

Eos a Universal Verifiable and Coercion Resistant Voting Protocol

Ştefan Patachi and Carsten Schürmann^(✉)

DemTech, IT University of Copenhagen,
Copenhagen, Denmark
{stpa, carsten}@itu.dk

Abstract. We present the voting protocol Eos that is based on a conditional linkable ring signatures scheme. Voters are organized in rings allowing them to sign votes anonymously. Voters may assume multiple pseudo identities, one of which is legitimate. We use the others to signal coercion to the Election Authority. Eos uses two mixing phases with the goal to break the connection between the voter and vote, not to preserve vote privacy (which is given already) but to guarantee coercion resistance by making it (nearly) impossible for a coercer to follow their vote through the bulletin board. Eos is universally verifiable.

1 Introduction

Most well-known voting protocols use a form of mixing to break the connection between vote and voter. Prêt-à-Voter [19], Helios [1], JCJ [12], Civitas [14], encrypt the votes at the time of casting, and then mix them before decrypting them for tabulation. Under the assumptions that at least one mixer is honest, so the argument, it is impossible to link a decrypted vote back to the identity of its voter. BelniosRF [5] uses randomizable cipher-texts instead. Selene [20] follows a different approach. It uses tracker numbers and assigns them to voters, who eventually will be able to identify their votes on a bulletin board but only after the election authority has shared cryptographic information with the voter. In Selene, voters can fool a coercer into believing that *any* and not just *one* vote on the bulletin board was theirs. Both JCJ and Selene are receipt-free and coercion-resistant.

To our knowledge not much work has been done to leverage the power of ring signatures [18] to the design of voting protocols, besides perhaps the mention in [7, 13, 22]. Here, a ring refers to a group of participants who have the ability to sign messages anonymously by hiding their respective identities within the group. Assuming that the message is sent over an anonymous channel, the receiver of

C. Schurmann—This work was funded in part through the Danish Council for Strategic Research, Programme Commission on Strategic Growth Technologies under grant 10-092309. This publication was also made possible by NPRP grant NPRP 7-988-1-178 from the Qatar National Research Fund (a member of Qatar Foundation). The statements made herein are solely the responsibility of the authors.

the message will not be able to trace the signature back to the signer. Ring signatures bring also other advantages. For example, an election authority will be able to publish all ballots and their signatures on a public bulletin board for public scrutiny without revealing the voters' identities. Every voter can check their vote by accessing the bulletin board. But there are also challenges: First, the administration of the voter's identities, i.e. private keys, and second the protection of such identities from misuse, for example, for the purpose of vote-selling or voter coercion.

For the first challenge, we do not provide a solution in this paper, we merely assume that an effective, trusted identity infrastructure is in place that allows a voter to access a private key on a secure platform, for example, by means of a trusted Hardware Security Module (HSM), such as for example, a Bitcoin wallet called Trezor. This may seem like a controversial assumption, but it really is not: Our experiments have shown that suitable HSMs exist. They may not be totally secure, but they are reasonable well designed to protect private keys against malicious agents and against malicious firmware. Hacking such an HSM is possible and requires fiddling with firmware and hardware, but we argue that this is difficult to do on a large scale in practice.

In this paper, we tackle a second challenge. We devise a ring signature scheme that allows voters to assume different pseudo identities, and that provides mechanisms for linking such identities from the point of view of the signature verifier. This scheme is a generalization of the so called linkable spontaneous anonymous group (LSAG) signatures [13], where all signatures from the same signer are linked. Correspondingly, we refer to this scheme as a *conditional-linking ring* (CLR) signature scheme. Using CLR signatures it is up to the voter to allow linking or not. Linking does not break anonymity. Compared to BeliniosRF [5], which is not based on conditional linkability but on signatures based on randomizable cipher-texts [3], anonymity follows directly from the definition of ring signatures and does not have to be added by another mechanism.

Based on CLR signatures, we develop the Eos voting protocol. Every voter is assumed to have physical access to an HSM and a way to authenticate. The HSM guards the voter's private key. The authentication method yields additional entropy, such as, for example, a PIN number, a picture, or a fingerprint. When casting a vote, before submission, it is either put into a "green envelope" marking it as valid, or a "red envelope" marking it as invalid or possibly coerced.

The entropy collected during the authentication procedure determines which pseudo-identity is used. We encrypt the color of the envelope and the electoral identity of the voter (which is unique), alongside the vote itself. All envelopes together with their corresponding CLR signatures are recorded and shared on a public bulletin board. Pseudo identities are malleable, which means that from the point of view of the voter or the coercer, it will be "discrete logarithm hard" to link any two ballots from the bulletin board. Voters and coercers will be able to check whether the ballot was recorded as cast but neither the voter nor the coercer will be able to link their ballots.

The protocol proceeds and executes two verifiable decryption mixes. The first mix-net shuffles the ballots and decrypts color and the electoral identity, but leaves the votes encrypted. All information is posted on a public bulletin board accompanied by zero-knowledge proofs of correct shuffle and correct decryption. Red envelopes are discarded and green envelopes are opened. The encrypted votes contained within (and only those) are passed to the second and last mixing stage. If two green envelopes are cast from two different pseudo-identities that correspond to the same electoral identity, this indicates a malicious attempt to double vote. In this case we discard all votes associated with this electoral identity. In the second and final mixing stage, we use a verifiable decryption mix to shuffle and decrypt votes. Also here, the resulting bulletin board is published including zero-knowledge proofs for public scrutiny.

For the subsequent security analysis of Eos, we consider an adversarial environment where only voters and their respective HSM devices have to be trusted. Election authorities, mixing nodes, and tellers may be assumed to be under the adversary's control. The adversary is assumed to have the usual adversarial capabilities, which includes deletion, addition or modification of any kind of information, including votes, logs entries, or messages. We show that Eos is universally verifiable.

Closely related to the adversary is the coercer, an agent that aims to exert undue influence onto the voter, for example by forcing the voter to vote in a particular way. In addition also election authorities, mix nodes and tellers, may be under the coercer's control. A coercer may be physically present in the room when voting takes place, or he may observe the voting procedure remotely. We consider a coercion attempt successful, if the coercer succeeds to force voter to vote a particular way and the coerced vote is then included in the final tally, in a way that is observable by the coercer. For the analysis of coercion resistance to make sense, we must assume that the coercer and the voter are different principals, and we do this by assuming that only the voter can authenticate successfully to the HSM device. In this paper we consider three different attacks against the coercion-freeness: (1) A coercer steals a voter's identity, (2) a coercer obtains a receipt by tracking a vote through the system and (3) a coercer steals the HSM and impersonates the voter.

