration Tellers, but they are kept separate here due to the different trust assumptions.

Our e-voting scheme consists of the following phases.

Setup. In the setup phase, the parties generate secret and public keys. Each voter creates a designated verifier key. The Tally Tellers generate a public key pk_T for a threshold encryption scheme.

Registration. In the registration phase, a voter V_i and the registration tellers run a protocol. The designated verifier key dvk_i of V_i and pk_T are used as inputs. As a result of the protocol, the voter obtains a credential C_i. Additionally, the voter identifier V_i, the key dvk_i and an encryption of C_i under pk_T are published on the web bulletin board (BB).

Tracker Preparation. In this phase, the Tracker Tellers and the voters run a protocol. A set of trackers $\{n_i\}_{i=1,\ldots,n}$, the designated verifier keys of the voters and pk_T are used as inputs. As a result of the protocol, each voter obtains a Pedersen commitment to its tracker. Additionally, an encryption under pk_T of the tracker associated with a voter V_i is appended to the row for voter V_i on the BB. In this protocol, the association between trackers and voters is not revealed to any party.

Vote Casting. In this phase, a voter V_i computes a ballot and publishes it on the BB. In our construction, the ballot contains an encryption under pk_T of the credential C_i and of the vote vote_i.

Tallying. In this phase, the Tally Tellers take as input the ballots published on the BB and run a protocol to output pairs (vote_a, n_a), which associate each valid vote with the tracker of the voter that cast that vote. Those pairs are published on the BB.

Tracker Retrieval. In this phase, a voter V_i and the tracker tellers run a protocol as a result of which V_i learns the tracker n with which it became associated in the tracker preparation phase.

Once a voter learns her tracker n_a, the voter can verify on the BB that the pair (vote_a, n_a) is correct.

Setup. Let G be a cyclic group of prime order q and g be a generator of G. In the setup phase, each voter creates designated verifier key $\mathsf{dvk}_i = g^{x_i}$. The designated verifier keys are used to provide deniability in the registration phase in JCJ and the implementation Civitas. We assume that the designated verifier key system is well setup, which includes that the voters have proven that they know their secret key. Additionally, we use the same designated verifier key in the Selene construction as the public key for the ElGamal encryption scheme. JCJ also suggest an alternative registration with an erasure function. In that case we need a PKI as in Selene where dvk_i is the voter's public key.

The Tally Tellers run the distributed key generation algorithm of the threshold encryption scheme to generate a public key pk_T and obtain each a private share of the secret key.

4 Description of the E-Voting Protocol

In this section, we describe the protocol combining JCJ and Selene in detail.

4.1 Registration

The registration is quite similar to JCJ/Civitas. Each voter has a designated verifier key dvk_i. The voter must prove (interactively in ZK) that she knows the secret key corresponding to dvk_i over an untappable channel during registration. This is to prevent a coercer from making a voter register a designated verifier key for which the secret key is unknown by the voter. For each eligible voter V_i, each Registration Teller RT_j randomly picks $C_{ij} \leftarrow_{\$} G$ and publishes $\{C_{ij}\}_{\mathsf{pk}_T}$ on BB in a row marked for voter V_i. As discussed in Sect. 7.1, we could instead use pseudonyms PV_i if everlasting privacy is desired.

For each voter, the encryptions are multiplied together to homomorphically obtain a single credential C_i. On BB, we now have the following row for each voter

$$V_i, \mathsf{dvk}_i, \{C_{ij}\}_{\mathsf{pk}_T}, \prod_j \{C_{ij}\}_{\mathsf{pk}_T} = \{C_i\}_{\mathsf{pk}_T}$$

Here we deviate from JCJ/Civitas by also including the public key dvk_i in the row. This key will be the public key used by the voter in Selene.

The voter now receives the credential shares and designated verifier proofs from the Registration Tellers

$$RT_i \rightarrow V_i : C_{ij}, \pi_{ij}$$

where π_{ij} is a designated proof to the key dvk_i proving that $\{C_{ij}\}_{\mathsf{pk}_T}$, appearing on BB, is an encryption of C_{ij}. The voter can now calculate C_i and check the proofs.

If the voter is coerced, she chooses at random an alternative value C_i' in G and shows this to the coercer. The proofs can be faked with her designated verifier key. It is here of course essential that the voter knows the secret key, but a coerced voter can even reveal this secret key to the coercer, as long as the coercer does not cooperate with the registration tellers. It is important that the coercer is not present by the reception of all C_{ij}'s.

The credential can be reused for several elections, and could, in principle, be obtained in booth by the registration authorities.

4.2 Tracker Preparation

Whereas the previous part was very similar to JCJ/Civitas, we now add the main ingredient of Selene, namely, the personal voting trackers that each voter can use to check her tallied vote.

The trackers $\{n_i\}_{i=1,\ldots,n}$ should be a negligible set of $\mathbb{Z}_q$ (i.e. the chance of a random element in $\mathbb{Z}_q$ being a tracker is negligible).

The Tracker Tellers first publish

$$n_i, \{g^{n_i}\}_{\mathsf{pk}_T}$$

on BB, where the encryption is with trivial randomness. The trackers are sent through a re-encryption mix and one anonymised tracker is added to each of the voters' rows to obtain

$$V_i, \mathsf{dvk}_i,\ \{C_i\}_{\mathsf{pk}_T}, \{g^{n_{\pi(i)}}\}_{\mathsf{pk}_T}$$

where π is the permutation used for mixing. In the following we will suppress π for easier notation. Note that, whereas credentials can be used for several elections, this tracker mixing needs to be renewed for each election.

