



# Towards Everlasting Privacy and Efficient Coercion Resistance in Remote Electronic Voting

Panagiotis Grontas<sup>1(✉)</sup>, Aris Pagourtzis<sup>1</sup>, Alexandros Zacharakis<sup>1</sup>,  
and Bingsheng Zhang<sup>2</sup>

<sup>1</sup> School of Electrical and Computer Engineering,  
National Technical University of Athens, Athens, Greece  
{pgrontas, azach}@corelab.ntua.gr, pagour@cs.ntua.gr

<sup>2</sup> School of Computing and Communications, Lancaster University, Bailrigg, UK  
b.zhang2@lancaster.ac.uk

**Abstract.** In this work, we propose a first version of an e-voting scheme that achieves end-to-end verifiability, everlasting privacy and efficient coercion resistance in the JCJ setting. Everlasting privacy is achieved assuming an anonymous channel, without resorting to dedicated channels between the election authorities to exchange private data. In addition, the proposed scheme achieves coercion resistance under standard JCJ assumptions. As a core building block of our scheme, we also propose a new primitive called *publicly auditable conditional blind signature* (PACBS), where a client receives a token from the signing server after interaction; the token is a valid signature only if a certain condition holds and the validity of the signature can only be checked by a designated verifier. We utilize this primitive to blindly mark votes under coercion in an auditable manner.

**Keywords:** Electronic voting · End-to-end verifiability · Coercion resistance · Everlasting privacy · Publicly auditable conditional blind signatures

## 1 Introduction

The cryptographic research on electronic voting spans almost four decades. Despite the proliferation of proposed schemes, few have been implemented and used in actual elections. This can be attributed to the fact, that e-voting systems must reconcile conflicting properties with *integrity* and *privacy* being the most important ones. Integrity is usually achieved through verifiability, which can be individual, universal or administrative, allowing the voter, the public

---

This is work in progress; some properties rely on assumptions which should be lifted in order to lead to a fully functional practical solution.

B. Zhang was partially supported by EPSRC grant EP/P034578/1 and Petras PRF.

© International Financial Cryptography Association 2019

A. Zohar et al. (Eds.): FC 2018 Workshops, LNCS 10958, pp. 210–231, 2019.

[https://doi.org/10.1007/978-3-662-58820-8\\_15](https://doi.org/10.1007/978-3-662-58820-8_15)

or some trusted authorities, respectively, to check that all participants followed the protocol. Privacy protection comes in many layers: At the most basic level the secrecy of the vote is protected from the talliers. *Everlasting privacy* aims to protect against future and more powerful adversaries modelling the fact that theoretical and practical advances (e.g. quantum computing) might render obsolete the cryptographic assumptions upon which privacy rests. *Receipt Freeness* protects from a dishonest voter wanting to sell her vote to a passive adversary. *Coercion Resistance* is the essential property that will enable remote electronic voting, where the lack of a controlled environment for vote casting, leaves the voters vulnerable to active adversaries that can ‘look over their shoulder’ and dictate their behavior.

Juels, Catalano and Jakobsson proposed in [12] a framework to defend against coercion attacks, which was implemented in Civitas [19]. Their main idea was, that in order to achieve coercion resistance, the aspiring coercer must not be able to tell whether his attempt succeeded or not. This can be done by allowing the voter, to cast many ballots accompanied by anonymous credentials. Specifically, she obtains a valid credential through a one-time use of an untappable channel. Moreover she is assumed to have the capability to generate many different but indistinguishable ones. Under coercion, she uses fake (unregistered) credentials, indistinguishable from the valid one, which is employed when the voter has a moment of privacy, a necessary condition for coercion resistance. To correctly count the votes, however, the system must filter out the ballots that correspond to false credentials. This is done by comparing (in encrypted form) the supplied credentials with the valid ones that are published in a master list - the voter roll - after registration ends. Moreover, since the JCJ scheme allows for multiple votes per voter, duplicate ballots must be removed before counting. In principle, these operations are of quadratic complexity with respect to the number of ballots cast, which is typically quite larger than the number of voters, making the scheme impractical for large scale elections.

The motivation for this work stems from the reasonable assumption, first stated in [15], that the importance of integrity peaks during and immediately after the voting process, but diminishes as more parties are convinced about the result and concede. On the other hand, privacy is important during both voting and counting but retains its importance long after the announcement of the results, since the votes serve as evidence of one’s political beliefs. Verifiability implies that such evidence is available to many third parties, making it in effect undeletable. This can have dire consequences in the case of a future oppressive regime. Therefore, our focus is a voting protocol that enhances privacy, both during and after voting without sacrificing verifiability.

**Our Contribution.** We propose a first version of a voting protocol based on the architecture of FOO [4], one of the most privacy aware voting schemes in the literature, augmented with an efficient implementation of the coercion resistance properties of JCJ [12]. In particular, we take advantage of the fact that in FOO, voting occurs in two phases, namely authorization and counting, and

use it to overcome the performance bottleneck of JCJ. We achieve this by using the idea of [38], i.e. marking the fake credentials during the authorization phase where voter identification is available. By using the voter ID the correct credential can be efficiently retrieved and compared to the supplied one with no need to check all credentials. Of course, during this phase the ballot contents must be blinded, as they can be correlated with the voter ID. The fact that the credential is invalid is conveyed to the counting phase by applying a publicly auditable variation of a novel cryptographic primitive, *Conditional Blind Signatures* (CBS) [41]. The counter receives the ballot and an authorization in the form of a blind signature, that contains a bit that specifies if the vote is valid or under coercion. The perfect blindness property of the CBS scheme combined with an anonymous channel enable us to achieve the everlasting privacy property, without residing to dedicated channels between the authorities. Our protocol achieves verifiability, coercion resistance and everlasting privacy with minimal assumptions.

**Related Work.** Various efforts in the literature have tried to overcome the performance bottleneck of the JCJ scheme. In [13,17] linear complexity is achieved, by blinding the credentials and then stripping off the encryption randomization. As a result, they could be efficiently compared through a hashtable. Both schemes, however, were later found by [16] not to be coercion resistant by using a classic tagging attack. In the same paper, a new approach improves the quadratic complexity of identifying invalid credentials, by representing them as tuples with an underlying mathematical structure and not as mere random group elements. Hashing could then be used on part of the credential, without affecting the coercion resistance property. This approach was improved in [20,29,34] by utilizing different forms of credential structure and renewal methods, in order to enable use in multiple elections and minimize reliance on the untappable channel. In a different line of work, in [26] it is pointed out that the tagging attack is irrelevant in the duplicate removal subphase. As a result, the blind hashtables can be used there, thus achieving the goal of linearity in the number of votes. For linear fake removal the voter retrieves through the untappable channel, during registration, the index where her real credential is stored in the voter roll. Through this index, coerced credentials are identified. In an alternate approach, [23] and [27] employ the idea of anonymity sets. During vote casting the voter presents the real credential mixed with reencryptions of some other credentials from the voter roll. Finally, [37] with the Selene system and [32] with coercion evidence less strict definitions of coercion resistance are offered which might prove easier to reconcile with the other conflicting properties of voting systems. Our work utilizes ideas from all these works and integrates them in an efficient manner via a new cryptographic primitive, conditional blind signatures.

