

(Universal) Unconditional Verifiability in E-Voting without Trusted Parties

Gina Gallegos-Garcia¹, Vincenzo Iovino², Alfredo Rial³, Peter B. Rønne², and
Peter Y. A. Ryan²

¹ Instituto Politecnico Nacional, Mexico

`ggallegosg@ipn.mx`

² University of Luxembourg

`vinciovino@gmail.com`, `{alfredo.rial@uni.lu, peter.roenne,
peter.ryan}@uni.lu`

Abstract. In e-voting protocol design, cryptographers must balance usability and strong security guarantees, such as privacy and verifiability. In traditional e-voting protocols, privacy is often provided by a trusted authority that learns the votes and computes the tally. Some protocols replace the trusted authority by a set of authorities, and privacy is guaranteed if less than a threshold number of authorities are corrupt. For verifiability, stronger security guarantees are demanded. Typically, corrupt authorities that try to fake the result of the tally must always be detected.

To provide verifiability, many e-voting protocols use Non-Interactive Zero-Knowledge Proofs (NIZKs). Thanks to their non-interactive nature, NIZKs allow anybody, including third parties that do not participate in the protocol, to verify the correctness of the tally. Therefore, NIZKs can be used to obtain universal verifiability. Additionally, NIZKs also improve usability because they allow voters to cast a vote using a non-interactive protocol.

The disadvantage of NIZKs is that their security is based on setup assumptions such as the common reference string (CRS) or the random oracle (RO) model. The former requires a trusted party for the generation of a common reference string. The latter, though a popular methodology for designing secure protocols, has been shown to be unsound.

In this paper, we address the design of an e-voting protocol that provides verifiability without any trust assumptions, where verifiability here is meant without eligibility verification. We show that Non-Interactive Witness-Indistinguishable proofs (NIWI) can be used for this purpose. The e-voting scheme is private under the Decision Linear assumption, while verifiability holds unconditionally. To our knowledge, this is the first private e-voting scheme with *perfect* universal verifiability, i.e. one in which the probability of a fake tally not being detected is 0, and *non-interactive* protocols that do not rely on trust assumptions.

Keywords: e-voting, verifiability, witness indistinguishability, bilinear maps.

Table of Contents

1	Introduction.....	3
1.1	Background and Statement of the Problem	3
1.2	Our Results	5
2	Our Model and Definitions	6
2.1	Privacy	7
2.2	Verifiability.....	8
2.3	On the Need of a Stronger Correctness Property	9
3	Warm-up: Our Weakly Verifiable eVote (Sketch)	9
3.1	Intuition	10
3.2	Weak Verifiability of the Construction	11
3.3	Weak Privacy of the Construction	12
4	Our Fully Verifiable eVote (Sketch).....	14
4.1	Sketch of the Construction	14
4.2	(Full) Verifiability of the Construction	15
4.3	(Full) Privacy of the Construction	16
5	eVote with Multiple Authorities and Threshold Privacy	18
5.1	Sketch of the Construction	19
5.2	Verifiability of the Construction	19
5.3	Privacy of the Construction.....	20
6	On The Reusability of the Public Parameters.....	21
7	Related Work	22
8	Definitions	25
8.1	E-Voting Schemes	25
	Correctness and verifiability.....	27
	Privacy.....	28
9	Building Blocks	30
10	Our Weakly Verifiable eVote	33
10.1	Correctness and Weak Verifiability of the Construction.....	35
10.2	Weak Privacy of the Construction	37
11	Our (Fully) Verifiable eVote	39
11.1	Correctness and (Full) Verifiability of the Construction.....	43
11.2	Privacy of the Construction.....	45
12	Future Directions	58

1 Introduction

1.1 Background and Statement of the Problem

The parties participating in a standard e-voting protocol are multiple voters and one authority. First, the authority sets up a public key. A voter computes a ballot on input her intended vote and sends the ballot to a write-only public bulletin board (PBB), which records it in an entry associated with that voter. In case of abstention, a special symbol \perp is recorded on the PBB. The authority uses its secret key to compute the tally on input all the ballots on the PBB, which could possibly be \perp in case of abstention. Finally, the correctness of the tally can be checked by running a verification algorithm.¹

E-voting protocols must provide two security properties: privacy and verifiability. Privacy should protect the secrecy of the votes. Verifiability should prevent a corrupt authority from faking the tally. We will provide a formal definition of verifiability that is stronger than previous ones in some respects.