Further each Tracker Teller TT_j randomly chooses $r_{ij} \leftarrow_{\$} \mathbb{Z}_q$ for each voter and publishes

$$\{\mathsf{dvk}_i^{r_{ij}}\}_{\mathsf{pk}_T}, \{g^{r_{ij}}\}_{\mathsf{pk}_T}, \Pi_{ij}$$

where Π_{ij} is a non-interactive zero-knowledge proof that this is done correctly. The proof is presented in the Selene protocol, see [RRI15], Appendix A. As in Selene, the terms from each Teller are now homomorphically combined with the encryption of the tracker, and we obtain a trapdoor commitment to the tracker

$$\{g^{n_i}\}_{\mathsf{pk}_T} \prod_j \{\mathsf{dvk}_i^{r_{ij}}\}_{\mathsf{pk}_T} = \{g^{n_i}\,\mathsf{dvk}_i^{r_i}\}_{\mathsf{pk}_T}$$

with $r_i = \sum_j r_{ij}$. This is appended to each voter's row. Finally, the Tally Tellers decrypt the trapdoor commitment to the tracker, $g^{n_i}\,\mathsf{dvk}_i^{r_i}$, for each voter.

4.3 Vote Casting

Vote casting is done like in JCJ. Here we follow Civitas. If voter V_i wants to vote "vote$_i$", she anonymously sends to BB

$$(\{C_i\}_{\mathsf{pk}_T}, \{\text{vote}_i\}_{\mathsf{pk}_T}, \pi)$$

where π is a zero-knowledge proof that the vote is well-formed together with a proof that C_i and vote$_i$ are simultaneously known, which prevents vote copying. To cast a vote in presence of a coercer, the fake credential given to the coercer is simply used in place of the real one.

4.4 Improving the Coercion Resistance of JCJ

The JCJ protocol has a tally procedure which leaves room for certain coercion attacks. Let us first remind ourselves how the tally procedure works. It relies heavily on the Tally Tellers performing Plaintext Equivalence Tests (PETs) on the encryption of the credentials, see e.g. [CCM08].

1. Zero-knowledge proofs of the cast ballots are checked, and invalid ballots are removed.
2. Duplicates, i.e., ballots that use the same credential, are removed according to the existing vote update policy. This is done using PETs among the ciphertexts of the credentials in the cast ballots. This means that the coercer cannot mark the vote with a chosen number of duplicates.
3. The list of ciphertexts of registered credentials is anonymized using a mix-net. Further, from the list of valid ballots after duplicate removal, we likewise use a parallel mix-net to anonymize the pairs of ciphertexts of credentials and votes.
4. Unauthorized votes, i.e., ballots that do not use a registered credential, are removed by performing PETs of the credentials from the cast votes with the list of registered credentials.
5. The remaining valid votes can now be decrypted to reveal the tally.

The duplicate removal can in certain quite special situations give the coercer unwanted information and correspondingly hinders coercion-resistance, as we will now see. This was discovered, but not analyzed, in [Roe16]. The problem appears when the coercion happens dynamically or across elections. Consider an uncoerced voter who has already voted. The coercer now detects this somehow, say by overhearing this or seeing this in the browsing history of the voter.[1] The coercer can now coerce the voter just before voting ends. The coerced voter now gives the coercer a fake credential, and they can sit down and cast an, in fact, invalid vote. However, in the duplicate removal phase, it will then be evident that the credential was fake, since no duplicates are detected for the fake vote. To circumvent this, all voters should start by casting fake votes if they want to be prepared for later coercion threats, which seems pretty complicated. Note that the protocol in [KHF11] actually does something similar to prevent board flooding attacks on JCJ, but the cost is a statistical coercion-resistance.

Another case is a voter which was coerced in an earlier election and gave the coercer a fake credential. At a later election, the coercer can now cast a vote using this credential and check whether this will have duplicates in the duplicate removal phase. If this does not happen, the coercer can conclude that either the credential was fake, or that the voter did not vote in the latter election, which might be improbable. This means that the coerced voter also needs to cast votes using the fake credential even at elections after being coerced to be on the safe side.

[1] We can assume that coerced voters are careful to use only devices out of the reach of the coercer or to delete browsing history, but this is more unlikely for uncoerced voters.

Note that it does not help to do a mix before performing the duplicate elimination since the groups of ballots could still be marked by a certain number of duplicates.

If vote updating is not intended, we can sidestep the issue by simply dropping the step of duplicate removal. After anonymizing both the registered credentials and the cast ballots, PETs are performed for each registered credential against the cast ballots until the first match comes up. We then pick this as the vote for the given credential. For the set of cast votes for a given valid credential this will pick one in the set at random. The method thus reveals a minimum amount of information, but makes vote updating harder to implement. Further it also decreases verifiability as discussed below, but Selene helps here.

4.5 Tallying with Selene

Tallying with Selene requires a minor modification. First, all proofs are checked and invalid votes are discarded. Then all cast pairs

$$(\{C_a\}_{\mathsf{pk}_T}, \{\mathrm{vote}_a\}_{\mathsf{pk}_T}) \mapsto (\{C_{\pi(a)}\}_{\mathsf{pk}_T}, \{\mathrm{vote}_{\pi(a)}\}_{\mathsf{pk}_T})$$

are re-encryption mixed.