Eos is constructed to be resistant against all three kinds of coercion attacks: Because of the use of CLR signatures, a coercer can always be assigned an alternate valid identity (1). Under a commonly accepted assumption, it is always possible to guarantee that not all mixing nodes are under the coercer's control, which means that neither voter nor coercer will be able to track the votes, but will have to revert to checking zero-knowledge proofs of knowledge (2). And finally, the authentication mechanisms provided by an HSM can be made arbitrarily complex, making it increasingly hard for a coercer to brute-force his way in (3).

Contributions. The contributions of this paper are (1) a Conditional Linkable Ring (CLR) signature scheme, described in Sect. 3, (2) an efficient mixer that satisfies special soundness but still produces a proof of correct shuffle ($2n$ modexp operations to generate and $4n$ modexp operations to verify shuffle proofs),

described in Sect. 4, (3) the Eos voting protocol, described in Sect. 5, and (4) proofs that the Eos is vote secrecy and integrity preserving, universally verifiable, receipt-free and coercion resistant, discussed in Sect. 6.

2 Basic Notations

We define the following notations that we will use throughout our paper. We work with an ElGamal cryptosystem that is defined over a cyclic group \mathbb{G} of prime p order q generated by g . All mathematical operations presented in this paper are done modulo p .

We use standard notation and write $\{m\}_y^r = (g^r, y^r m)$ for the ElGamal tuple that one obtains by encrypting message m with randomness r and a public key y . We use r to denote randomness, and write $r \in_{\mathbb{R}} \mathbb{Z}_q$ to express that we choose r from \mathbb{Z} modulo q at random using a uniform distribution. We will also use sequences of n elements over \mathbb{Z}_q , written as $\langle x_1, \dots, x_n \rangle \in \mathbb{Z}_q^n$. We define $[n]$ to denote the index set of natural numbers up to n as $\{1, \dots, n\}$. Furthermore, we use Greek letters σ for signatures and π to denote permutations. We write \mathbb{P}_n as the set of all permutation of n elements. Concatenation of two elements $a, b \in \{0, 1\}^*$ is written as $a||b$.

3 Conditional-Linkable Ring Signatures

We begin the technical exposition of this paper by defining the concept of Conditional Linkable Ring (CLR) signatures. In a linkable ring (LR) signature scheme [13], a verifier can learn which signatures originate from the same signer. Note that this does not mean that the verifier learns something about the identity of the signer — ring signatures always guarantee the anonymity of the signer. For our application however, linkability is overly restrictive — if we were to use LR signatures naively, we could not achieve coercion-resistance. Therefore, we relax the notion of linkability, and introduce *conditional linkability* giving the signer the ability to link (revealing that the signatures originate from the same signer) or not to link (making it look like as if two signatures were produced by two different signers) and *claimability* allowing the signer to claim the authenticity of a signature by proving his position in the ring in zero knowledge. When a signature is claimed, the signer reveals his position in the ring and loses anonymity.

Preparation Phase: Every prospective member of the ring creates a secret key $x_i \in \mathbb{Z}_q$, and shares the public key $y_i = g^{x_i}$ with a designated election authority.

Set-Up Phase: The election authority produces the set of ring members $L = \langle y_1, \dots, y_n \rangle$, which represents eligible voters.

Identity Selection Phase: Assume that the signer is the member of the ring at position α . The signer selects a pseudo-identity by choosing $\phi \in_{\mathbb{R}} \mathbb{G}$ and by forming the pair (ϕ, θ) where $\theta = \phi^{x_\alpha}$. The pair (ϕ, θ) is called a pseudo identity. If the signer wishes his signature to be linkable, he will always choose the same value of ϕ , otherwise, he will choose a different value for ϕ for every signature.

CLR signatures give us quite a lot of freedom to choose ϕ . To see that LR signatures proposed in [13] are simply an instance of CLR signatures, choose ϕ to be the cryptographic hash value of L (written as $h = H_2(L)$ in their paper) and compute $\theta = h^{x_\alpha}$. Liu et al. denote this value as \tilde{y} . The security of linkability reduces therefore to a secure choice of ϕ , for which it is sufficient to require that finding $\log_g \phi$ is a hard problem.

Signing Algorithm. We begin now with the presentation of the signing algorithm. Our algorithm follows closely the original LR signing algorithm described in [13], the most notable difference being that we use ϕ and θ instead of h and \tilde{y} , respectively. Given the message to be signed $m \in \{0, 1\}^*$, for each element in the ring L , a cipher text is computed, starting from the position of the signer in the ring, α , to the end of the list and from the beginning of the list back to α .

The first computed cipher text is therefore

$$c_{\alpha+1} = \mathcal{H}(m \| g^u \| \phi^u)$$

where, \mathcal{H} is a cryptographic hash function that returns a number from \mathbb{Z}_q (referred to as H_1 in [13]) and $u \in_{\mathbb{R}} \mathbb{Z}_q$. Next, for each element in L from $i = \alpha + 1$ to n and from $i = 1$ to $\alpha - 1$, the signer computes:

$$c_{i+1} = \mathcal{H}(m \| g^{s_i} \cdot y_i^{c_i} \| \phi^{s_i} \cdot \theta^{c_i})$$

where each $s_i \in_{\mathbb{R}} \mathbb{Z}_q$ is a random number assigned for each entity in L . Note that at step $i = n$, we generate $c_1 = c_{i+1}$. The signer computes:

$$s_\alpha = u - x_\alpha \cdot c_\alpha \pmod q$$

Finally, the output of the CLR signing algorithm is the signature σ on message m with the pseudo-identity (ϕ, θ) that has the following structure:

$$\sigma(m) = (c_1, \langle s_1, \dots, s_n \rangle)$$

Verification Algorithm. After having received a message m and the corresponding signature, $\sigma(m) = (c_1, \langle s_1, \dots, s_n \rangle)$ from the pseudo-identity (ϕ, θ) , anybody can now verify the signature by executing the following steps of the verification algorithm, which is computationally linear in terms of size of L , that should output either the signature is valid or not, i.e. it was generated by a ring member or not. For each element in L starting from $i = 1$ to n compute:

$$c_{i+1} = \mathcal{H}(m \| g^{s_i} \cdot y_i^{c_i} \| \phi^{s_i} \cdot \theta^{c_i})$$

The algorithm validates the signature if and only if the last $c_{n+1} = c_1$, where c_1 is contained as the first argument in the signature.