Privacy protection assumes the existence of a trusted authority in many e-voting systems [Cha81,CGS97,DJ01,CCC⁺09,Adi08,RS06,RT09,JCJ10]. As for schemes that distribute the trust among several authorities, privacy protection still requires that not all of the authorities are corrupt. Nevertheless, *verifiability* (also called integrity) should be guaranteed even if the authorities are corrupt.

Many e-voting systems make use of Non-Interactive Zero-Knowledge Proofs (NIZK) [BFM88,DMP88,Gol01,RS92,DDO⁺01] to provide verifiability. NIZK must provide two properties: soundness and zero-knowledge. Soundness prevents a corrupt prover from proving a false statement, i.e., a statement for which no witness exists. Zero-knowledge ensures that the verifier does not learn any information about the witness.

Zero-knowledge is defined following the simulation paradigm, i.e., it requires the existence of a simulator that computes a valid proof without knowledge of the witness. However, if such a simulator existed, soundness would not hold. This apparent contradiction is solved by resorting to trust assumptions like the Common Reference String (CRS) model [BFM88]. In the CRS model, a trusted party generates a CRS that is used by both provers and verifiers. The simulator is given the additional power of computing the CRS. Thanks to that, the simulator knows trapdoor information that allows it to simulate proofs for all statements.

For some applications of NIZK, the CRS model is not problematic. For instance, in IND-CCA public key encryption schemes [NY90,DDO⁺01,CS03b], zero-knowledge does not need to hold for the receiver of ciphertexts because the receiver must be able to decrypt anyway. Therefore, the CRS is computed by the receiver, while the NIZK proofs are computed by the sender of ciphertexts. However, in e-voting, the authority cannot compute the CRS because it must compute proofs that show the correctness of the tally.

An alternative to the CRS model is the Random Oracle (RO) model [BR93]. The RO model assumes the availability of a perfect random function available

¹ In this description we skipped some details (e.g., eligibility and authentication) that are not relevant to our setting. See below for more discussion.

to all parties. NIZKs that use the RO model are constructed following the Fiat-Shamir heuristic [FS87]. To prove that a NIZK proof constructed following this heuristic is zero-knowledge, we need programmability of the RO, i.e., the ability of the simulator to change the input/output of the RO.

To compute a proof, in practice, the prover replaces the RO by some “secure” hash function. Therefore, this hash function must be chosen honestly for zero-knowledge to hold. Consequently, all the parties must trust the implementation of a concrete hash function (e.g., SHA-3 [BDPA11]). We note that a hash function could have been designed in a malicious way (e.g., “programmed” like in the simulation) to allow the computation of a proof for a false statement. Currently, this needed trust on the implementation of hash functions does not exist. In fact, different political entities have developed their own hash functions because they do not trust the hash functions designed by others. For instance, the Russian government discourages the use of SHA-3 and encourages the use of its own hash function [Fed12].

Moreover, even when programmability is not needed, the RO methodology has been shown to be unsound [CGH98]. Further problems are known regarding the programmability of the RO in the context of NIZK [GK03,Kal06,BDSG⁺13]. The current techniques to avoid the need of programmability resort to the CRS model [DFN06,Lin15,CG15,CPSV16].

This motivates our main question: is it possible to design an e-voting scheme that is verifiable without assuming any trust assumption (like CRS and RO)?

In a survey [Lip05], Lipmaa asks whether Non-Interactive Witness Indistinguishable Proofs (NIWI) can be used to replace NIZKs. NIWIs can be constructed without using any trust assumptions [GOS06,DN00,BOV03,BP15].²

NIWI is a non-interactive proof/argument system that provides weaker security guarantees in comparison to NIZKs. While NIZKs ensure that a proof does not reveal any information about the witness, NIWIs only guarantee that, for any two witnesses w_1 and w_2 for the same statement, a proof computed with w_1 is computationally indistinguishable from a proof computed with w_2 . Note that this notion only makes sense for languages with multiple witnesses for each statement, which is not always the case.

To our knowledge, it was not known how to use NIWI to construct an e-voting scheme (eVote, in short) that is both private and verifiable. Usually, it is very difficult to use NIWI because of its weaker security guarantee. Nonetheless, inspired by a recent result on functional encryption [BSW11,GGH⁺13] of Badrinarayanan, Goyal, Jain and Sahai [BGJS16], we are surprisingly able to profitably use NIWI to answer our main question affirmatively.