Further the pairs of registered credentials and tracking numbers

$$(\{C_i\}_{\mathsf{pk}_T}, \{g^{n_i}\}_{\mathsf{pk}_T}) \mapsto (\{C_{\pi'(i)}\}_{\mathsf{pk}_T}, \{g^{n_{\pi'(i)}}\}_{\mathsf{pk}_T})$$

from each voter's column are re-encryption mixed in parallel. From each entry in this anonymised list of credential-tracker pairs, the Tally Tellers do PETs against the credentials from the anonymised list of cast votes. The first time we get a positive match, the corresponding vote is decrypted (verifiably) together with the corresponding tracker. If wanted, one can also do more elaborate PETs (like in JCJ-Civitas), first removing all duplicate votes, possibly with some vote update policy, as explained in Sect. 4.4.

The end result (after taking the discrete log of the trackers) is the Tally Board of valid vote-tracker pairs (since the set of trackers is small and known, it is easy to go from g^n to n)

$$(\mathrm{vote}_a, n_a).$$

4.6 Tracker Retrieval

Finally, the tracker retrieval happens like in Selene. Each Tracker Teller provides each voter with their share $g^{r_{ij}}$.

$$TT_j \rightarrow V_i : g^{r_{ij}}.$$

This happens according to some random time distribution a suitable time after the tally has been published, see [RRI16] and via, for the coercer, untappable channels.

The voter (or rather her device) combines these shares to get g^{r_i}. Together with the public trapdoor commitment $g^{n_i}\,\mathsf{dvk}_i^{r_i}$, the term g^{r_i} forms an ElGamal encryption of the tracker under the key dvk_i. The voter can now decrypt and directly check that her vote appears correctly on the Tally Board.

Trackers can be faked in the case of coercion, just like in Selene. That is, the voter finds the wanted fake tracker, n^*, on BB for the coercer's choice of vote and calculates

$$\left(g^{-n^*} g^{n_i}\,\mathsf{dvk}_i^{r_i}\right)^{x_i^{-1}}$$

as the fake term to give to the coercer instead of g^{r_i}. Here x_i is the secret key of $\mathsf{dvk}_i = g^{x_i}$.

A potential attack would raise if an adversary, possibly colluding with all the Tracker Tellers, could make a voter get a fake g^{r_i} term that the voter decrypts to a valid tracker different from the true tracker of the voter with non-negligible probability. In [RRI15] it is proven that this is hard under a standard computational assumption.

5 More Efficient Zero-Knowledge Proofs in Selene

In the tracker preparation phase, the Tracker Tellers publish

$$\{\mathsf{dvk}_i^{r_{ij}}\}_{\mathsf{pk}_T}, \{g^{r_{ij}}\}_{\mathsf{pk}_T}, \Pi_{ij}$$

where the zero-knowledge proof was of the correctness of this construction, i.e. that the two generators are raised to the same known power. However, the term $\{g^{r_{ij}}\}_{\mathsf{pk}_T}$ is not really needed. In principle, it could be used for accountability if the Tracker Teller tries to send a wrong $g^{r_{ij}}$ to the voter. However, for deniability, the Tracker Teller sends this term without any proof to the voter. This means that there is no proof that the Teller sent a wrong message to the voter. Thus we suggest to only publish

$$\{\mathsf{dvk}_i^{r_{ij}}\}_{\mathsf{pk}_T}, \Pi'_{ij}$$

where Π'_{ij} is a shorter zero-knowledge proof, showing that the ciphertext indeed encrypts the key dvk_i to a known power. In a long version of this note we present this proof in details; it consists of 8 group elements in some group of prime order p and of 6 elements of $\mathbb{Z}_p^\star$. We also prove in the long version that the adversary, also in this case, even when colluding with all Tellers, only has a negligible chance of constructing a fake term $g^{r_{ij}}$ that makes the voter decrypt to a valid tracker different from her real tracker.

6 Security Assumptions and Arguments for Security

In this section we will briefly mention the trust assumptions for the voting authorities and give brief explanations of why the different security properties hold.

6.1 Trust Assumptions for the Tellers

- The Registration Tellers are trusted individually for coercion-resistance and collectively for verifiability. For everlasting privacy via pseudonyms (see Sect. 7.1) they are individually trusted for everlasting privacy.
- The Tally Tellers are trusted collectively for privacy (and hence coercion-resistance) and verifiability. A threshold version follows directly. We will here assume that the verifiable reencryption mixes done in the protocol are performed by the Tally Tellers, and that these are private if at least one Teller is honest.
- The Tracker Tellers are trusted collectively for privacy. They are trusted individually for coercion-resistance since the voter needs to know which $g^{r_{ij}}$ to fake for the coercer (like for the Registration Tellers).

6.2 Verifiability

For verifiability, we assume that the voters keep their private designated verifier keys secret. An adversary colluding with all the Registration Tellers can however still obtain the credential of a voter, and cast votes on her behalf, violating at least eligibility verifiability. The same can happen if the adversary and all the Tally Tellers collude, see also Sect. 7.2 below how to mitigate this risk.

However, if such grand collusion do not happen, the only ballots on BB with a given voters correct credential are with overwhelming probability cast by the voter herself. That the correct vote is now chosen in the tally is secured by checking the zero-knowledge proofs of the verifiable PETs and verifying the correctness of the mixes. Finally, the actual decryption of the vote can also be verified.

The correctness of the individual verifiability of the Selene trackers, is very similar to the original Selene construction. The verification of the first mix of the trackers ensures that each voter gets a unique tracker, from the set of trackers. The pairwise mix of registered credentials and trackers, together with verification of the PETs ensure that this tracker is assigned to the voter's cast vote. Again, the correct decryption of the trackers can be verified. That the voters receive the correct trackers with overwhelming probability is discussed above.