Discussion. The idea to consider other ways to compute ϕ was already discussed in [13]. What is new in our work is to allow ϕ to be drawn at random from \mathbb{G} . We shall see in Sect. 5 how to choose ϕ to encode pseudo-identities.

The definition, properties and proofs of the signature scheme presented in [13], such as existential unforgeability and signer ambiguity, and the rewind-on-success lemma, carry over verbatim to the CLR signature scheme. CLR signatures are conditionally linkable and claimable. Once generated, the signer and only the signer can claim responsibility of a signature generated by him by providing in zero knowledge a discrete logarithm equality proof between $\log_{\phi} \theta = \log_g y_{\alpha}$.

4 Mix-Net

Next, we describe the mix-net the we will be using. The goal of the mixing is to shuffle ballots in such a way that it is impossible to correlate outputs to inputs.

The mixing task is standard, the literature on mix-nets that emit proofs of correct (parallel) shuffle is mature. In essence, any state-of-the-art mix-net could be used as long as it supports parallel shuffle, for example, Bayer [2], Groth [10,11], Wikström [23], or more recently Fauzi and Lempaa [8] whose security proof no longer relies on the Random Oracle Model.

While analyzing the different mix-net protocols, in particular [9,15,16,21], we observed simplifications to Neff’s protocol that we describe next. These make the proof of parallel shuffle more efficient while still satisfying special soundness, justifying why we have included a description of the *simplified Neff’s protocol* in this paper. The protocol is mathematically elegant, intuitive, and easier to understand than Neff’s original protocol as it follows closely the classic discrete logarithm equality zero knowledge proof.

Our proof of shuffle also relies on the Random Oracle Model. In comparison with Groth’s work, our proof has only three rounds and requires only $2n$ modexps for generating a proof and $4n$ modexps for verifying a proof. On the other hand, Bayer’s [2] protocol requires only sub-linear operations but needs 9 rounds to finish.

Below, we provide a semi-formal analysis that our proposed protocol is a zero-knowledge proof of correct shuffle.

4.1 Proof of Correct Shuffle

Our mix-net consists of several mixing servers, each of which receives as input a bulletin board of n ElGamal tuples (a_i, b_i) and produces an output bulletin board, where all tuples are re-encrypted and shuffled:

$$(c_i, d_i) = (a_{\pi(i)} \cdot g^{s_{\pi(i)}}, b_{\pi(i)} \cdot y^{s_{\pi(i)}}) \text{ for } i \in [n]$$

where y is the encryption key, $\langle s_1, \dots, s_n \rangle \in_{\mathbb{R}} \mathbb{Z}_q^n$ the randomness used for re-encryption, and $\pi \in_{\mathbb{R}} \mathbb{P}_n$ the permutation underlying the shuffle.

The challenge when defining a mix-net is how each mixing server can prove the correctness of the shuffle to a public verifier, without revealing information about the permutation π or the randomness $\langle s_1, \dots, s_n \rangle$ used for re-encryption.

The following proof of correct shuffle is inspired by the protocol developed by Sako and Kilian [21], where they say that the proof should show that the

\mathcal{P} secretly generates: $k \in_{\mathbb{R}} \mathbb{Z}_q$ and $\langle m_1, \dots, m_n \rangle \in_{\mathbb{R}} \mathbb{Z}_q^n$ and publishes commitment:

$$A = g^k \cdot \prod_{i \in [n]} a_i^{m_i} \qquad B = y^k \cdot \prod_{i \in [n]} b_i^{m_i}$$

\mathcal{V} sends challenge: $\langle e_1, \dots, e_n \rangle \in_{\mathbb{R}} \mathbb{Z}_q^n$

\mathcal{P} publishes response:

$$r_i = m_i + e_{\pi^{-1}(i)} \pmod q \text{ for } i \in [n]$$

$$t = k + \sum_{i \in [n]} (e_i \cdot s_{\pi(i)}) \pmod q$$

\mathcal{V} accepts the proof if the following verification calculations match:

$$g^t \cdot \prod_{i \in [n]} a_i^{r_i} = A \cdot \prod_{i \in [n]} c_i^{e_i} \qquad y^t \cdot \prod_{i \in [n]} b_i^{r_i} = B \cdot \prod_{i \in [n]} d_i^{e_i}$$

Fig. 1. Protocol: proof of correct shuffle

output of the mixer could be generated in some manner from the input. Finally, the aggregation of the entire set of ElGamal pairs, is inspired by Ramchen’s work [17]. Our proof follows the natural flow of a classic discrete logarithm equality zero knowledge proof depicted in Fig. 1, i.e. the mix server publishes a commitment of the input, a verifier challenges the output of the mixer and then mixer generates a response, which convinces the verifier that the shuffle was correct.

Let (a_i, b_i) be the n ElGamal tuples that form the input for the mix server. Let (c_i, d_i) be n ElGamal tuples, computed as above, be the output of the mix server. To prove the correctness of the shuffle, the mix server \mathcal{P} and a public verifier \mathcal{V} have to follow the protocol that is described in Fig. 1.