The term everlasting privacy was introduced by [15]. In [21], a protocol that uses perfectly hiding commitments and homomorphic encryption was proposed. Their main idea was that the public votes are protected perfectly using the commitment scheme, while the openings are computationally protected but are

exchanged through private channels. As a result they are not publicly available. This idea is presented as a primitive that can be integrated in any homomorphic tallying scheme in [31], while in [30] it is applied to mixnets providing everlasting privacy towards the public. This is further expanded and formalized as practical everlasting privacy in [28] by noting that such a future adversary might be more powerful in terms of computing power, but will have less information to operate on, since ephemeral data generated by the protocol will be unavailable in the long run. More recently, in [40], the authors implement the commitments with verifiable secret sharing and present a scheme that provides privacy and integrity against unbounded adversaries. However they assume untappable channels between the authorities and deny voters the ability to individually verify their votes. Our scheme differs from this line of work, since we do not assume or use specific channels between the authorities. All information is exchanged through the public Bulletin Board. This mode of communication is more realistic since a future regime with advanced cryptographic capabilities will have access to information exchanged by former governmental agencies in a private manner. Moreover our scheme also provides coercion resistance.

Coercion resistance combined with everlasting privacy seems to be an important desideratum in recent works. However this has not yet been possible in the JCJ framework, which is what our scheme accomplishes. In [39], a version of Selene enhanced for JCJ coercion resistance is equipped with everlasting privacy towards the public with the use of pseudonyms. However the creation process of pseudonyms and their relationship to real voter IDs and credentials requires trust assumptions and private channels between the members of the registration authority. Our work requires only the use of an anonymous channel and provides the same guarantees to both insiders and outsiders. In [36], everlasting privacy is achieved by using perfectly hiding commitments to registered identity credentials along with an anonymous channel. To achieve coercion resistance, votes can be overwritten and only the last one counts. As a result a voter under coercion can save her real vote for the end. This is a much stronger assumption than a simple moment of privacy required by our scheme and the JCJ framework; for example, an adversary who is able to cast a last minute vote achieves coercion.

## 2 Preliminaries

We begin by describing the agents, the functional components and the cryptographic primitives that make up our scheme.

- The main participants are naturally the  $n$  voters. In our protocol, like in [23], we assume that there exist pro democratic organizations that cast extra votes for registered voters in an effort to increase the size of the anonymity set.
- The *registration authority*  $RA$  registers the identities of the voters and provides credentials. We assume this occurs offline using an untappable channel.
- The *tallying authority*  $TA$  authorizes which ballots are accepted for counting by using a blind signature with an implicit validity bit. Later in the tallying phase it counts the valid ballots and announces the result.

In our scheme the registration authority and the tallying authority can be the same physical election authority *EA*. In reality they both consist of many members with conflicting interests for the election outcome. For clarity, however, we shall refer to them henceforth as if they consist only of a single member.

**Bulletin Board (BB).** A standard component of most electronic voting schemes. It is an authenticated broadcast channel with memory. It is meant to be implemented with Byzantine agreement algorithms. We do not provide implementation details here, following the vast majority of the electronic voting literature. We assume that whenever the voters use the BB, they are doing so through an *anonymous* channel that reveals no information about the identity of the sender of a message and that all the messages (requests, votes, proofs etc.) produced by our protocol can be found on the BB.

**Homomorphic Encryption Scheme.** We assume all cryptographic operations are performed by a JCJ compatible cryptosystem, i.e. one that supports reencryption and verifiable threshold decryption. In order to prove coercion resistance for our scheme we will use the Modified El Gamal (M-El Gamal) cryptosystem as presented in JCJ. It operates in a group  $\mathbb{G}$  of prime order  $q$  where the DDH Problem is hard. Two generators  $g_1, g_2$  are chosen randomly. The secret key is an element  $x \in \mathbb{Z}_q$  and the public key is  $h = g_1^x$ . Encryption<sup>1</sup> is performed as:  $E_h(m, r) = (g_1^r, g_2^r, mh^r)$  while decryption as:  $D_x(a, b, c) = c \cdot a^{-x}$ .

**Proofs of Knowledge.** We make extensive use of non-interactive zero knowledge (NIZK) proofs of knowledge. For instance a (M-El Gamal) ciphertext  $(a, b, c)$  is accompanied by proof that  $a, b$  have the same secret discrete logarithm relative to  $g_1, g_2$ . This can be implemented with the Chaum - Pedersen protocol [6] and made non interactive with the Fiat - Shamir heuristic [2]. We denote this by the functionality NIZK in the following way:  $\text{NIZK}\{(g_1, g_2, a, b), (x) : a = g_1^x \wedge b = g_2^x\}$ . We also use proofs that a value is a member of a set. We achieve this by using OR compositions of Chaum - Pedersen proofs as described in [7]. We also use proofs of knowledge of a discrete logarithm [3]. Finally designated verifier proofs [8], denoted as DVP, convince the voter but not the coercer that his credential was correctly encrypted, as in JCJ.

**Verifiable Shuffles.** We assume a functionality *Shuffle*, like the one proposed in [25], that takes as input a list of encrypted values and outputs a random permutation and reencryption of these values along with NIZK proofs that these operations were correctly performed. We use shuffles in tallying as in JCJ.

---

<sup>1</sup> For compactness we omit the encryption randomness, except when it is absolutely necessary for the operation of our scheme. We also use the plain ElGamal to describe the protocol and refer to M-El Gamal only in the coercion resistance analysis.

**Blind Signatures.** They allow a signer to sign messages without having access to their contents [1]. To this end the user *blinds* the message and the signer signs it in this blinded form. The user subsequently *unblinds* the signature, and retrieves a valid signature for the plain message. Their security properties [9, 24] are *blindness* or *unlinkability* which states that the signer cannot retrieve the signed message or associate signatures with protocol executions. *Unforgeability* states that the user cannot generate more message-signature pairs than those obtained by the signer. In the proposed protocol we use a variation of blind signatures, Conditional Blind Signatures, to enable everlasting privacy and move the marking of coerced votes from the tallying to the authorization phase.

**Plaintext Equivalence Test.** The functionality PET is a primitive introduced in [11] to convince a distributed set of entities, who share a decryption key that two ciphertexts indeed encrypt the same plaintext. It works by first having the participants blind the ciphertexts and then employing the homomorphic properties of the underlying cryptosystem to compute a function on them, such that a joint decryption of the result indicates if the two initial ciphertexts encrypt the same message or not. We use PET to mark duplicate votes and also embed them in the signatures that mark coerced votes in the authorisation phase.

### 3 Publicly Auditable Conditional Blind Signatures

Our voting scheme is built on a variation of Conditional Blind Signatures (CBS) [41]. This primitive allows a signer  $\mathcal{S}$  to blindly generate signatures on messages submitted by the user  $\mathcal{U}$ . These signatures however are verifiable only by a designated verifier  $\mathcal{V}$ , like in [8]. Furthermore, their validity depends on a secret information bit ‘injected’ into the signature along with the possession of a secret key by the designated verifier. In this way the signer can ‘instruct’ the verifier to accept the signature or not. The secret bit cannot be learned by the user however, since both cases are indistinguishable to her. Note that the roles of  $\mathcal{S}$  and  $\mathcal{V}$  can be played by the same entity, thus allowing the signer to send information regarding attributes of blinded messages to herself in the future.

The security of CBS extends the standard security properties of blind signatures such as blindness and protection against One More Forgery to account for the secret bit  $b$ . Additionally, to formally express the idea that  $b$  controls the validity of the signature an extra property, *Conditional Verifiability*, is defined in [41]. We must observe here that the user cannot validate the signature she receives, since she does not have knowledge of  $b$ . Although this seems counter-intuitive with respect to traditional signatures, in our setting it is essentially the exact property we need to achieve coercion resistance.

In [41] an instantiation of CBS is given by extending the well known three-round Okamoto-Schnorr blind signatures [5] (Appendix B). This instantiation is proved to have perfect blindness, computational resistance to *Strong* One More Forgery under the Computational Diffie Hellman assumption and Conditional Verifiability under the Decisional Diffie Hellman assumption.

In practice, the scheme’s round complexity can be reduced by randomly generating the initial commitment in a preagreed manner. Moreover it can also be combined with a multiplicatively homomorphic encryption scheme as the one we assume in our voting protocol. We present this modified version in Appendix C, where the signer and verifier are the same entity.