6.3 Vote Privacy

If the Tally Tellers or Tracker Tellers collude they can easily break privacy. Otherwise privacy of the mixes and encryptions will ensure privacy. In general, ballot independence is ensured by the construction (at least if we do not do the duplicate weeding) if we check the proofs of the PETs. This also means that even if the Registration Tellers collude and can cast valid votes on behalf of voters, this does not violate privacy.

6.4 Coercion-Resistance and Coercion-Mitigation

Coercion-resistance and, related, receipt-freeness is a harder problem. The point is that even in the ideal version of the scheme, the voters will know exactly which vote is theirs in the final tally by checking their unique tracker. This is intended and gives the voter a reassurance of the correctness of the vote. However, each voter knowing their unique tracker does constitute a piece of information, not obtainable in standard voting schemes, and which is not foreseen in standard definitions of coercion-resistance and receipt-freeness.

Coerced voters however still have good options to *mitigate* coercion. They have algorithms to both fake their credential and the term to obtain their tracker number. The difference to standard coercion-resistance crystallizes when the voter shows a fake tracking number to the coercer, and it turns out to be the coercer's own tracker. This was analyzed in Selene [RRI15] where also several alternative versions without this drawback were discussed, but at the cost of a less clear Tally Board.

Another problem comes from a slight lack of coercion-resistance in the JCJ construction itself, which is then magnified by the addition of individual verifiability. JCJ, and any scheme which has voting authorised by a token which can be faked and provided to the coercer, is not strictly coercion-resistant at least if used across multiple elections. The point is that the coercer can cast votes using the token obtained from the voter for special candidates which are expected to get a low number of votes. In the extreme case where this candidate does not get any votes in the election, the coercer knows he was provided with a fake token. In less extreme cases, the coercer only gains statistical information of the validity of the token. Used across elections the statistical certainty can be improved. This is not so different from the coercer actually directly demanding the voter to vote for the corresponding candidate. However, the point is that the coercer can choose to follow this strategy completely without the voter's knowledge and thus without the voter having the choice to remain undetected by following the coercer's wish. The individual verifiability here worsens the situation[2] since the coercer can demand to know which candidate he chose, which the voter should know via the individual verifiability mechanism (here via the tracking number). Of course, the voter cannot reply with certainty due to the covert strategy of the coercer. However we will see in Subsect. 7.4 that we can change the tracker retrieval mechanism to allow the voter to answer the coercer's demand with a valid tracker thus successfully defending against this attack (unless the chosen candidate did not get any votes, in which case no defense would ever be possible for the chosen result function).

7 Extensions and Alternative Protocols

7.1 Everlasting Privacy via Pseudonyms

Privacy is easy to break for a future adversary who is able to break the employed encryption, e.g. because the DDH assumption happens to be broken or simply

[2] Thanks to Véronique Cortier for pointing this out.

by the expected increase in computational power over time. In general, we think about the future adversary as having unlimited computational power, but only being active after the election using the data from BB.

A quick and dirty way to obtain everlasting privacy is to use pseudonyms (see [LHK16] for a more advanced approach to everlasting privacy and coercion-resistance, however with an efficiency drawback). I.e., instead of labelling the rows on the bulletin board with the voter IDs, we use pseudonyms. We assume that only the Registration Tellers and the Tracker Tellers know the relation between the pseudonyms, designated verifier keys and the actual voter IDs. Especially, this information will not be public and not available to the future adversary.

Of course, pseudonyms are not the best way to preserve privacy, especially across elections. However, they are easy to implement with not too big usability costs. In particular, the JCJ construction works with credentials, which to the future adversary are just like pseudonyms labelling the voters, even though they only appear under encryption. As we show now, we can also use the Selene mechanism in this case with some modifications.

Registration with Pseudonyms. In the registration phase, we mark the voter's row on BB with the pseudonym PV_i instead of V_i

$$PV_i, (\mathsf{dvk}_i)^{s_{ij}}, \{C_{ij}\}_{\mathsf{pk}_T}, \prod_j \{C_{ij}\}_{\mathsf{pk}_T} = \{C_i\}_{\mathsf{pk}_T}$$

Note that each Registration Teller also takes the public key of the voter dvk_i and raises it to the random power s_{ij} before publishing it. For each voter, we can now collect the terms

$$PV_i, (\mathsf{dvk}_i)^{s_i} = \prod_j (\mathsf{dvk}_i)^{s_{ij}}, \prod_j \{C_{ij}\}_{\mathsf{pk}_T} = \{C_i\}_{\mathsf{pk}_T},$$

with $s_i = \sum_j s_{ij}$. The Registration Tellers now send both the credential shares C_{ij}, the random exponents r_{ij}, the pseudonym and the designated verifier proofs to the voter

$$RT_i \rightarrow V_i : C_{ij}, s_{ij}, \pi_{ij}, PV_i$$

where π_{ij} is a designated proof to the key dvk_i proving that $\{C_{ij}\}_{\mathsf{pk}_T}$, appearing on BB, is an encryption of C_{ij}. The voter checks the proofs and the validity of the values s_{ij}. Further, the voter can now calculate C_i and s_i. For internal purposes the voter can update her key to be $(\mathsf{dvk}_i)^{s_i}$.

The reason for raising dvk_i to s_i is two-fold. The first reason is to blind the public key from the future adversary. From $(\mathsf{dvk}_i)^{s_i}$, it is information-theoretically impossible to infer dvk_i. The second reason is to prevent the following verifiability attack. Suppose that all registration tellers collude. They could then point two or more voters to the same pseudonym and credential, which would only be detected if the attacked voters unlikely compare pseudonyms.