Theorem 1. *The protocol described in Fig. 1 is complete.*

Proof. To show that our protocol is correct, we have to prove that the equations that \mathcal{V} verifies hold when the response $(\langle r_1, \dots, r_n \rangle, t)$ is computed correctly.

$$g^t \cdot \prod_{i \in [n]} a_i^{r_i} = A \cdot \prod_{i \in [n]} c_i^{e_i}$$

$$g^{k + \sum_{i \in [n]} e_i \cdot s_{\pi(i)}} \cdot \prod_{i \in [n]} a_i^{m_i + e_{\pi^{-1}(i)}} = g^k \cdot \prod_{i \in [n]} a_i^{m_i} \cdot \prod_{i \in [n]} (a_{\pi(i)} \cdot g^{s_{\pi(i)}})^{e_i}$$

$$g^k \cdot g^{\sum_{i \in [n]} e_i \cdot s_{\pi(i)}} \cdot \prod_{i \in [n]} a_i^{m_i} \cdot \prod_{i \in [n]} a_i^{e_{\pi^{-1}(i)}} = g^k \cdot \prod_{i \in [n]} a_i^{m_i} \cdot \prod_{i \in [n]} (a_{\pi(i)} \cdot g^{s_{\pi(i)}})^{e_i}$$

$$\prod_{i \in [n]} g^{e_i \cdot s_{\pi(i)}} \cdot \prod_{i \in [n]} a_i^{e_{\pi^{-1}(i)}} = \prod_{i \in [n]} a_{\pi(i)}^{e_i} \cdot \prod_{i \in [n]} g^{s_{\pi(i)} \cdot e_i}$$

$$\prod_{i \in [n]} a_i^{e_{\pi^{-1}(i)}} = \prod_{i \in [n]} a_{\pi(i)}^{e_i}$$

The last equation in the proof is true because the product aggregation happens through the entire set of n elements. This means we can compute the product aggregation of $a_i^{e_{\pi^{-1}(i)}}$ in a permuted way, namely $a_{\pi(i)}^{e_{\pi(\pi^{-1}(i))}} = a_{\pi(i)}^{e_i}$.

In the same way, the second equations that \mathcal{V} has to verify can be proven to hold if the response $(\langle r_1, \dots, r_n \rangle, t)$ is computed correctly. \square

Theorem 2. *The protocol described in Fig. 1 satisfies special soundness.*

Proof. Each transcript of our protocol has the following form:

$$\text{View}[\mathcal{P} \leftrightarrow \mathcal{V}] = (A, B, \langle e_1, \dots, e_n \rangle, \langle r_1, \dots, r_n \rangle, t)$$

where A and B represent the initial commitment, sequence $\langle e_1, \dots, e_n \rangle$ is the random challenge picked by the verifier, and the sequence $\langle r_1, \dots, r_n \rangle$ together with the value t represent the response to the challenge.

For any cheating prover \mathcal{P}^* (that does not know the permutation $\pi(i)$ and the re-encryption coefficients $\langle s_1, \dots, s_n \rangle$), given two valid conversations between \mathcal{P} and the verifier \mathcal{V} , $(A, B, \langle e_1, \dots, e_n \rangle, \langle r_1, \dots, r_n \rangle, t)$ and $(A, B, \langle e'_1, \dots, e'_n \rangle, \langle r'_1, \dots, r'_n \rangle, t')$ that have the same commitment but different challenge $e_i \neq e'_i$, the permutation $\pi(i)$ used for shuffling the board can be computed in polynomial time in the following way:

$$\forall i \in [n], \exists p \text{ such that } \pi(i) = p, \text{ where } r_p - r'_p = e_i - e'_i$$

The permutation $\pi(i)$ is the actual secret the mixing server has to hide. The re-encryption coefficients $\langle s_1, \dots, s_n \rangle$ are also assumed to be kept secret.

This is precisely what our choice of re-encryption mechanism guarantees. \square

Theorem 3. *The protocol described in Fig. 1 is honest verifier zero knowledge.*

Proof. We prove that for any cheating verifier \mathcal{V}^* , there exists a simulator \mathcal{S} that can produce a computationally indistinguishable transcript of the protocol that would take place between \mathcal{P} and \mathcal{V}^* if it knew the challenge in advance.

Our simulator \mathcal{S} gets as input: the initial set of n ElGamal tuples (a_i, b_i) , the mixed set of ElGamal tuples (c_i, d_i) and a challenge in the form of a random sequence $\langle e_1, \dots, e_n \rangle$. \mathcal{S} proceeds by picking a random response of the transcript:

$$\begin{aligned} \langle r_1, \dots, r_n \rangle &\in_{\mathbb{R}} \mathbb{Z}_q^n \\ t &\in_{\mathbb{R}} \mathbb{Z}_q \end{aligned}$$

\mathcal{S} computes the initial commitment:

$$A = g^t \cdot \prod_{i \in [n]} (a_i^{r_i} \cdot c_i^{-e_i}) \qquad B = y^t \cdot \prod_{i \in [n]} (b_i^{r_i} \cdot d_i^{-e_i})$$

\mathcal{S} outputs the transcript: $(A, B, \langle e_1, \dots, e_n \rangle, \langle r_1, \dots, r_n \rangle, t)$.

It is obvious that the transcript \mathcal{S} outputs will always pass the equations that \mathcal{V} has to verify. Note that this transcript was generated independently of the permutation $\pi(i)$ and the re-encryption coefficients $\langle s_1, \dots, s_n \rangle$ used for mixing, thus is zero knowledge. \square

4.2 Proof of Correct Parallel Shuffle

The proof of shuffle for mixing individual cipher texts can be extended to a proof of correct parallel shuffle for sequences of ElGamal tuples. Such a parallel mixer expects as input a matrix over ElGamal tuples with n rows and ℓ columns: $(a_{i,j}, b_{i,j})$, where $i \in [n]$ and $j \in [\ell]$, the Mixer then outputs a mixed and re-encrypted matrix where only the rows are shuffled. This matrix is defined as

$$(c_{i,j}, d_{i,j}) = (a_{\pi(i),j} \cdot g^{s_{\pi(i),j}}, b_{\pi(i),j} \cdot y^{s_{\pi(i),j}})$$

where y is the encryption key, $(s_{1,1}, \dots, s_{n,\ell}) \in_{\mathbb{R}} \mathbb{Z}_q^{n \times \ell}$ are the re-encryption coefficients and $\pi \in_{\mathbb{R}} \mathbb{P}_n$ is a permutation.

The proof of correct parallel shuffle depicted in Fig. 2 is designed to convince a public verifier that the same permutation $\pi(i)$ was applied to each column. The proof, inspired by [17], deviates slightly from the construction that we have presented for the simple case in the previous section. By applying the same challenge e_i to all columns in the matrix, the verifier will be assured that the same permutation $\pi(i)$ was applied consistently across all columns.