**Public Auditability.** The purpose of the CBS in the proposed voting scheme is to “mark” a ballot as valid if it is accompanied by an encryption of the same credential  $\sigma$  that was generated during registration. If any other credential  $\sigma'$  is used, the ballot is marked as invalid, indicating that the voter is under coercion. The CBS scheme can be used to convey this bit of information to the verifier, but by design it hides it from the user. As a result we cannot apply it as-is, since this will lead to loss of verifiability. We overcome this by introducing *Publicly Auditable* CBS, which adds auditability using NIZK proofs of correctness during signing and verification. In particular the CBS conditional bit is implicitly

**Common input:**  $\text{params}_{\text{CBS}}, \text{pk}_{\text{CBS}}, C_1, C_2 \in \mathbb{G}^2, \mathcal{H}_1 : \mathbb{G}^4 \rightarrow \mathbb{G}, \mathcal{H}_2 : \mathcal{M} \times \mathbb{G} \rightarrow \mathbb{Z}_q$   
 **$\mathcal{U}$ 's private input:**  $m \in \mathcal{M}$   
 **$\mathcal{S}$ 's private input:**  $s \in \mathbb{Z}_q$  s.t.  $k = g_1^s$

**$\mathcal{U}$  executes the Blind Algorithm:**

- Compute  $x := \mathcal{H}_1(C_1, C_2)$ ;
- Pick random  $u_1, u_2, d \leftarrow_R \mathbb{Z}_q$  and compute  $x^* := x g_1^{u_1} g_2^{u_2} v^d$ ,  $e^* := \mathcal{H}_2(m, x^*)$  and  $e := e^* - d$ ;
- Send  $e$  to  $\mathcal{S}$ .

**$\mathcal{S}$  executes the Sign Algorithm:**

- Compute  $x := \mathcal{H}_1(C_1, C_2)$ ;
- Pick random  $y_2 \leftarrow_R \mathbb{Z}_q$  as the second part of the CBS;
- Compute  $\nu := x g_2^{-y_2} v^{-e}$ ;
- Pick random  $t \in \mathbb{Z}_q$  and compute  $N := \mathbf{E}_h(\nu; t)$ ;
- Pick random blinding factor  $\alpha \in \mathbb{Z}_q$  and compute  $W := (C_2/C_1)^\alpha$  and apply signing key to compute  $B := (N \cdot W)^s$  with:

$$\begin{aligned} \pi_1 &\leftarrow \text{NIZK}\{(h_1, h, \nu, N), (t) : N = \mathbf{E}_h(\nu; t)\} \\ \pi_2 &\leftarrow \text{NIZK}\{(C_1, C_2, W), (\alpha) : W = (C_2/C_1)^\alpha\} \\ \pi_3 &\leftarrow \text{NIZK}\{(h, k, N, W, B), (s) : B = (N \cdot W)^s \wedge k = g_1^s\} \end{aligned}$$

- Set  $\text{bsig} := (B, N, W, y_2, \pi_1, \pi_2, \pi_3)$  and send  $\text{bsig}$  to the  $\mathcal{U}$ .

**$\mathcal{U}$  executes the Unblind Algorithm:**

- Verify  $\pi_1, \pi_2, \pi_3$ ;
- Unblind by computing  $\text{sig}_1 := B \cdot \mathbf{E}_h(k^{u_1})$  and  $\text{sig}_2 := y_2 + u_2$ .
- Set  $\text{sig} := (x^*, e^*, \text{sig}_1, \text{sig}_2)$  and output  $(m, \text{sig})$ .

**Fig. 1.** The publicly auditable CBS sign protocol PACBS Sign


computed and embedded in the signature by applying the PET functionality on the registered and voting credentials.

The PACBS scheme operates in a group  $\mathbb{G}$  of prime order  $q$ , where the DDH is hard. During the parameter generation phase random group elements  $(g_1, g_2, v, h_1)$  are selected. These elements are the public parameters of the protocol and are denoted as  $\mathbf{params}_{\text{CBS}}$ . A signing key  $s \in \mathbb{Z}_q$  and an encryption key  $z \in \mathbb{Z}_q$  are also selected. These secret keys are collectively denoted as  $\mathbf{sk}_{\text{CBS}}$ . The corresponding public keys are  $k := g_1^s$  and  $h := h_1^z$ , denoted as  $\mathbf{pk}_{\text{CBS}}$ .

The PACBS signing protocol in Fig. 1 assumes two random oracles  $\mathcal{H}_1, \mathcal{H}_2$ .

The signer obtains after registration an encryption of the valid credential  $C_1 := \mathbf{E}(\sigma)$ . During voting the voter provides an encryption of the voting credential  $C_2 := \mathbf{E}(\sigma')$  along with a blinded version  $e$  of the message being signed (the vote). The value  $(C_2/C_1)$  is blinded with a random  $\alpha \in \mathbb{Z}_q$  and multiplied with the signature. The essence of this procedure is that  $(C_2/C_1)$  is an encryption of a random element unless  $\sigma = \sigma'$  in which case it is an encryption of 1. In the former case the random element will ‘corrupt’ the signature. In the latter case the signature will be valid, since it is homomorphically multiplied by 1. Every interested entity can verify that the signer did not deviate from the protocol by checking the transcript and the proofs. Thus, an honest voter who knows the input and in particular whether  $C_1, C_2$  are encryptions of the same plaintext knows that his output corresponds to a valid signature.

The PACBS verification algorithm is given in Fig. 2. The verifier  $\mathcal{V}$ , given a message a signature and key pair, outputs whether the signature is valid or not



**Fig. 2.** The publicly auditable CBS verify protocol PACBS Verify



in a way that every other entity can be convinced about it. In reality,  $\mathcal{V}$  embeds the PET functionality inside the signature verification equation, which will again hold only if the credentials supplied are the same.

### 4 The Voting Protocol

Our scheme builds on the variation of FOO presented in [10] that is based on public key encryption instead of commitments, thus reducing the number of communication rounds. We use an extra authorization phase, where the issued credentials are secretly marked as valid or invalid using PACBS. Our main idea is that the validity checks for the credentials are (implicitly) done during vote authorization and not during tallying. The protocol has a one-time *Registration* phase and in each election there are three phases, namely *Authorization*, *Voting* and *Tallying*. To achieve the security properties we take advantage of the separation between Authorization and Voting phases. We stress that each phase starts only after the previous one has ended. A simplified view of the workflow of our protocol is depicted in Fig. 3. In the Authorization phase the EA checks for the validity of the supplied voter credentials. This can be done in constant time as the voter identity is known (but not the vote) and the EA has access to the voter roll. As a result, it compares the voter supplied credential  $E_h(\sigma')$  with the voter roll version  $E_h(\sigma)$  and finds out if the voter is under coercion. Since the identity of the voter is known at this point, the EA can use it to check eligibility by inspecting if there is a corresponding credential in the voter roll. Moreover it can be used to group all the ID and credential pairs so that only one is kept,




Fig. 3. Framework architecture and workflow

according to some predefined rule (e.g. last credential counts). Finally the voter and the *EA* interact according to the PACBS Sign protocol and obtain a validity token on the blinded ballot. While the honest voter knows that  $\sigma = \sigma'$  and can be sure that the signature will be valid, a coercer without this information cannot know if the ballot will be counted.