This would only give one vote to the two voters. Note that this verifiability is outside the scope of the JCJ assumptions, assuming at least one Registration Teller is honest. However, in [Roe16], it was shown that we can do better (see also below). However, by knowing the exponentials, the registration tellers would need to know a discrete logarithm relation between the attacked voters, which is infeasible by the hardness of the discrete logarithm problem, if we assume that the PKI has been set up properly.

The remainder of the protocol can now proceed as above with dvk_i replaced by $(\mathsf{dvk}_i)^{s_i}$. The future adversary will be able to relate a vote to the pseudonym and $(\mathsf{dvk}_i)^{s_i}$, but not directly to the voter. Note that the Tracker Tellers need to know the relation between pseudonyms and voters to return the random terms $g^{r_{ij}}$ to the voters. Like the Registration Tellers they are thus also assumed not to be colluding with the future adversary. This trust is one of the reasons to distinguish them from the Tally Tellers.

7.2 Stronger Verifiability

In [Roe16], a version of JCJ-Civitas was presented which has stronger security guarantees, and only changes the registration and voting procedure slightly. The main point is that the voters know the discrete logarithm of their credential, and this can be seen as a secret key. The cast ballots containing the encrypted credential are basically anonymously signed using this secret key. This prevents verifiability attacks where either all Tally Tellers or Registration Tellers are corrupted. In that case, they know the secret credentials, and could cast valid votes on behalf of any voter. If we use the duplicate removal step, which had slight coercion-resistance problems, as discussed above, this attack could be detectable by alert voters. However, even so, it could lead to unsolvable disputes about the validity of the election, see [Roe16].

Selene can also be added to this version of JCJ just as for standard JCJ. However, we can also create a new combination of JCJ and Selene where, post-registration, the voters only have to handle a single key (actually, coerced voters, of course, also need to handle the fake keys).

The registration works as follows. For a given voter V_i, all Registration Tellers RT_j choose random values $c_{ij} \in \mathbb{Z}_q$ and publish $\{g^{c_{ij}}\}_{\mathsf{pk}_T}$ on BB. The voter gets c_{ij} from RT_j together with a designated zero-knowledge proof to dvk_i, proving the correct encryption of $g^{c_{ij}}$.

The ciphertexts of the credential shares can now be multiplied together, but are further multiplied by $\{\mathsf{dvk}_i\}_{\mathsf{pk}_T}$, which for verifiability is encrypted with trivial randomness. Since ElGamal is homomorphic, the final ciphertext is an encryption of the voter credential $C_i = g^{c_i} := g^{\sum_j c_{ij} + x_i}$. However, in this case the voter, and only the voter, knows the discrete logarithm, since the Registration Tellers do not know the secret key of dvk_i.

In case of coercion, the voter will present the coercer with a random number c'_i and corresponding group element $C'_i = g^{c'_i}$ and claim this is the real credential – just like in JCJ, but now working with the discrete logarithms instead of the group elements.

After registration, BB contains

$$V_i, \{C_i\}_{\mathsf{pk}_T}$$

and the uncoerced voter only needs to store the discrete logarithm of C_i. We do not demand now V_i to store dvk_i separately. The Tracker Tellers can mix and add $\{g^{n_i}\}_{\mathsf{pk}_T}$ to each voter as above, but the Tracker Tellers can now only work with $\{C_i\}_{\mathsf{pk}_T}$. Due to the homomorphic property of ElGamal, this is however enough. To create the trapdoor commitment, the Tracker Teller TT_j randomly chooses $r_{ij} \leftarrow_{\$} \mathbb{Z}_q$, and publishes for each voter

$$\{C_i\}_{\mathsf{pk}_T}^{r_{ij}} = \{C_i^{r_{ij}}\}_{\mathsf{pk}_T}, \mathit{\Pi}_{ij}$$

where again $\mathit{\Pi}_{ij}$ is a NIZKPoK that this is done correctly. We have here chosen the version without publishing the encryption of $g^{r_{ij}}$, however this only changes for the proof.

Observe that we need a proof that an ElGamal ciphertext is raised to some known power and this accounts to a proof of knowledge of the randomness r in a DH-tuple. A NIZKPoK for it can be obtaining by applying the Fiat-Shamir's heuristic to the Chaum-Perdersen's proofs [CP93]. The coercion-resistance of the public information follows from the DH-assumption observing what follows. Let us assume for simplicity that there is only one teller. Then, the coercer can see $C_i^r g_i^n$ along with the ciphertext raised to r but not C_i and note also that the voter does *not* know r. Thus, under the DH-assumption we can conclude that this information consists of just random group elements.

By homomorphically multiplying $\{g^{n_i}\}_{\mathsf{pk}_T}$ with all the $\{C_i^{r_{ij}}\}_{\mathsf{pk}_T}$, we get the trapdoor commitment $\{g^{n_i} C_i^{r_i}\}_{\mathsf{pk}_T}$ where the trapdoor key now is c_i. The Tally Tellers decrypt these commitments verifiably.

Vote casting follows [Roe16] and works like before. The voter casts

$$(\{C_i\}_{\mathsf{pk}_T}, \{\text{vote}_i\}_{\mathsf{pk}_T}, \pi)$$

anonymously to BB. The difference is that the zero-knowledge proof now also contains a proof of knowledge of the discrete logarithm in the encrypted credential, i.e. like an anonymous signature.