² Note that in the literature there exist both NIWIs in the CRS model [GS08] and one-message NIWIs without CRS (see the citations above). Henceforth, unless specified otherwise, we refer only to one-message NIWIs without CRS, and in particular to the NIWIs for *CircuitSat* of Groth *et al.* [GOS06].

1.2 Our Results

First, we define correctness, privacy and verifiability properties for an eVote. We define two flavors of privacy and verifiability: weak and full. We propose an eVote that is (fully) private and (fully) verifiable. Its privacy can be reduced to the Decision Linear assumption [BBS04]. Its verifiability is *perfect* (see below) and thus is not based on any assumption. Moreover, its verifiability is *universal*, i.e., even a third party who did not participate in the election process should be able to verify the correctness of the tally. As a warm-up, we also describe an eVote that fulfills the weak privacy and weak verifiability properties.

Our eVote uses as building blocks a NIWI proof system, a public key encryption scheme with perfect correctness and unique secret key, and a perfectly binding commitment scheme. It can be instantiated by using just bilinear groups [Jou04,BF03]. For instance, we can instantiate our construction with the NIWI of Groth *et al.* [GOS06] and the Decision Linear encryption scheme of Boneh *et al.* [BBS04]. When instantiated with those building blocks, our construction is the first eVote with *non-interactive* algorithms for casting and verifying ballots and for computing and verifying the tally that provides *perfect* (weak and full) verifiability (as defined in Def. 3) and that fulfills the (weak and full) privacy property under the Decision Linear assumption [BBS04]. The Decision Linear assumption is a well-studied assumption over bilinear groups. Our construction attains *universal* verifiability, i.e., even third parties who did not participate in the election process are able to verify the tally.

We prove that our weakly verifiable eVote fulfills the weak verifiability and weak privacy properties in Corollary 3, and we prove that our (fully) verifiable eVote fulfills the (full) verifiability and (full) privacy properties in Corollary 6. We remark that the computational assumption is only needed to prove that our eVotes fulfill the (weak or full) privacy properties. In contrast, no assumption at all is necessary to prove that they fulfill the (weak or full) verifiability properties.

Therefore, our eVote with non-interactive algorithms is the first eVote whose perfect verifiability is not based on any trust and that is provably secure under a well-studied and falsifiable assumption [Nao03]. The latter is a key point of our results because otherwise one could just claim that an eVote in the RO model is secure when instantiated with any hash function. However, even when using such “unfalsifiable” assumptions, *perfect* verifiability cannot be achieved against unbounded adversaries because, in practice, for any hash function whose domain is larger than the range, the probability of finding a collision is not 0.

In Section 5, we outline how to adapt our (fully) verifiable construction to a model with multiple authorities. In this model, the tally evaluation algorithm is run by a set of authorities and the privacy property must hold if at least one authority is honest. (As this is not the main focus of our work, we do not present formal definitions and details for its construction.) An important advantage of our construction is that no interaction among the authorities is required. In this respect, our techniques completely diverge from previous approaches to the problem and may be of independent interest. We stress that the multi-string

model of Groth and Ostrovsky [GO14], though conceptually appealing in this scenario, fails to provide a solution.

In this work, we use cryptographic primitives to demonstrate the achievability of perfect verifiable systems. However, we are not concerned about usability and “human-friendly” verifiability, as dealt with in [Riv06,RR06,RRI16]. Furthermore, we only consider traditional e-voting systems and hence we neglect other approaches [KY02,DJ03,Gro04,HRZ10,KSRH12,GIR16].

Our privacy definition is inspired by the one of Benaloh [Ben87], also called “PRIV” in [BCG⁺15], which we reformulate by using modern terminology and we modify conveniently to withstand the attacks shown in [BCG⁺15]. We believe that our (fully) verifiable eVote can be proven secure according to other definitions of security, like for instance the one of Chase *et al.* [CKLM13], but we did not investigate the details because it is out of the scope of this initial work.