In the Voting phase, the voter casts the (unblinded) signature and the ballot to the BB. In the Tallying phase the *EA* must act as the verifier in PACBS and counts the votes only if the signature is valid. However, the ballot and result pairs must also be shuffled, so that the coercer loses track. Only then can they be decrypted to yield the final tally. This cannot be achieved directly by the PACBS Verify algorithm in Fig. 2 because shuffling must occur before decryption. Consequently the *EA*, actually uses a slight variation of PACBS Verify, which does not decrypt the final result nor create the proof of correct decryption, since these take place “outside” of the algorithm. We denote this alternate verification procedure EncVerify. In this stage neither the ID nor any credential information is present so the ballot cannot be linked to a voter.

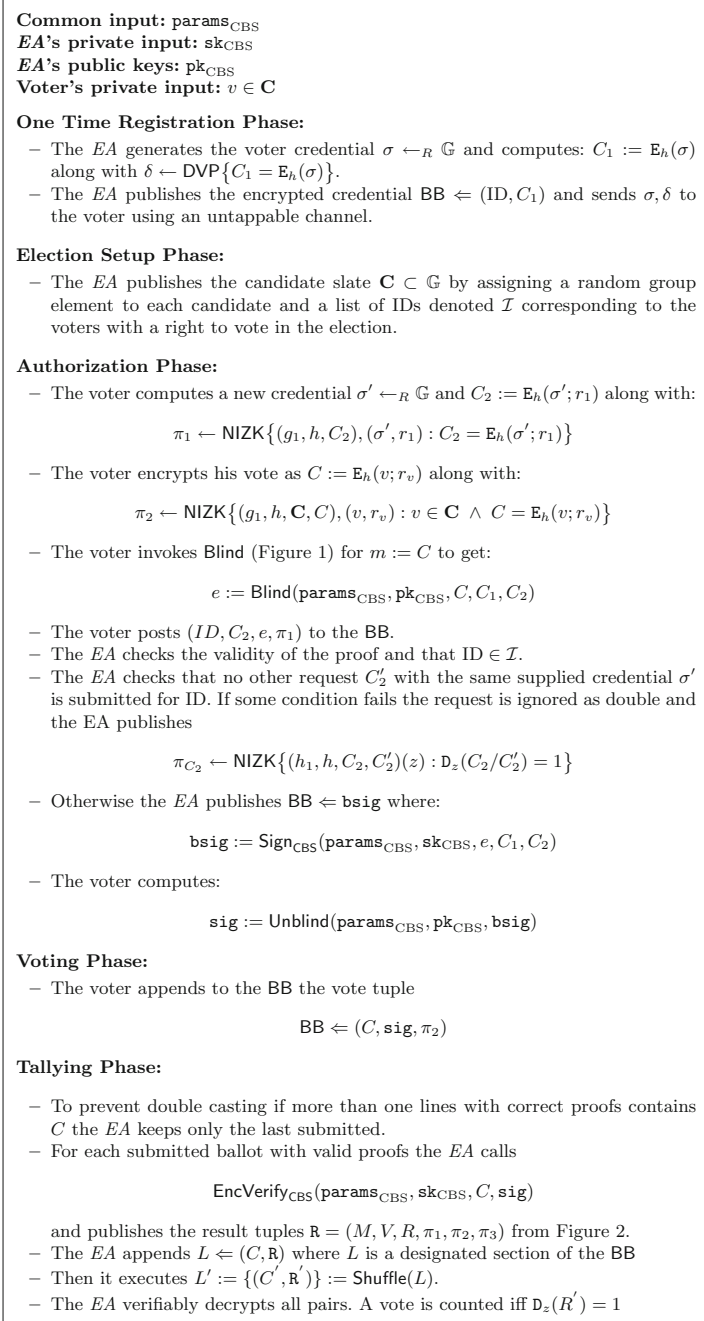
A detailed description is given in Fig. 4. Correctness follows by inspecting Figs. 1, 2 and 4. We assume that honest voters intentionally issue invalid votes during authorization to thwart forced abstention attacks.

**Distributed EA.** The *EA* can be modeled as a set of mutually distrustful parties executing secure protocols. In particular, the parameters for the protocol can be securely generated using standard techniques. The keys can be computed using a verifiable secret sharing scheme. The credentials can be generated as in [19]. Apart from the PACBS Sign and PACBS Verify all other actions performed by the *EA* (Decryption, Shuffle) are also standard and the checking for doubly issued credentials can be performed using PETs. The PACBS Sign and PACBS Verify can easily be extended by essentially performing the same protocols with each key share and combining the results.

**Performance.** Our analysis closely resembles [22]. Excluding the elimination of double votes all computations are linear in the number of votes. If  $|ID_i|$  denotes the number of votes cast with  $ID_i$  and  $m = \max_i |ID_i|$  then the number of computations is  $\mathcal{O}(m^2n)$ . This can be further reduced to  $\mathcal{O}(mn)$  using a method like the blind hashables of [17], since the tagging attack is not applicable in this phase. In any case, assuming that the number of duplicates per voter will be constant in practice - i.e.  $m = \mathcal{O}(1)$  - then the number of computations becomes linear in the number of voters  $n$ .

## 5 Security Analysis

**Threat Model.** Since our work is an extension of [12], our assumptions follow theirs closely. Firstly we require trusted implementation in software and hardware. While the amount of trust required can be decreased by using techniques



**Fig. 4.** The voting protocol

such as Benaloh challenges and code-voting as in [18,33,37], it cannot be completely disregarded. This is easier said than done, but it is a common practice in the vast majority of proposed voting protocols at our level of abstraction.

We assume two types of adversaries, one computationally bounded that acts during or shortly after the election and one that is computationally unbounded and acts in the future. The former models the security requirements which are vital during the election such as integrity, verifiability and coercion resistance while the latter models our requirement for everlasting privacy.

As far as present adversaries are concerned, we assume that they can perform only probabilistic polynomial time computations and for which our computational assumptions hold. To prove verifiability we assume that the adversary fully controls the election authorities and corrupts voters of his choice [33].

As far as coercion resistance is concerned, the adversary can fully control a subset of the voters by impersonating them, but there exists another subset with uncertain behavior. As in [12] each uncontrolled voter has a moment of privacy. The adversary can corrupt a subset of the voting authorities, which consist of mutually distrusting agents. Moreover he is capable of controlling the Bulletin Board and all other public channels, but there exist anonymous channels, where the identity of the sender of a message cannot be discovered. Finally there are honest participants, maybe nonprofit organizations, that cast invalid votes with valid voter IDs in order to thwart a forced abstention attack.

The future adversary is computationally unbounded and so can break any cryptographic assumption. Her goal is to gain information about the votes of a subset of the voters. We make the assumption that she can gain no information about the identity of the voter by the anonymous channel.

**Verifiability.** We follow the end-to-end verifiability definition proposed by Kiayias *et al.* [33], which can be viewed as a computational variant of the KTV framework as summarized in [35]. The adversarial goal against system’s integrity is to cause deviation from the intended tally of all the honest voters while election auditing remains successful without complains. We consider an adversary that controls the *EA* and a subset of the voters. All the voters who did not participate in the election are considered to be compromised. This is because a malicious (registration) authority can always impersonate absent voters without the PKI assumption.

Our scheme achieves end-to-end verifiability against a fully corrupted *EA* under the random oracle model. As a standard requirement, we assume the existence of a trusted *BB*. Although the voters’ clients are assumed to be honest for current protocol description, it is easy to add the Benaloh challenge mechanism [14] to prevent the malicious clients from tampering the ballot as the patch for Helios [18]. The voter needs to verify that her submitted ballot was recorded correctly on the *BB* and was taken as an input of the shuffle/mix-net.