Tallying is just like before, and retrieving the trackers likewise. However, for coerced voters, faking the random term $g^{r_{ij}}$ is now different from standard Selene. The point is that, whereas in the standard case, the coerced voter will hand out the real secret key of dvk_i to the coercer, in this case the coercer will get a fake key $C'_i = g^{c'_i}$. The fake term g^{r_i} is thus calculated as

$$\left(g^{-n^*} g^{n_i} C_i^{r_i}\right)^{{c'}_i^{-1}}$$

since, when combining this with the commitment on BB, we get a ciphertext which decrypts to the wanted tracker n^* when we decrypt with the fake credential key given to the coercer. Actually, this construction is mildly better than standard Selene for coercion. The reason is that, if the coercer somehow manages to see the real term g^{r_i}, this will decrypt to the voter's correct tracker in standard Selene, but here it will decrypt to a random number, since the coercer is in the possession of a fake key. The voter can thus still claim that something must have gone wrong, or the system is corrupted, whereas in standard Selene the chance of this would be negligible. In real life, this is probably not a very usable defense for coerced voters.

Note that, if Tracker Tellers are corrupted, they can reveal relations on the credentials between voters from the decrypted commitments, since they know the random coins used in the commitments. This is however less of a problem in this version of JCJ since the discrete log of the credential is needed to break verifiability, and the Tracker Tellers are anyway trusted for coercion-resistance.

7.3 On the Secure Platform Problem

One of the main problems of e-voting is the secure platform problem. Very often this problem is ignored and the voter's computing platform is considered safe. An alternative useful approach is to use an out-of-band channel, e.g. using vote codes on paper, see e.g. Pretty Good Democracy [RT09].

Instead of resorting to out-of-band channels, one can also try to secure the device used by the voter, see e.g. [NV12] [NFVK13b] where simple smart cards are used. These are further used to improve usability for the voter. One drawback of dedicated hardware might be forced abstention attacks from local coercers, who simply seize the device from the coerced voter.

Instead, we can try to spread the risk of malware attacks to two independent devices, assuming that the adversary will not be able to control both. Further, we keep these devices general, i.e., it could be smartphones or laptops and not dedicated hardware. Keys could have backups on more devices if the voter is afraid of forced coercion. Due to the setup with two different credential/keys, the combination of JCJ and Selene (with two credential/keys) seems ideal for this task.

Let us assume that the voter has two computing devices D_1 and D_2. We store the secret key of the designated verifier key dvk_i on D_2. The voter now uses device D_1 for the registration where the voter gets the credential from the registration. The credential is then stored on D_1, and possibly with secure backups. Note that, during registration, only the public key dvk_i is needed, thus device D_2 can be excluded from this process.

A coerced voter can provide fake proofs without using device D_1, i.e., by only using the secret key on device D_2. Thus device D_2 does not learn the credential.

Vote-casting can be done on device D_1 since it holds the credential, but does not need device D_2. Finally, tracker retrieval and vote verification can be done on D_2 without using D_1.

In order to perform an undetected change of the vote, an adversary needs to infect both device D_1 to get the correct credential, and device D_2 in order to fake the verification of the final tallied vote with the Selene mechanism. Since the devices could be very independent, e.g. the check of the final vote could even be done on some public PC (with a threat of a privacy attack, of course), this seems to greatly reduce the danger from malware.

7.4 Using JCJ to Improve Selene

The combination of JCJ and Selene cannot only be used to add extra verifiability to JCJ, but can also provide a more secure tracker retrieval in Selene. The point is that the voter can authenticate herself with her credential. We can use this to make the tracker retrieval active. That is, instead of the Tracker Tellers sending out the $g^{r_{ij}}$ terms, with the risk of the coercer intercepting the message, the voter contacts the Tracker Tellers to obtain the terms. We will here briefly sketch the idea.

The voters can identify themselves to the Tracker Tellers with a ciphertext of the credential. Here and in the following such encrypted credentials should be followed by zero-knowledge proofs of plaintext knowledge of a special form that makes sure that it is not copied from e.g. already cast election ballots, or reused for ballots or authentication in later elections. For clarity we will suppress these proofs in the following. The Tally Tellers can now perform a PET with the registered credential (while also checking the zero-knowledge proof) to check the authenticity. After authentication, the terms $g^{r_{ij}}$ are handed out.

Coerced voters need to have a time window between the publication of the tally board and the start of the tracker retrieval, where they can upload a fake $g^{r_{ij}}$ term to each Tracker Teller. They do this via an anonymous channel

$$V_i \rightarrow TT_j : V_i, \{C_1\}_{\mathsf{pk}_T}, \{C_2\}_{\mathsf{pk}_T}, \{(g^{r_{ij}})_{\mathsf{fake}}\}_{\mathsf{pk}_T}.$$

The first plaintext is supposed to be the real credential, the second plaintext the faked credential (different from C_1) and the third plaintext is the faked term that will be shown, when someone with credential C_2 tries to retrieve their tracker share. The Tally Tellers need to be invoked to get this term. The fake term could also be sent in plain, if the channel is considered untappable for the coercer.

Now, if a coercer tries to retrieve the random term, the voter should have made a faking request beforehand, and the coercer gets the faked term.

However, we need to be careful since the coercer should not be able to use the update mechanism to discover that he is in the possession of a faked credential. We thus proceed as follows. After the time window, each Tracker Teller now has a database for each voter with rows of faking requests (which might come from the coercer as well). For understandability, we assume that each voter has maximally one coercer, and we can then weed this list so that the value of the first credential C_1 can only appear once, copies are removed via PETs. A retrieval request now takes the form of $V_i, \{C\}_{\mathsf{pk}_T}$. The Tracker Teller now processes this request via the following algorithm which has two memory slots. Before beginning, the ciphertext of the registered credential is loaded to memory slot 1 and the real value $g^{r_{ij}}$ is loaded to slot 2. For the given request TT_j requests a PET of the submitted ciphertext with the value in memory slot 1. If the PET is successful, it hands out the stored value in memory slot 2 and exits the algorithm. If not, it requests PETs against the database C_1s and the value in memory slot 1. If there is no success, it hands out a random number and exits the algorithm, but if there is a success it stores the corresponding C_2 ciphertext in memory slot 1, and the fake value in memory slot 2, suppresses the corresponding database row for the current session and reiterates the algorithm, now essentially acting as if the fake credential was a real credential, but with the fake value in memory slot 2. The algorithm stops since the database is finite.