\mathcal{P} secretly generates: $(k_1, \dots, k_\ell) \in_{\mathbb{R}} \mathbb{Z}_q^\ell$ and $(m_1, \dots, m_n) \in_{\mathbb{R}} \mathbb{Z}_q^n$ and publishes commitment:

$$A_j = g^{k_j} \cdot \prod_{i \in [n]} a_{i,j}^{m_i} \text{ for } j \in [\ell] \qquad B_j = y^{k_j} \cdot \prod_{i \in [n]} b_{i,j}^{m_i} \text{ for } j \in [\ell]$$

\mathcal{V} sends challenge: $(e_1, \dots, e_n) \in_{\mathbb{R}} \mathbb{Z}_q^n$

\mathcal{P} publishes response:

$$r_i = m_i + e_{\pi^{-1}(i)} \text{ mod } q \text{ for } i \in [n]$$

$$t_j = k_j + \sum_{i \in [n]} (e_i \cdot s_{\pi(i),j}) \text{ mod } q \text{ for } j \in [\ell]$$

\mathcal{V} verifies for each $j \in [\ell]$ and accepts the proof if all calculations match:

$$g^{t_j} \cdot \prod_{i \in [n]} a_{i,j}^{r_i} = A_j \cdot \prod_{i \in [n]} c_{i,j}^{e_i} \qquad y^{t_j} \cdot \prod_{i \in [n]} b_{i,j}^{r_i} = B_j \cdot \prod_{i \in [n]} d_{i,j}^{e_i}$$

Fig. 2. Protocol: proof of correct parallel shuffle

Our proof has the same security properties as the simple proof presented in the previous section. Completeness holds as it follows a slightly more generalized version of the calculation done in the proof of Theorem 1, as now we need to take into account index j for each ElGamal tuple in a sequence. Special soundness follows exactly the same arguments as in the proof of Theorem 2. The proof of honest verifier zero knowledge is an elegant generalization of the proof of

Theorem 3: The simulator \mathcal{S} is modified in such a way that it outputs a sequence of initial commitments for each $j \in [\ell]$:

$$A_j = g^{t_j} \cdot \prod_{i \in [n]} (a_{i,j}^{r_i} \cdot c_{i,j}^{-e_i}) \quad B_j = y^{t_j} \cdot \prod_{i \in [n]} (b_{i,j}^{r_i} \cdot d_{i,j}^{-e_i})$$

where $\langle r_1, \dots, r_n \rangle \in_{\mathbb{R}} \mathbb{Z}_q^n$ and $\langle t_1, \dots, t_\ell \rangle \in_{\mathbb{R}} \mathbb{Z}_q^\ell$ represent the response of the challenge $\langle e_1, \dots, e_n \rangle$.

This proof might be seen as an ℓ -run of the simple protocol, to which we feed the same challenge sequence $\langle e_1, \dots, e_n \rangle$. Note that our proof of correct parallel shuffle does not break the honest verifier zero knowledge property because in each run, the prover picks a different value k_j . Moreover, each run of the protocol is applied on a different partial board $(a_{i,j}, b_{i,j})$, for $i \in [n]$. We summarize these findings in form of a theorem.

Theorem 4. *The proof of correct shuffle satisfies completeness, special soundness, and honest verifier zero knowledge.*

As for complexity, the computation cost for the proof of correct parallel shuffle is summarized as follows. To generate a proof of correct shuffle of the entire matrix, a prover will require $2n\ell + 2\ell$ exponentiations and $3n\ell$ multiplications. In contrast, the verifier will be more expensive, because it requires $4n\ell + 2\ell$ exponentiations and $4n\ell$ multiplications.

5 Eos Protocol

CLR signatures and mix-nets are the building blocks of the Eos Voting protocol that we define next. The hallmark characteristics of the protocol is that voters are organized in rings and they can sign their ballots anonymously. The use of mix-nets make it impossible for coercers to trace the coerced ballot. Each voter has the possibility to assume one out of many pseudo identities. This mechanism can be used, as least in theory, to signal to the election authority that the ballot was submitted as coerced. More research is necessary to study how to do this in practice. For the purpose of this paper, we can assume that voter has access to green and red envelopes. A green envelope means that the vote contained within reflects the voter's intention while a red envelope signals coercion. The color of the envelope will be encrypted.

A *voter* is a person that can participate legitimately in the election process. All voters together generate a set of CLR signed ballots as input. Every ballot cast must be signed by an eligible voter, but not every eligible voter is required to cast a ballot. A voter may be under the influence of a *coercer*, who may be colluding with the election authorities, mix nodes, or tellers, to steal a voter's identity, tries to track the vote, or steals the voter's HSM.

The *election authority* administrates the election. Its role is to initialize the election, form the ring for CLR signing and collect the signed ballots cast by the voters. Each ballot is recorded on a public bulletin board allowing the voter to check that it was recorded as cast. The election authority is responsible for starting and supervising the mixing phase. Eos assumes that all bulletin boards are

append-only, but other variants are possible (although not discussed here). Votes are cast for one *candidate* only. The result of Eos is a public bulletin board with all votes recorded in clear text. Eos supports distributed and threshold decryption schemes, which entails that shares of the decryption key are distributed among different tellers.

Election Set-Up Phase. The election authority prepares $L = \langle y_1, \dots, y_n \rangle$, the list of all eligible voter public keys that will form the ring. In addition, the election authority prepares the set of candidates as vote choices $\mathbb{V} \subset \mathbb{Z}_p^*$. We assume that there are several mixing servers maintained by different non-colluding principals, each with access to a good source of entropy. We call a mixing server *honest*, if it is correct and not controlled by either the adversary or the coercer. An honest mixing server does not reveal the permutation used for mixing.

As it is common practice, we use a (t, e) -threshold cryptosystem, as described in [6], to encrypt and decrypt ballots. All ballots will be encrypted with a public key Y , while the corresponding private key is split and shared among e tellers. Recall that in threshold encryption, it takes the shares of at least t tellers in order to decrypt successfully. Decryption will fail, if less than t shares are available.

Voting Phase. The voter commits to the color of the envelope, using the respective private key $x_i \in \mathbb{Z}_q$ associated with a public key $y_i = g^{x_i} \in L$ and some entropy generated during the authentication process. We use both, private key and entropy to derive (deterministically) using a secure hashing function, the randomness needed in the ElGamal encryption. Once the ballot is generated and signed, it is sent to the election authority that publishes it on the (append only) public bulletin board.