During registration (Fig. 3 - step 2), the consistency between the voter’s credential  $\sigma$  and the published  $E_h(\sigma)$  is guaranteed by the DVP, which is intentionally not universally verifiable to enable coercion resistance. In the authorization

phase, the signatures for the validity of the credential (Fig. 3 - step 10) are verifiable due to the design of PACBS Sign and the proof published for each unprocessed request. More specifically, the *EA* shows that the produced signature is valid if and only if the submitted credential matches the recorded one. In the tallying phase, the public auditability property of the PACBS Verify protocol, the verifiable shuffle and the proof of correct decryption prevent the authority from deviating in any way from the protocol specification. Finally, everyone can check if the total number of valid signatures are less than or equal to the number of voters,  $n$ . This would prevent the malicious *EA* from inserting additional valid signatures. Since the honest voters' signatures are all cast, recorded, and tallied correctly, the rest valid signatures can be viewed as the adversarial ones. Hence, the malicious *EA* cannot add more votes even if she has the signing key.

**Eligibility.** The eligibility property is based on the resistance to Strong One More Forgery property of the signature scheme, since a valid signature is required for the vote to be counted. This implies the strong assumption that the adversary is restricted to a polylogarithmic number of honest authorizations.

**Privacy.** Our protocol satisfies vote privacy. In the authorization phase the encrypted vote is blinded when posted to the BB. As a result there is no way to recover the selection of the voter even if the *EA* is fully corrupted. In the voting phase the privacy of the vote depends on the privacy of the actions performed by the *EA* (Decryption, Shuffle). Without the assumption of an anonymous channel our system offers Helios [18] level privacy under similar trust assumptions. However, assuming an anonymous channel privacy protection becomes complete.

**Everlasting Privacy.** Our scheme easily meets the requirements for practical everlasting privacy set in [28], despite the fact that there are no private channels between the participants. A future adversary with access to the data in the BB, but without access to the untappable channels and to network related information will not be able to associate authorization requests and votes because of the blindness of the signatures. Moreover, the tallying phase where there are neither voter identities nor credentials present, matches the Helios without identities case of [28] which is proved to satisfy practical everlasting privacy. However if we assume an anonymous channel as in [4, 36] our scheme has complete everlasting privacy. In the authorization phase, the perfect blindness of the signature scheme ensures that no information regarding the vote is leaked. Furthermore, when the ballots are posted in the BB, despite being only computationally protected, they are cast through an anonymous channel and contain no information about the identity of the voter. Finally, the encrypted vote and signatures cannot be associated with any particular execution of the signing protocol that validated it. As a result a semi-honest unbounded adversary, watching all the public interactions cannot associate any voter with his vote.

**Coercion Resistance.** Our scheme is Coercion Resistant. In particular if a coercer requests a credential  $\sigma$  from a voter, its validity cannot be proved. As a result the validity of the signature issued for this credential is unknown to the coercer, due to the properties of PACBS. Moreover multiple votes cast with the same ID in the authorization phase, protect from a forced abstention attack. The reasoning is similar to [23, 27]. In the tallying phase, shuffling and PACBS Verify ensure the coercer loses track of his submitted vote and the only information he gets is the final tally. A detailed analysis is given in Appendix A.

## 6 Conclusion

In this paper we presented a new approach to provide coercion resistance in an efficient manner and combine it with everlasting privacy. Our protocol is based on minimal assumptions: a single use of an untappable channel and the existence of an anonymous channel. We utilized Conditional Blind Signatures [41], a recent primitive that allows a signer to inject a bit of secret information to a blind signature that controls if it should validate or not, which we improved for our purposes. Our scheme is proved secure under the JCJ [12] coercion resistance framework. The perfect blindness provided by CBS allows for stronger privacy guarantees; combined with a perfectly anonymous channel it provides the everlasting privacy property. In a future version of this work we plan to augment the intuitive security analysis presented here, using rigorous definitions and proofs.

**Acknowledgements.** The authors would like to thank Peter Browne Roenne and the anonymous reviewers for their helpful comments and suggestions.

## A Analysis of Coercion Resistance

We prove the coercion resistance property of the proposed voting scheme by closely following the JCJ techniques. We slightly modify the games  $c$ -resist and  $c$ -resist-ideal of JCJ to account for the extra authorization phase. We treat the auth functionality as a function that provides a valid or invalid ballot in relation to its private input.

Firstly, we examine the options of a coerced voter. Such a voter, can simply supply the adversary with a fake random element of  $\mathbb{G}$ . Having a fake credential, the signature she receives will be invalid, a fact undetectable by  $\mathcal{A}$  due to the design of the PACBS protocol. The voter can authenticate and cast her real vote during her moment of privacy. The coercer cannot decrypt the corresponding entry in the voter roll and find out if he was presented with the real one or not. If he tries to vote, when he receives the signature he will not be able to validate it himself. In the tallying phase the shuffle will make him lose track of the vote.

*The Games c-resist and c-resist-ideal.* The only change in the c-resist game of JCJ is the extra authorization phase. During this phase voters supply the essential information, based on their secret credential, and they get a ballot which can be either valid or invalid. We assume that all messages are exchanged via the bulletin board and thus are available to the adversary.

In the c-resist-ideal game an extra ideal functionality is needed, which we call *idauth* and its purpose is to authenticate and cast the ballots to the **BB**. This functionality gets the inputs of the players and computes the ballots as follows:

- It gives only one valid ballot for each valid credential.
- For honest voters it extracts the underlying credential and responds with a valid/invalid ballot based on it.
- Adversarial requests with credentials of corrupted players are answered normally. For the credential of the coerced voter the validity of the ballot computed is determined by the coin toss.
- The outputs are written directly to the **BB**.

In both games,  $\lambda$  is the security parameter,  $n$  is the number of voters and  $n_V$  is the number of voters the adversary can corrupt. Further inputs include the candidate slate **C** and the distribution  $D$  that governs the behaviour of the uncontrolled voters. In c-resist-game, the adversary corrupts a subset  $V$  of the voters and obtains their credentials after registration. We denote by  $U$  the rest of the voters. They perform the registration process and receive their credential. Subsequently the coercer chooses the voter  $j$  to coerce and her uncoerced vote  $\beta$ . Of course the voter should not be already corrupted and the vote should be a valid choice. A coin is flipped and if it is 0, the coerced voter fools  $\mathcal{A}$  by invoking the functionality *fakekey* and generating a fake credential while voting with her real one. If it is 1 she obeys and gives her real credential  $sk_j$  to  $\mathcal{A}$ . Honest voters perform the authorization phase and cast their votes, while  $\mathcal{A}$  invokes the *auth* functionality with any input of his choice using the information gathered so far.

The tallying phase is then performed and  $\mathcal{A}$  tries to guess the result of the coin, using the data on the **BB**, the final tally **X** and the proofs provided. In c-resist-ideal-game the same things happen with minor differences. The keys obtained by  $\mathcal{A}$  do not assist him with the choice of the voter to coerce. He is always given the real credential and *auth* is replaced with *idauth*. Finally, the only data  $\mathcal{A}$  can use to guess the coin flip is the final tally **X** and the number of invalid votes  $\Gamma$ .

We now present the simulation for the proof that our scheme is coercion resistant.

1. **Input:** The simulator  $\mathcal{S}$  takes as input the elements  $g_1, g_2, h_1, h_2$  of a group  $\mathbb{G}$  of order  $q$  and a vector  $w$  from a distribution  $D$ , which mirrors  $\mathcal{A}$ 's uncertainty. Each element of  $w$  is a set of valid and invalid votes, taking into account that each voter casts more than one ballot.  $\mathcal{S}$  tries to answer whether  $(g_1, g_2, h_1, h_2)$  is a DH quadruple or not.