Organization. In Section 8 we present the definitions of an eVote, its verifiability and privacy; an overview of them is presented in Section 2. In Section 9 we present the building blocks we will use in our constructions. In Section 3 (resp. Section 4) we include all major details needed to understand our construction for a weakly verifiable eVote (resp. fully verifiable eVote) and its security properties. In Section 10 (resp. Section 11) we present the full details of our construction for a weakly verifiable eVote (resp. fully verifiable eVote) and of its security properties.

In Section 5, we outline how to adapt our (fully) verifiable eVote to a model with multiple authorities and threshold privacy. In this model, the tally evaluation algorithm is run by a set of authorities and privacy must hold if at least one of the authorities is honest.

In Section 6 we make some additional remarks about our definitions and in particular about the possibility of re-using the parameters through different elections. In Section 7 we discuss relevant related works.

Finally, in Section 12 we discuss some future directions in cryptography and e-voting that our work opens up.

2 Our Model and Definitions

In this section, we introduce our definitions of privacy and verifiability. We use a simple e-voting model with a single authority. We remark that, even for this model, it was not known how to avoid the use of CRSs or ROs. In Section 5, we outline how to adapt our constructions to a model with multiple authorities. Formal definitions of an eVote and of its privacy and verifiability are given in Section 8.1.

We use a general tally function $F : (\mathcal{M} \cup \{\perp\})^N \rightarrow \{0, 1\}^* \cup \{\perp\}$, where \mathcal{M} is the message space. The special symbol \perp denotes either an invalid vote or a blank ballot, when it is input to the function, or an error, when it is output by the function. A voter casts \perp to denote a blank ballot, i.e., a valid ballot where

no candidate is chosen. A voter abstains from voting by not casting any ballot or by casting an invalid ballot, which is replaced by \perp in the evaluation phase.

Our general tally function F must satisfy a very natural property given in Def. 2. The messages belong to a message space \mathcal{M} that is not specified. As byproduct, our constructions can be instantiated, for instance, to the case of a YES/NO election with the sum as tally function. As shown in [BCG⁺15], care has to be taken when considering general tally functions.

2.1 Privacy

Our *privacy* definition is indistinguishability-based and states that no PPT adversary can win the following game with non-negligible advantage. The adversary receives the public key generated by a challenger and chooses two tuples of strings that encode either valid votes in the message space $\mathcal{M} \cup \{\perp\}$ or arbitrary ballots, which are cast by possibly corrupt voters. We require that the tally function outputs the same result on input any of the tuples of strings.

The challenger chooses at random one of the two tuples. The challenger runs the ballot verification algorithm on input each of the arbitrary ballots and replaces the arbitrary ballot in the tuple by \perp if verification is unsuccessful. The challenger runs the cast algorithm on input each of the valid votes in the message space to compute a ballot and replaces the valid vote in the tuple by the ballot. Then the challenger computes the tally and a proof of correctness of the tally.

The new tuple, which replaces valid votes by ballots and invalid arbitrary votes by \perp , is given to the adversary along with a proof of the correctness of the tally. The adversary guesses which of the two tuples was chosen by the challenger.

More formally, the adversary sends two tuples $V_0 = (m_{0,1}, \dots, m_{0,N})$ and $V_1 = (m_{1,1}, \dots, m_{1,N})$ and a set $S \subset [N]$. The set S contains the indices of the strings of arbitrary ballots. For each $j \in S$, $m_{0,j} = m_{1,j}$ must hold. For each $j \notin S$, $m_{0,j}, m_{1,j} \in \mathcal{M} \cup \{\perp\}$ must hold. Moreover, we require that for all $d_1, \dots, d_N \in \mathcal{M} \cup \{\perp\}$, $F(m'_{0,1}, \dots, m'_{0,N}) = F(m'_{1,1}, \dots, m'_{1,N})$ must hold, where, for each $j \in S$, $m'_{0,i} = m'_{1,i} = d_i$ must hold, and for each $j \notin S$, $b \in \{0, 1\}$, $m'_{b,j} = m_{b,j}$ must hold.

Our definition can be viewed as a variant of Benaloh’s ballot privacy definition [Ben87] (also called “PRIV” in [BCG⁺15]) reformulated by using modern terminology and corrected to rule out some known attacks [BCG⁺15].

We also define *weak privacy*. The difference between the definitions of *weak privacy* and *privacy* is that, in *weak privacy*, the set S must be empty, i.e., the adversary cannot submit arbitrary ballots.