Ballot Generation. A ballot consists of three ElGamal tuples: an encryption of the color of the envelope, an encryption of the electoral identity of the signer and an encryption of the vote. *Encryption of the color:* Recall from Sect. 3 the definition of h and \tilde{y} . A green envelope is formed as an encryption of h , whereas the red envelope is an encryption of 1. *Encryption of the electoral identity:* In the case of a green envelope, there will be an encryption of \tilde{y} , while in the case of a red envelope, there will be again an encryption of 1. *Encryption of the vote:* The vote, to be encrypted, will be represented as value $v \in \mathbb{V}$.

The ballot generation algorithm is depicted in Fig. 3. Formally, the ballot generation algorithm for voter α depends on the following inputs, the authentication entropy \mathcal{E} (such as PIN, a picture, a fingerprint), the private key x_α , and an election specific generator h . The first two ElGamal tuples (F, ϕ) and (T, θ) of the generated ballot play an important role in forming the pseudo identity. Let ϕ be the second projection (trap door commitment) of the encryption of the color of the envelope, and θ is the second projection of the encryption of the electoral identity. Together, (ϕ, θ) form the pseudo identity of the signer.

Due to the deterministic nature of the randomness used for ElGamal encryption, a voter is able to generate the same pseudo identity deterministically multiple times. If an implementation of Eos uses, for example, a PIN code as entropy

The algorithm of Ballot Generation starts by the device computing:

$$\begin{aligned} f &= \mathcal{H}(\mathcal{E} \| x_\alpha \| h) \\ t &= f \cdot x_\alpha \bmod q \\ d &\in_{\mathbb{R}} \mathbb{Z}_q \\ (D, \delta) &= \{v\}_Y^d = (g^d, Y^d \cdot v) \end{aligned}$$

If authentication was successful:

$$(F, \phi) = \{h\}_Y^f = (g^f, Y^f \cdot h) \quad (T, \theta) = \{\tilde{y}\}_Y^t = (g^t, Y^t \cdot \tilde{y})$$

If authentication was unsuccessful:

$$(F, \phi) = \{1\}_Y^f = (g^f, Y^f) \quad (T, \theta) = \{1\}_Y^t = (g^t, Y^t)$$

The generated ballot is: $((F, \phi), (T, \theta), (D, \delta))$.

Fig. 3. Algorithm: ballot generation

for authentication, the pseudo identity of the voter is uniquely defined by the choice of PIN. The valid pseudo identity is selected locally on the voting device by correct authentication, i.e. by using the correct PIN. If the same coercer forces the same voter to vote multiple times, Eos will do so, as it computes the same coerced pseudo identity.

In addition, to guarantee the internal consistency of an encrypted ballot, the signer proves in zero knowledge that the encryptions of the color of the envelope and of the electoral identity are correct by providing a proof of the discrete logarithm equality between $\log_F T = \log_\phi \theta$. This means that there will be only one pseudo identity per device, for each value of f . Note that $\log_\phi \theta = x_\alpha$ (i.e. private key of an eligible voter) is enforced by the CLR signature verification algorithm. Together with the encrypted vote, one must include also a proof of knowledge of the discrete logarithm of $\log_g D$. This will protect against vote copying.

A malicious user might also try to cast multiple countable votes by encrypting his electoral identity with different values of f . Obviously, this could happen in theory, but practically this attack would require the malicious voter to tamper with software or hardware to trick the protocol. This however, would be noticed as we discuss later in the description of the Ballot Verification Phase Sect. 5 where the value of \tilde{y} will be decrypted and duplicates will be visible. We suggest, in this case, to discard the multiple votes from the same electoral identity.

Note that only during the Tallying Phase (Sect. 5), the vote v will be visible in plain text. There the public can scrutinize and validate each plain text and

if it represents a valid candidate v , such that $v \in \mathbb{V}$. Otherwise, the vote should be disregarded.

Signing a Ballot. The CLR signature of a ballot is computed as described in Sect. 3. Concretely, the pseudo-identity (ϕ, θ) is embedded in the ballot and the message to be signed is publicly computable.

$$m = \mathcal{H}(D \parallel \delta)$$

The CLR signature will then be computed as:

$$\sigma(m) = (c_1, \langle s_1, \dots, s_n \rangle).$$

Beside the ballot and the signature, a voter has to send also the two zero knowledge proofs described above: one for proving the correct encryptions and the second for proving the knowledge of the vote.

Public Bulletin Board. The public bulletin board is a public file, to which only the election authority is allowed to append to. Each entry on the board contains a ballot, its corresponding CLR signature and two zero knowledge proofs. Note that no ballots will ever be removed from the public bulletin board, only added. Each voter and coercer is able to check that their respective vote have been appended on the bulletin board after submission, hence individually verifiable. Ballots from the same pseudo identity can be related on the board as they have pseudo identity (ϕ, θ) . Assuming that voter and coercer use different pseudo identities, their votes can only be related with a negligible probability.

Ballot Verification Phase. Once the Voting Phase finished and all votes have been collected, the election authority no longer accepts signed ballots and seals the public bulletin board cryptographically. The election authority performs a cleansing operation on the public bulletin board and only copies those ballots (without signatures and zero knowledge proofs) to a new board, for which both zero knowledge proofs are checked and the CLR signature is validated. In the case multiple ballots were cast from the same pseudo identity, only the most recent ballot is copied. The cleansing operation is publicly verifiable. This procedure is visible in Fig. 4, as some of the ballots get crossed out and disregarded.

Parallel Mixing. Before the election authority commences with decrypting the ballots, it first uses the parallel shuffle described in Sect. 4.2 to shuffle the entire bulletin board by re-encrypting all three ElGamal tuples of each entry. We assume that there are multiple mixing servers, at least one of which is honest, which protects confidentiality and prevents a coercer to follow a vote through the protocol. Each mixing server performs a mixing operation on the output of the previous mix server and constructs a mixed board together with a Proof of Correct Parallel Shuffle, which is subsequently validated by the election authority. In case a proof fails, an error has occurred, and the output of this particular server is disregarded and another mixing server is used. After the shuffle is complete neither voters nor coercers will be able to identify their ballots on the mixed board, unless all mixing servers collude.