2. **Parameter generation:** Initially the  $\mathcal{S}$  creates the M-El Gamal encryption key by randomly choosing  $x_1, x_2 \in \mathbb{Z}_q$  and computing  $h = g_1^{x_1} g_2^{x_2}$ . The public key is  $(g_1, g_2, h)$ . He then creates a signing key pair for the CBS scheme by choosing  $g_3, g_4, y \leftarrow_R \mathbb{G}$ ,  $s \in \mathbb{Z}_q$  and  $k = g_3^s$ . The secret key is  $s$  and the public key is  $(g_3, g_4, y, k)$ .
3. **Registration:** Each voter is assigned a random  $\sigma_i \leftarrow_R \mathbb{G}$ . Using the public key,  $\mathcal{S}$  publishes the voter roll. Finally, the candidate slate  $\mathbf{C}$  is published.
4. **Corruption:**  $\mathcal{A}$  corrupts voters.
5. **Coercion:**  $\mathcal{A}$  chooses the player to coerce and her honest vote  $(j, \beta)$ . The appropriate tests are performed in  $(j, \beta)$  according to the games' definitions.
6. **Coin Flip:**  $\mathcal{S}$  chooses  $b \leftarrow_R \{0, 1\}$ . If  $b = 0$ ,  $\mathcal{A}$  is given a random group element  $\sigma^* \leftarrow_R \mathbb{G}$ , else she is given the real credential  $\sigma^* \leftarrow \sigma_j$ .
7. **Authorization Requests:**  $\mathcal{S}$  issues the signature requests for the honest voters according to  $w$ . For each element of  $w$  she issues  $(E_h(\sigma_i), ID_i, \text{PoK}_1)$  where  $E_h(\sigma_i) = (h_1^{u_i}, h_2^{u_i}, h_1^{u_i x_1} h_2^{u_i x_2} \sigma_i)$  for random  $u_i$  and the proof  $\text{PoK}_1$  is simulated by the programmability of the random oracle by using standard techniques.  $\mathcal{A}$  issues his authorization requests.
8. **Double requests elimination:** Using the secret key  $x_1, x_2$ ,  $\mathcal{S}$  decrypts and eliminates double requests with the same credential.
9. **Authorization:**  $\mathcal{S}$  simulates this phase using her PACBS signing key. The messages are encrypted votes according to  $w$ . Encryptions are done in the same way as before.  $\mathcal{A}$  is given signatures in a straightforward manner.
10. **Vote Casting:**  $\mathcal{S}$  submits ballots for the honest voters.  $\mathcal{A}$  submits ballots for the corrupt and the coerced voters.
11. **Tallying:** Using his secret keys and standard techniques for proofs,  $\mathcal{S}$  simulates tallying in a straightforward manner.
12. **Guess:**  $\mathcal{A}$  decides  $b'$ .
13. **Output:**  $\mathcal{S}$  outputs 1 iff  $b = b'$ .

Let's examine the view of  $\mathcal{A}$ . Apart from the data he produces, in the authorization phase he sees the encrypted credentials with the proofs that accompany them, and the signatures given. These include a message  $x$  uniformly distributed in  $\mathbb{G}$ , an encrypted first part of a signature and the second part of the signature which is a uniformly distributed element in  $\mathbb{Z}_q$ . In the tallying phase he sees the encrypted ballots, their proofs and the signatures. The signatures include two random elements  $x^*, sig_2$  and an encrypted first part. Finally he gets the intermediate results and the tally with the proof. Apart from the encrypted messages and the proofs, all other data are random and do not assist him in deciding  $b$ .

Suppose that the input of  $\mathcal{S}$  is a Diffie-Hellman (DH) tuple. Then all the encryptions done by  $\mathcal{S}$  are valid and the view of  $\mathcal{A}$  is the same as the c-resist experiment. If the input is not a DH tuple then every encryption  $\mathcal{S}$  did results in uniformly distributed elements in  $\mathbb{G}^3$ .  $\mathcal{A}$ 's view is the same as in the c-resist-ideal experiment.



These imply that

$$\mathbf{Adv}_{ES,A}^{c\text{-resist}} = |\Pr[S = 1 | \text{DH}(g_1, g_2, h_1, h_2)] - \Pr[S = 1 | \neg \text{DH}(g_1, g_2, h_1, h_2)]|$$

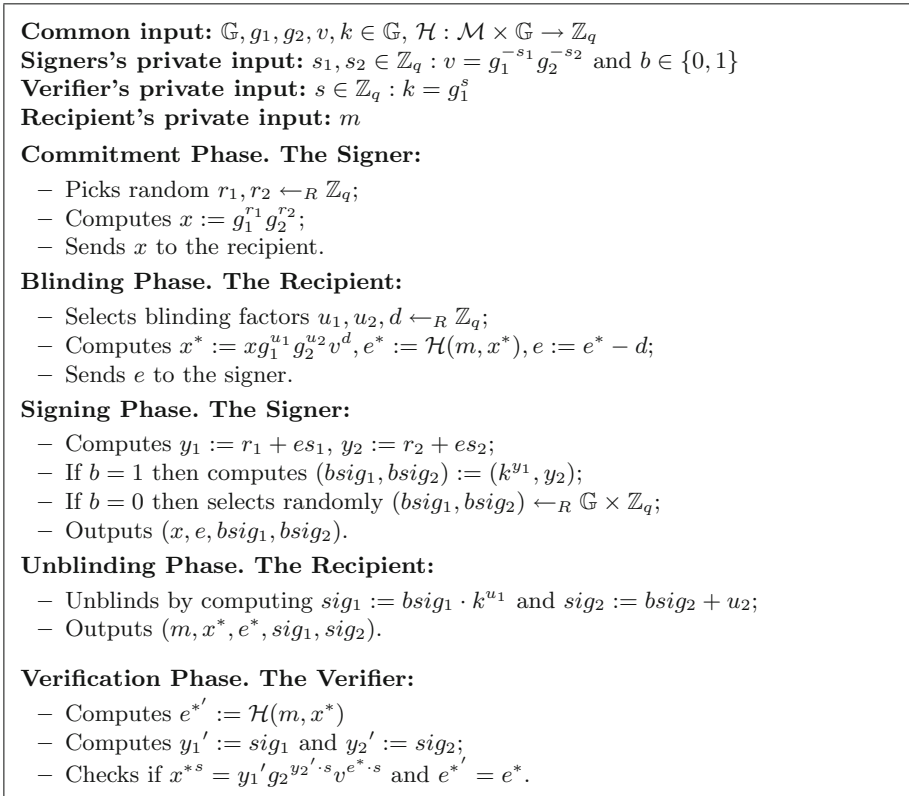
which is equal to  $\mathbf{Adv}_S^{DDH}$  and so it is negligible if the DDH assumption holds.