The coercer only has a negligible chance of guessing the real credential. Thus with overwhelming probability, if the coercer asks for tracker retrieval, the algorithm will after its first step simulate that the credential, handed to the coercer by the voter, is the real one with a corresponding faked term. In this way the retrieval mechanism will act as if the coercer has a real credential. Note that timing might be a side channel attack here, so some default delay is required in the response time.

The advantage of the system is that also coerced voters can safely do verification of their votes, the disadvantage is a rather complicated system, and the voter still needs to be active to fake their trackers.

Another advantage is that the mechanism can also be used to defend against the attack mentioned in Subsect. 6.4. To do this, the retrieval mechanism also performs PETs of the registered fake credential (which was given to the coercer) and the credentials contained in the cast ballots. For a positive PET the corresponding vote can be disclosed to the voter, who can then calculate the fake term. This would require another step to update the fake term. Alternatively, the voter can even beforehand fake terms for all possible candidates and further PET checks between the vote cast by the coercer and these faked terms can then directly update the fake value.

A more careful description of this retrieval mechanism and corresponding security proofs are postponed for future work.

8 Conclusions and Future Work

We have shown that it is possible to use the Selene mechanism in JCJ, providing an e-voting protocol where voters can individually check that their vote was counted as intended, while still preserving a good level of coercion-resistance. Further, several alternatives were presented providing: better verifiability (while only handling a single key), everlasting privacy, a more secure tracker retrieval and better protection against malware on the voters' computing devices. Also improvements to Selene, in terms of efficiency, and JCJ, in terms of coercion-resistance, were presented.

This paper did not provide formal proofs of the security guarantees. These are currently under consideration for the classical Selene protocol in the UC framework, and should later be extended to also include this work.

Two main problems of JCJ were not touched upon, namely, efficiency and usability. Regarding usability, Selene, in some sense is a step backwards. The users (in the first version of the protocol at least) needs to handle two keys post-registration. And coerced voter have to careful when they retrieve the trackers. Further investigations are necessary to determine to which extent this can be handled by the voter assisting devices, and if the extra clarity and trust given by the check of the final vote will outweigh this. We however, also plan to increase the usability of JCJ in the future by allowing the voters to use short codes. Finally regarding efficiency, the versions of JCJ presented here still suffer from the tally time being quadratic in the number of voters, a problem we will also try to solve in future a work.

Acknowledgements. Vincenzo Iovino is supported by the Luxembourg National Research Fund (FNR grant no. 7884937). Further, this work is also supported by the INTER-Sequoia project from the Luxembourg National Research Fund, which is joint with the ANR project SEQUOIA ANR-14-CE28-0030-01.

References

[ACKR13] Arapinis, M., Cortier, V., Kremer, S., Ryan, M.: Practical everlasting privacy. In: Basin, D., Mitchell, J.C. (eds.) POST 2013. LNCS, vol. 7796, pp. 21–40. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36830-1_2

[Adi08] Adida, B.: Helios: web-based open-audit voting. In: Proceeding of 17th USENIX Security Symposium, pp. 335–348 (2008)

[Ben06] Benaloh, J.: Simple verifiable elections. In: Wallach, D.S., Rivest, R.L. (eds.) 2006 USENIX/ACCURATE Electronic Voting Technology Workshop, EVT 2006, Vancouver, BC, Canada, 1 August 2006. USENIX Association (2006)

[BFM88] Blum, M., Feldman, P., Micali, S.: Non-interactive zero-knowledge and its applications (extended abstract). In: 20th Annual ACM Symposium on Theory of Computing, pp. 103–112. ACM Press, May 1988

[BR93] Bellare, M., Rogaway, P.: Random oracles are practical: a paradigm for designing efficient protocols. In: Ashby, V. (ed) ACM CCS 1993: 1st Conference on Computer and Communications Security, pp. 62–73. ACM Press, November 1993

[CCM08] Clarkson, M.R., Chong, S., Myers, A.C.: Civitas: a secure voting system. In: IEEE Symposium on Security and Privacy (2008)

[Cha01] Chaum, D.: Surevote: technical overview. In: Proceeding of Workshop on Trustworthy Elections (WOTE 2001) (2001)

[CMR+16] Clark, J., Meiklejohn, S., Ryan, P.Y.A., Wallach, D., Brenner, M., Rohloff, K. (eds.): FC 2016. LNCS, vol. 9604. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53357-4

[CP93] Chaum, D., Pedersen, T.P.: Wallet databases with observers. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 89–105. Springer, Heidelberg (1993). https://doi.org/10.1007/3-540-48071-4_7

[CPP13] Cuvelier, É., Pereira, O., Peters, T.: Election verifiability or ballot privacy: do we need to choose? In: Crampton, J., Jajodia, S., Mayes, K. (eds.) ESORICS 2013. LNCS, vol. 8134, pp. 481–498. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40203-6_27

[Fre00] Freedman, M.J.: Design and analysis of an anonymous communication channel for the free haven project (2000). http://www.freehaven.net/doc/comm.ps