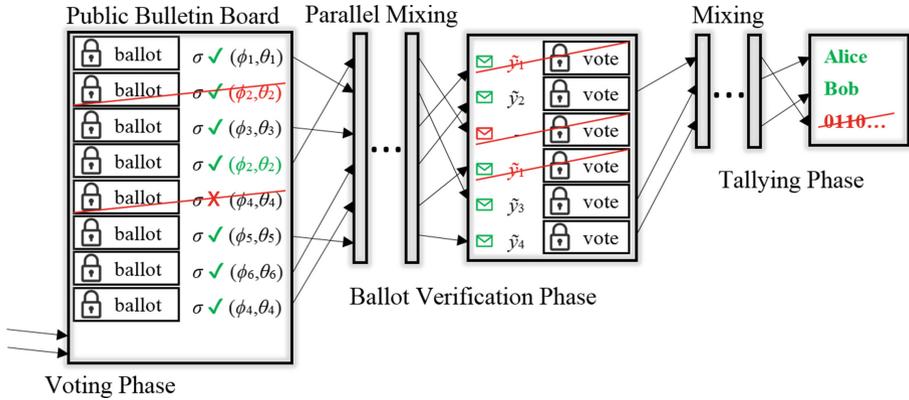


Fig. 4. Protocol overview

Pseudo Identity Decryption. To decrypt the pseudo identities for each entry, t tellers out of e must come together and decrypt the contents of the mixed board using the threshold decryption scheme. At this stage, only the first two ElGamal tuples will be decrypted, i.e. the color of the envelope and the electoral identity of each entry. The vote itself will not be decrypted and this is guaranteed assuming that strictly more than $e - t$ tellers are honest. All ballots whose color of envelope do not decrypt to the value of h or 1 will be disregarded as they are not well-formed. All ballots whose color decrypts to 1 will be discarded because they are coerced. The remaining ballots should all have unique values for the electoral identity. In case there are multiple ballots whose electoral identity decrypts to the same value \tilde{y} , these ballots should be disregarded as they represent an attempt to cast multiple votes. This scenario might happen in case of a malicious voter misusing the Eos protocol. These examples can be seen in Fig. 4 as some ballots are crossed out in the Ballot Verification Phase.

Tallying Phase. The remaining ballots are those that encrypt valid votes that must be counted. To extract the votes from these ballots, we drop the encryptions of the color and electoral identity and create a new bulletin board to undergo another round of mixing before decryption. The bulletin board only contains encrypted votes of the form (D, δ) . This way, we assure that the link between an electoral identity and the vote is broken.

Mixing. Recall that CLR signatures are claimable as the voter can prove in zero knowledge the discrete logarithm equality between $\log_g y_i = \log_h \tilde{y}$. By mixing the list of encrypted votes once more, a voter might only prove to a potential coercer, that he voted but not who he voted for. For this phase, the simple mixing protocol described in Sect. 4.1 is used.

Vote Decryption. Finally, the tellers get together once more and perform a threshold decryption of the board of remaining encrypted votes and produce a proof

of correct decryption. After decryption, each value of v should be counted as a valid vote for the respective candidate if $v \in \mathbb{V}$.

6 Analysis

Eos is individually verifiable, because every voter (and also every coercer) can check if his ballot was correctly recorded by checking the public bulletin board. Eos is also universally verifiable, because all zero knowledge proofs generated by the mixing servers and the tellers are public. Eos is designed to assure that every ballot published on any of the public bulletin boards is mixed and decrypted correctly.

The integrity of the Eos protocol follows from the proofs of correct decryptions for each of the two bulletin boards. Mixing and decryption operations are applied to the entire bulletin board and can be challenged by any public verifier. The removal of invalid and coerced votes from any of the bulletin boards is verifiable, because it can also be challenged by a public verifier as the color (red or green) of the envelopes will be made visible during the Ballot Verification Phase.

The secrecy of the vote is guaranteed by the ElGamal cryptosystem and the use of a cryptographically secure hashing function. The anonymity of the voter is guaranteed by the CLR signature scheme, which protects the voter's true identity. At the same time, we have to assume that there will be at least one honest mixing server that will not disclose its choice of permutation. This assures that a coercer is not able to trace his ballot all the way to the decrypted board and learn if the coerced vote was cast in a green or red envelope. Last but not least, we assume that there will be at least $e - t + 1$ honest tellers to participate in the threshold decryption. This means that we assume that t dishonest tellers will never collude to decrypt the ballots from the bulletin board before the final step of the protocol as this will represent an attack to the fairness of the election.

In terms of receipt-freeness, Eos guarantees that neither a voter nor a coercer can identify his ballot on the decrypted board. This is achieved through two mixing phases which break the connection between the ballot on the public bulletin board and the one on the decrypted board. In addition, a coercer may force a voter to cast a particular vote. In this case, the voter will use one of the alternate pseudo identities to sign the ballot, which will subsequently be discarded during the Ballot Verification Phase.

Eos is constructed in such a way that the pseudo identity used for a coerced vote is computationally indistinguishable from the real pseudo identity of the voter. Even if the coercer had stolen the voter's identity, he would not be able to use it to identify the signer of the CLR signature, because this identity is encrypted and will only be decrypted after the first round of mixing. As Eos is receipt-free, the coercer will not be able to track a coerced vote through the system. And lastly, if a coercer steals the voter's HSM, he might in principle be able to cast a vote for each pseudo-identity. However, clever authentication schemes, possibly even using a combination of pin numbers and biometrics, can be devised to make the space of pseudo identities prohibitively large. We conclude

that Eos is resistant against the three coercion attacks we have outlined in the introduction of this paper.

One bit of power that the coercer has over the voter is that of an abstention attack, to force a vote for a particular candidate for which he knows will receive only this one vote, something like an invalid write-in vote. All the coercer has to do is to check that this vote appears on the final decrypted board of votes. If it does, this would mean that the coercer forced the voter to cast an invalid vote, spoiling the ballot. This situation can be mitigated by the voter proving that his vote is part of the valid set of votes \mathbb{V} without revealing what the vote is, for example using a disjunctive zero knowledge proof protocol as described in [4]. These votes could be cleansed earlier, and would therefore never appear on the final board.

7 Conclusion and Future Work

We have described in this paper a verifiable, privacy preserving coercion-resistant voting protocol that was inspired by Conditional-Linkable Ring (CLR) signatures. Furthermore, we argued for why this protocol protects the integrity of the election, how it guarantees the secrecy of the vote, receipt freeness and is coercion resistant as long as one of the mixing servers is honest. In future work, we plan to reduce the size of CLR signatures from linear to constant size, for example using “accumulators” such as described in [7]. These constant sized signatures can also be made linkable [22].