<b>Algorithm 1.</b> c-resist	<b>Algorithm 2.</b> c-resist-ideal
<b>Input</b> : $n, n_V, \mathbf{C}, D, \lambda$ <b>Output</b> : $result \in \{0, 1\}$ $(V, U) \leftarrow \mathcal{A}(\text{corrupt})$ $\{(sk_i, pk_i) \leftarrow \text{reg}(sk_R, ID_i, \lambda)\}_{i \in [n]}$ $(j, \beta) \leftarrow \mathcal{A}(\{sk_i\}_{i \in V}, \mathbf{Coerce})$ <b>if</b> $\beta \notin \mathbf{C}$ or $j \notin U$ <b>then</b>     output 0 <b>end</b> $b \leftarrow_R \{0, 1\}$ <b>if</b> $b = 0$ <b>then</b>   $sk^* \leftarrow \text{fakekey}(pk_T, sk_j, pk_j)$   $\text{BB} \leftarrow$   $\text{auth}(sk_j, pk_j, sk_T, pk_T, \mathbf{C}, \beta, \lambda)$ <b>else</b>   $sk^* \leftarrow sk_j$ <b>end</b> $\text{BB} \leftarrow$ $\text{auth}(sk_i, pk_i, sk_T, pk_T, \mathbf{C}, D, \lambda)_{i \in U \setminus \{j\}}$ $\text{BB} \leftarrow$ $\mathcal{A}^{\text{auth}(\cdot)}(\{sk_i\}_{i \in V}, sk^*, pk_T, \mathbf{C}, \text{BB})$ $(\mathbf{X}, P) \leftarrow$ $\text{tally}(sk_T, \text{BB}, \mathbf{C}, \{pk_i\}_{i \in V \cup U}, \lambda)$ $b' \leftarrow \mathcal{A}(\mathbf{X}, P, \text{BB}, \mathbf{Guess})$ output $b == b'$	<b>Input</b> : $n, n_V, \mathbf{C}, D, \lambda$ <b>Output</b> : $result \in \{0, 1\}$ $(V, U) \leftarrow \mathcal{A}(\text{corrupt})$ $\{(sk_i, pk_i) \leftarrow \text{reg}(sk_R, ID_i, \lambda)\}_{i \in [n]}$ $(j, \beta) \leftarrow \mathcal{A}(\mathbf{Coerce})$ <b>if</b> $\beta \notin \mathbf{C}$ or $j \notin U$ <b>then</b>     output 0 <b>end</b> $b \leftarrow_R \{0, 1\}$ $sk^* \leftarrow sk_j$ <b>if</b> $b = 0$ <b>then</b>   $\text{BB} \leftarrow$   $\text{idauth}(sk_j, pk_j, sk_T, pk_T, \mathbf{C}, \beta, \lambda)$ <b>end</b> $\{\text{BB} \leftarrow$ $\text{idauth}(sk_i, pk_i, sk_T, pk_T, \mathbf{C}, D, \lambda)\}_{i \in U \setminus \{j\}}$ $\text{BB} \leftarrow$ $\mathcal{A}^{\text{idauth}(\cdot)}(\{sk_i\}_{i \in V}, sk^*, pk_T, \mathbf{C})$ $(\mathbf{X}, P) \leftarrow$ $\text{tally}(sk_T, \text{BB}, \mathbf{C}, \{pk_i\}_{i \in V \cup U}, \lambda)$ $b' \leftarrow \mathcal{A}(\mathbf{X}, P, \Gamma, \mathbf{Guess})$ output $b == b'$

Finally, we must note that the exact level of protection each voter receives depends on the size of the anonymity set, i.e. the number of decoy votes cast with their ID by other honest voters or organizations. We plan to incorporate this analysis in future versions of our work.

## B Plain Okamoto-Schnorr CBS Scheme

We briefly present the simple Okamoto Schnorr CBS Scheme from [41]. The secret signing key consists of the values  $s_1, s_2 \in \mathbb{Z}_q$  as in [5] with corresponding public verification key  $v = g_1^{-s_1} g_2^{-s_2}$ . During the signing and unblinding phases the public key  $k$  of the verifier is used. For the verification algorithm, the verifier checks the verification equation using the hash of the message and the commitment using the secret key  $s \in \mathbb{Z}_q$ . If the secret signer bit is 1, then the signature will be valid, otherwise the verification equation will not hold. Thus the verifier will learn the secret bit of the signer. We also assume the existence of a random oracle  $\mathcal{H}$ .



**Fig. 5.** The original CBS Scheme based on Okamoto Schnorr signatures

## C Modified Okamoto-Schnorr CBS Scheme

The protocol in Fig. 5 can be combined with a multiplicatively homomorphic encryption scheme. It can also be made more practical if the parties agree in a common method to randomly generate the commitment message  $x$ . Moreover, we can let the signer play the role of the verifier, as a way to send the secret bit to oneself in the future.

We present this modified version in Fig. 6. Note that the unblinding of the first part of the signature, still occurs on the exponent, but this time in encrypted form.

**Common input:**  $\mathbb{G}, g_1, g_2, h_1, v, x \in \mathbb{G}$ , public keys  $k, h, \mathcal{H} : \mathcal{M} \times \mathbb{G} \rightarrow \mathbb{Z}_q$

**Signers's private input:**  $s \in \mathbb{Z}_q : k = g_1^s, z \in \mathbb{Z}_q : h = h_1^z$  and  $b \in \{0, 1\}$

**Recipient's private input:**  $m$

**Blinding:**

- $\mathcal{U}$  chooses random  $u_1, u_2, d \leftarrow_R \mathbb{Z}_q$
- $\mathcal{U}$  computes  $x^* := x g_1^{u_1} g_2^{u_2} v^d, e^* := \mathcal{H}(m, x^*)$  and  $e := e^* - d$ ;
- $\mathcal{U}$  sends  $e$  to  $\mathcal{S}$

**Signing:**

- $\mathcal{S}$  randomly selects  $bsig_2 \leftarrow_R \mathbb{Z}_q$ ;
- If  $b = 1$  then compute  $bsig_1 := \mathbf{E}_h((x \cdot g_2^{-bsig_2} \cdot v^{-e})^s)$ ;
- If  $b = 0$  then select randomly  $bsig_1 \leftarrow \mathbb{G}^2$ ;
- Output  $(x, e, bsig_1, bsig_2)$ .

**Unblinding:**

- $\mathcal{U}$  computes  $sig_1 := bsig_1 \cdot \mathbf{E}_h(k^{u_1})$  and  $sig_2 := bsig_2 + u_2$ ;
- Output  $(m, x^*, e^*, sig_1, sig_2)$ .

**Verification:**

- $\mathcal{V}$  recomputes  $e^*$  as  $\mathcal{H}(m, x^*)$ . If  $e^* \neq \mathcal{H}(m, x^*)$   $\mathcal{V}$  outputs 0.
- $\mathcal{V}$  decrypts the first part of the signature to receive  $y_1' := \mathbf{D}_z(sig_1)$
- $\mathcal{V}$  sets  $y_2' := sig_2$
- $\mathcal{V}$  checks if  $x^{*s} = y_1' g_2^{y_2' \cdot s} v^{e^* \cdot s}$  and outputs 1 iff it holds

**Fig. 6.** Modified conditional blind signatures

## References

1. Chaum, D.: Blind signatures for untraceable payments. In: Chaum, D., Rivest, R.L., Sherman, A.T. (eds.) *Advances in Cryptology*, pp. 199–203. Springer, Boston (1983). [https://doi.org/10.1007/978-1-4757-0602-4\\_18](https://doi.org/10.1007/978-1-4757-0602-4_18)
2. Fiat, A., Shamir, A.: How to prove yourself: practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) *CRYPTO 1986*. LNCS, vol. 263, pp. 186–194. Springer, Heidelberg (1987). [https://doi.org/10.1007/3-540-47721-7\\_12](https://doi.org/10.1007/3-540-47721-7_12)
3. Schnorr, C.P.: Efficient identification and signatures for smart cards. In: Brassard, G. (ed.) *CRYPTO 1989*. LNCS, vol. 435, pp. 239–252. Springer, New York (1990). [https://doi.org/10.1007/0-387-34805-0\\_22](https://doi.org/10.1007/0-387-34805-0_22)
4. Fujioka, A., Okamoto, T., Ohta, K.: A practical secret voting scheme for large scale elections. In: Seberry, J., Zheng, Y. (eds.) *AUSCRYPT 1992*. LNCS, vol. 718, pp. 244–251. Springer, Heidelberg (1993). [https://doi.org/10.1007/3-540-57220-1\\_66](https://doi.org/10.1007/3-540-57220-1_66)
5. Okamoto, T.: Provably secure and practical identification schemes and corresponding signature schemes. In: Brickell, E.F. (ed.) *CRYPTO 1992*. LNCS, vol. 740, pp. 31–53. Springer, Heidelberg (1993). [https://doi.org/10.1007/3-540-48071-4\\_3](https://doi.org/10.1007/3-540-48071-4_3)
6. 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](https://doi.org/10.1007/3-540-48071-4_7)