[Gam85] El Gamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. IEEE Trans. Inf. Theory **31**(4), 469–472 (1985)

[GRCC15] Grewal, G.S., Ryan, M.D., Chen, L., Clarkson, M.R.: Du-vote: remote electronic voting with untrusted computers. In: Fournet, C., Hicks, M.W., Viganò, L. (eds.) IEEE 28th Computer Security Foundations Symposium, CSF 2015, Verona, Italy, 13–17 July 2015 (2015)

[HH07] Hans, D., Helmut, K.: Intorduction to cryptography-principles and applications (2007)

[HK14] Haenni, R., Koenig, R.: Voting over the Internet on an insecure platform. In: Design, Development, and Use of Secure Electronic Voting Systems, chapter IGI Global, March 2014

[HL09] Heather, J., Lundin, D.: The append-only web bulletin board. In: Degano, P., Guttman, J., Martinelli, F. (eds.) FAST 2008. LNCS, vol. 5491, pp. 242–256. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-01465-9_16

[JCJ05] Juels, A., Catalano, D., Jakobsson, M.: Coercion-resistant electronic elections. In: Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society, WPES 2005, Alexandria, VA, USA, 7 November 2005, pp. 61–70 (2005)

[JJ00] Jakobsson, M., Juels, A.: Mix and match: secure function evaluation via ciphertexts. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 162–177. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-44448-3_13

[KHF11] Koenig, R., Haenni, R., Fischli, S.: Preventing board flooding attacks in coercion-resistant electronic voting schemes. In: Camenisch, J., Fischer-Hübner, S., Murayama, Y., Portmann, A., Rieder, C. (eds.) SEC 2011. IAICT, vol. 354, pp. 116–127. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21424-0_10

[KMST16] Küsters, R., Mueller, J., Scapin, E., Truderung, T.: Select: a lightweight verifiable remote voting system. In: IEEE 29th Computer Security Foundations Symposium, CSF 2016, Lisbon, Portugal, 27 June-1 July 2016, pp. 341–354. IEEE Computer Society (2016)

[KOKV11] Karayumak, F., Olembo, M.M., Kauer, M., Volkamer, M.: Usability analysis of helios - an open source verifiable remote electronic voting system. In: Proceeding of Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (EVT/WOTE 2011) (2011)

[KR16] Kremer, S., Rønne, P.B.: To du or not to du: a security analysis of du-vote. In: IEEE European Symposium on Security and Privacy, EuroS&P 2016, Saarbrücken, Germany, 21–24 March 2016, pp. 473–486. IEEE (2016)

[LHK16] Locher, P., Haenni, R., Koenig, R.E.: Coercion-resistant internet voting with everlasting privacy. In: Clark, J., Meiklejohn, S., Ryan, P.Y.A., Wallach, D., Brenner, M., Rohloff, K. (eds.) FC 2016. LNCS, vol. 9604, pp. 161–175. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53357-4_11

[MN06] Moran, T., Naor, M.: Receipt-free universally-verifiable voting with everlasting privacy. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 373–392. Springer, Heidelberg (2006). https://doi.org/10.1007/11818175_22

[NFVK13a] Neumann, S., Feier, C., Volkamer, M., Koenig, R.E.: Towards a practical JCJ/civitas implementation. IACR Cryptology ePrint Archive **2013**, 464 (2013)

[NFVK13b] Neumann, S., Feier, C., Volkamer, M., Koenig, R.E.: Towards A practical JCJ/civitas implementation. In: Horbach, M. (ed) Informatik 2013, 43. Jahrestagung der Gesellschaft für Informatik e.V. (GI), Informatik angepasst an Mensch, Organisation und Umwelt, 16–20 September 2013, Koblenz, Deutschland, vol. 220 of LNI, pp. 804–818. GI (2013)

[NV12] Neumann, S., Volkamer, M.: Civitas and the real world: problems and solutions from a practical point of view. In: Seventh International Conference on Availability, Reliability and Security, Prague, ARES 2012, Czech Republic, 20–24 August 2012, pp. 180–185 (2012)

[Ped91] Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 129–140. Springer, Heidelberg (1992). https://doi.org/10.1007/3-540-46766-1_9

[Roe16] Roenne, P.B.: JCJ with improved verifiability guarantees. In: The International Conference on Electronic Voting E-Vote-ID 2016, 18–21 October 2016, Lochau/Bregenz, Austria (2016)

[RRI15] Ryan, P.Y.A., Rønne, P.B., Iovino, V.: Selene: voting with transparent verifiability and coercion-mitigation. IACR Cryptology ePrint Archive, 2015:1105 (2015)

[RRI16] Ryan, P.Y.A., Rønne, P.B., Iovino, V.: Selene: voting with transparent verifiability and coercion-mitigation. In: Clark, J., Meiklejohn, S., Ryan, P.Y.A., Wallach, D., Brenner, M., Rohloff, K. (eds.) FC 2016. LNCS, vol. 9604, pp. 176–192. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53357-4_12

[RT09] Ryan, P.Y.A., Teague, V.: Pretty good democracy. In: Christianson, B., Malcolm, J.A., Matyáš, V., Roe, M. (eds.) Security Protocols 2009. LNCS, vol. 7028, pp. 111–130. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36213-2_15

[Sch94] Schneier, B.: Applied Cryptography (1994)

[SK95] Sako, K., Kilian, J.: Receipt-free mix-type voting scheme. In: Guillou, L.C., Quisquater, J.-J. (eds.) EUROCRYPT 1995. LNCS, vol. 921, pp. 393–403. Springer, Heidelberg (1995). https://doi.org/10.1007/3-540-49264-X_32