Our protocol is different from other coercion mitigating protocols, such as Selene [20] or JCJ [12]. In Selene tracker numbers are generated prior to the election, and once a vote is cast, only the trap-door commitment is shared with the voter. After the election is over, the randomness necessary to decrypt the tracker number is shared, allowing each voter to gain confidence in that his or her vote was recorded correctly. Moreover, this protocol allows every voter to trick a potential coercer into believing that he or she voted for the coercer’s choice. In JCJ, every voter has access to different kinds of credentials. One credential is there to be used to cast a valid vote, whereas as the other credentials are there to cast a vote that from the outset looks like a valid vote, but really is not. The election authority will be able to weed out coerced votes. A detailed comparison to Selene and JCJ is left to future work.

References

1. Adida, B.: Helios: web-based open-audit voting. In: Conference on Security Symposium, pp. 335–348. USENIX Association (2008)
2. Bayer, S., Groth, J.: Efficient zero-knowledge argument for correctness of a shuffle. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 263–280. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-29011-4_17](https://doi.org/10.1007/978-3-642-29011-4_17)
3. Blazy, O., Fuchsbauer, G., Pointcheval, D., Vergnaud, D.: Signatures on randomizable ciphertexts. In: Catalano, D., Fazio, N., Gennaro, R., Nicolosi, A. (eds.) PKC 2011. LNCS, vol. 6571, pp. 403–422. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-19379-8_25](https://doi.org/10.1007/978-3-642-19379-8_25)

4. Camenisch, J., Stadler, M.: Proof systems for general statements about discrete logarithms. Technical report, ETH Zurich (1997)
5. Chaidos, P., Cortier, V., Fuchsbauer, G., Galindo, D.: BeleniosRF: a non-interactive receipt-free electronic voting scheme. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS 2016), pp. 1614–1625. ACM, New York (2016)
6. Cramer, R., Gennaro, R., Schoenmakers, B.: A secure and optimally efficient multi-authority election scheme. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 103–118. Springer, Heidelberg (1997). doi:[10.1007/3-540-69053-0_9](https://doi.org/10.1007/3-540-69053-0_9)
7. Dodis, Y., Kiayias, A., Nicolosi, A., Shoup, V.: Anonymous identification in *Ad Hoc* groups. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 609–626. Springer, Heidelberg (2004). doi:[10.1007/978-3-540-24676-3_36](https://doi.org/10.1007/978-3-540-24676-3_36)
8. Fauzi, P., Lipmaa, H., Zajac, M.: A shuffle argument secure in the generic model. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016. LNCS, vol. 10032, pp. 841–872. Springer, Heidelberg (2016). doi:[10.1007/978-3-662-53890-6_28](https://doi.org/10.1007/978-3-662-53890-6_28)
9. Furukawa, J., Sako, K.: An efficient scheme for proving a shuffle. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 368–387. Springer, Heidelberg (2001). doi:[10.1007/3-540-44647-8_22](https://doi.org/10.1007/3-540-44647-8_22)
10. Groth, J.: A verifiable secret shuffle of homomorphic encryptions. *J. Cryptol.* **23**(4), 546–579 (2010)
11. Groth, J., Ishai, Y.: Sub-linear zero-knowledge argument for correctness of a shuffle. In: Smart, N. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 379–396. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-78967-3_22](https://doi.org/10.1007/978-3-540-78967-3_22)
12. Juels, A., Catalano, D., Jakobsson, M.: Coercion-resistant electronic elections. In: Workshop on Privacy in the Electronic Society, pp. 61–70. ACM (2005)
13. Liu, J.K., Wei, V.K., Wong, D.S.: Linkable spontaneous anonymous group signature for ad hoc groups. In: Wang, H., Pieprzyk, J., Varadharajan, V. (eds.) ACISP 2004. LNCS, vol. 3108, pp. 325–335. Springer, Heidelberg (2004). doi:[10.1007/978-3-540-27800-9_28](https://doi.org/10.1007/978-3-540-27800-9_28)
14. Myers, A.C., Clarkson, M., Chong, S.: Civitas: toward a secure voting system. In: Symposium on Security and Privacy, pp. 354–368. IEEE (2008)
15. Neff, C.A.: A verifiable secret shuffle and its application to E-voting. In: Proceedings of the 8th ACM Conference on Computer and Communications Security (CCS 2001), pp. 116–125. ACM, New York (2001)
16. Neff, C.A.: Verifiable mixing (shuffling) of elgamal pairs (2003). <http://www.votehere.org/vhti/documentation/egshuf.pdf>
17. Ramchen, K., Teague, V.: Parallel shuffling and its application to prêt à voter. In: Proceedings of the 2010 International Conference on Electronic Voting Technology/Workshop on Trustworthy Elections (EVT/WOTE 2010), pp. 1–8. USENIX Association, Berkeley (2010)
18. Rivest, R.L., Shamir, A., Tauman, Y.: How to leak a secret. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 552–565. Springer, Heidelberg (2001). doi:[10.1007/3-540-45682-1_32](https://doi.org/10.1007/3-540-45682-1_32)
19. Ryan, P.Y., Bismark, D., Heather, J., Schneider, S., Xia, Z.: The Prêt à Voter verifiable election system. *IEEE Trans. Inf. Forensics Secur.* **4**(4), 662–673 (2009)
20. 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). doi:[10.1007/978-3-662-53357-4_12](https://doi.org/10.1007/978-3-662-53357-4_12)

21. 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). doi:[10.1007/3-540-49264-X_32](https://doi.org/10.1007/3-540-49264-X_32)
22. Tsang, P.P., Wei, V.K.: Short linkable ring signatures for E-voting, E-cash and attestation. In: Deng, R.H., Bao, F., Pang, H.H., Zhou, J. (eds.) ISPEC 2005. LNCS, vol. 3439, pp. 48–60. Springer, Heidelberg (2005). doi:[10.1007/978-3-540-31979-5_5](https://doi.org/10.1007/978-3-540-31979-5_5)
23. Wikström, D.: A commitment-consistent proof of a shuffle. In: IACR Cryptology ePrint Archive: Report 2011/168 (2011)