7. Cramer, R., Damgård, I., Schoenmakers, B.: Proofs of partial knowledge and simplified design of witness hiding protocols. In: Desmedt, Y.G. (ed.) CRYPTO 1994. LNCS, vol. 839, pp. 174–187. Springer, Heidelberg (1994). [https://doi.org/10.1007/3-540-48658-5\\_19](https://doi.org/10.1007/3-540-48658-5_19)
8. Jakobsson, M., Sako, K., Impagliazzo, R.: Designated verifier proofs and their applications. In: Maurer, U. (ed.) EUROCRYPT 1996. LNCS, vol. 1070, pp. 143–154. Springer, Heidelberg (1996). [https://doi.org/10.1007/3-540-68339-9\\_13](https://doi.org/10.1007/3-540-68339-9_13)
9. Juels, A., Luby, M., Ostrovsky, R.: Security of blind digital signatures. In: Kaliski, B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 150–164. Springer, Heidelberg (1997). <https://doi.org/10.1007/BFb0052233>
10. Ohkubo, M., Miura, F., Abe, M., Fujioka, A., Okamoto, T.: An improvement on a practical secret voting scheme. ISW 1999. LNCS, vol. 1729, pp. 225–234. Springer, Heidelberg (1999). [https://doi.org/10.1007/3-540-47790-X\\_19](https://doi.org/10.1007/3-540-47790-X_19)
11. 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](https://doi.org/10.1007/3-540-44448-3_13)
12. Juels, A., Catalano, D., Jakobsson, M.: Coercion-resistant electronic elections. In: Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society, pp. 61–70. ACM (2005)
13. Smith, W.D.: New cryptographic voting scheme with best-known theoretical properties. In: Frontiers in Electronic Elections (FEE 2005), June 2005
14. Benaloh, J.: Simple verifiable elections. In: EVT 2006 (2006)
15. 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](https://doi.org/10.1007/11818175_22)
16. Araújo, R., Foulle, S., Traoré, J.: A practical and secure coercion resistant scheme for remote elections. In: Frontiers of Electronic Voting (2007)
17. Weber, S.G., Araujo, R., Buchmann, J.: On coercion-resistant electronic elections with linear work. In: ARES, pp. 908–916. IEEE (2007)
18. Adida, B.: Helios: web-based open-audit voting. In: Proceedings of the 17th Conference on Security Symposium, pp. 335–348. USENIX Association (2008)
19. Clarkson, M.R., Chong, S., Myers, A.C.: Civitas: toward a secure voting system. In: IEEE Security and Privacy Symposium (2008)
20. Araújo, R., Ben Rajeb, N., Robbana, R., Traoré, J., Youssfi, S.: Towards practical and secure coercion-resistant electronic elections. In: Heng, S.-H., Wright, R.N., Goi, B.-M. (eds.) CANS 2010. LNCS, vol. 6467, pp. 278–297. Springer, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-17619-7\\_20](https://doi.org/10.1007/978-3-642-17619-7_20)
21. Moran, T., Naor, M.: Split-ballot voting: everlasting privacy with distributed trust. ACM Trans. Inf. Syst. Secur. **13**(2), 16 (2010)
22. 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](https://doi.org/10.1007/978-3-642-21424-0_10)
23. Schlöpfer, M., Haenni, R., Koenig, R., Spycher, O.: Efficient vote authorization in coercion-resistant internet voting. In: Kiayias, A., Lipmaa, H. (eds.) Vote-ID 2011. LNCS, vol. 7187, pp. 71–88. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-32747-6\\_5](https://doi.org/10.1007/978-3-642-32747-6_5)
24. Schröder, D., Unruh, D.: Security of blind signatures revisited. IACR Cryptology ePrint Archive, p. 316 (2011)

25. 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). [https://doi.org/10.1007/978-3-642-29011-4\\_17](https://doi.org/10.1007/978-3-642-29011-4_17)
26. Spycher, O., Koenig, R., Haenni, R., Schläpfer, M.: A new approach towards coercion-resistant remote E-voting in linear time. In: Danezis, G. (ed.) FC 2011. LNCS, vol. 7035, pp. 182–189. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-27576-0\\_15](https://doi.org/10.1007/978-3-642-27576-0_15)
27. Clark, J., Hengartner, U.: Selections: internet voting with over-the-shoulder coercion-resistance. In: Danezis, G. (ed.) FC 2011. LNCS, vol. 7035, pp. 47–61. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-27576-0\\_4](https://doi.org/10.1007/978-3-642-27576-0_4)
28. 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](https://doi.org/10.1007/978-3-642-36830-1_2)
29. Araújo, R., Traoré, J.: A practical coercion resistant voting scheme revisited. In: Heather, J., Schneider, S., Teague, V. (eds.) Vote-ID 2013. LNCS, vol. 7985, pp. 193–209. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-39185-9\\_12](https://doi.org/10.1007/978-3-642-39185-9_12)
30. Buchmann, J., Demirel, D., van de Graaf, J.: Towards a publicly-verifiable mixnet providing everlasting privacy. In: Sadeghi, A.-R. (ed.) FC 2013. LNCS, vol. 7859, pp. 197–204. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-39884-1\\_16](https://doi.org/10.1007/978-3-642-39884-1_16)
31. 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](https://doi.org/10.1007/978-3-642-40203-6_27)
32. Grewal, G.S., Ryan, M.D., Bursuc, S., Ryan, P.Y.A.: Caveat coercitor: coercion-evidence in electronic voting. In: IEEE Security and Privacy Symposium. IEEE (2013)
33. Kiayias, A., Zacharias, T., Zhang, B.: End-to-end verifiable elections in the standard model. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 468–498. Springer, Heidelberg (2015). [https://doi.org/10.1007/978-3-662-46803-6\\_16](https://doi.org/10.1007/978-3-662-46803-6_16)
34. Araújo, R., Barki, A., Brunet, S., Traoré, J.: Remote electronic voting can be efficient, verifiable and coercion-resistant. In: Clark, J., Meiklejohn, S., Ryan, P.Y.A., Wallach, D., Brenner, M., Rohloff, K. (eds.) FC 2016. LNCS, vol. 9604, pp. 224–232. Springer, Heidelberg (2016). [https://doi.org/10.1007/978-3-662-53357-4\\_15](https://doi.org/10.1007/978-3-662-53357-4_15)
35. Cortier, V., Galindo, D., Kuesters, R., Mueller, J., Truderung, T.: SoK: verifiability notions for e-voting protocols. In: IEEE Security and Privacy Symposium, pp. 779–798 (2016)
36. 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](https://doi.org/10.1007/978-3-662-53357-4_11)
37. 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](https://doi.org/10.1007/978-3-662-53357-4_12)
38. Grontas, P., Pagourtzis, A., Zacharakis, A.: Coercion resistance in a practical secret voting scheme for large scale elections. In: ISPAN-FCST-ISCC 2017, pp. 514–519 (2017)

39. Iovino, V., Rial, A., Rønne, P.B., Ryan, P.Y.A.: Using selene to verify your vote in JCJ. In: Brenner, M., et al. (eds.) FC 2017. LNCS, vol. 10323, pp. 385–403. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-70278-0\\_24](https://doi.org/10.1007/978-3-319-70278-0_24)
40. Yang, N., Clark, J.: Practical governmental voting with unconditional integrity and privacy. In: Brenner, M., et al. (eds.) FC 2017. LNCS, vol. 10323, pp. 434–449. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-70278-0\\_27](https://doi.org/10.1007/978-3-319-70278-0_27)
41. Zacharakis, A., Grontas, P., Pagourtzis, A.: Conditional blind signatures. In: 7th International Conference on Algebraic Informatics (Short Version) (2017). <http://eprint.iacr.org/2017/682>