The privacy definitions that we use here are simple and do not capture vote replay attacks, see e.g. [CS10]. Such attacks are easily prevented by enforcing ballot independence. This can e.g. be done by appending a proof of knowledge of the plaintext in the ballots of the voters. Presently, this has not been done in the NIWI setting, so we will disregard this point for clarity. However, we stress that it is easy to change the schemes to satisfy full privacy definitions within the framework of having trust for privacy, but not for verifiability.

2.2 Verifiability

We define a ballot verification and a tally verification algorithm. In our definition of verifiability, we require two conditions to hold. The first condition states that, if each ballot and the proof of correctness of the tally issued by the authority are verified successfully by the respective algorithms, then each ballot C_i (possibly computed on input a maliciously generated public key) must be associated with a unique message $m_i \in \mathcal{M} \cup \{\perp\}$, and the result y claimed by the authority equals $F(m_1, \dots, m_n)$.

The second one requires that, even when the adversary generates the public key, if honest voters cast a ballot that is accepted by the ballot verification algorithm, then the ballot has to be “counted”. More concretely, consider that some ballots are computed by honest voters and are accepted by the ballot verification algorithm. (These ballots could be ill-formed if they are computed on input a public key generated by the adversary.) Consider also that the remaining ballots are computed by corrupt voters. In this situation, the tally evaluation algorithm outputs a tally y and a proof of correctness that, along with the public key and the ballots, is accepted by the tally verification algorithm. Then, it must be the case that the ballots sent by honest voters were counted to obtain y . For example, if the tally function is a sum function that sums binary votes and three honest voters cast three 1’s, then the authority should not be able to claim that $y < 3$.

Remarkably, our construction provides *perfect* verifiability. *Perfect* verifiability means that the probability that a malicious authority computes an incorrect tally and a proof that are accepted by the tally algorithm is *null*.

We also define *weak verifiability*. In *weak verifiability*, the authority can incorrectly claim that $y = \perp$. The second condition described above is still guaranteed for all the tallies $y \neq \perp$. For a weakly verifiable eVote, we only require weak privacy.

Although our weakly verifiable eVote satisfies weaker properties, it represents a worthwhile warm-up. Our (fully) verifiable eVote is based on it, though with some relevant modifications. Our weakly verifiable construction does not need a ballot verification algorithm, but for simplicity we use the same syntax for both the weakly verifiable and (fully) verifiable schemes. In Section 2.3, we describe a definition of correctness that is stronger than previous ones. This definition is needed to exclude the case that an eVote, that is intuitively not verifiable, fulfills formally the definition of verifiability.

In Section 5, we outline how to adapt our (fully) verifiable construction to a model with multiple authorities. In this model, the tally is computed by a set of authorities. Privacy must hold if at least one authority is honest. Because this model is not the main focus of our work, we do not present formal definitions or a detailed description of its construction. We note that such a construction would satisfy a different, but still without trust assumptions, definition of verifiability that essentially states that if there is at least one honest voter, the verifiability holds with overwhelming probability over the random coins of such voter, a very minimal assumption.

In this paper, for simplicity, we do not directly address issues of *eligibility*. We assume that a ballot is associated with a voter uniquely and that the adversary cannot submit a ballot on behalf of some voters. Unconditional eligibility verifiability seems hard or impossible to achieve, since we normally use some commitment, e.g. a PKI, and digital signatures on the ballots to prove the eligibility, but such an approach is not secure against an computationally unbounded adversary.

Our construction can however easily be extended to take into account such attacks by using digital signatures in a standard, but non-perfect, way, see e.g. [CGGI14]. The resulting construction would nonetheless satisfy a meaningful notion of verifiability secure against computationally bounded adversaries not based on any trust assumption and this constitutes an advancement to the state of the art in the field. In fact, to our knowledge, it is not even known how to construct an eVote protocol with computational verifiability without trusted parties.

2.3 On the Need of a Stronger Correctness Property

We justify here why a stronger correctness property is needed. Traditionally, the correctness property guarantees both that the ballot verification algorithm accepts the ballots computed by the cast algorithm, and (2) that the tally verification algorithm accepts the tally and the proof computed by the tally evaluation algorithm. In the latter, the ballots taken as input by the tally evaluation algorithm are computed by the cast algorithm. Therefore, it is not guaranteed that the tally verification algorithm accepts the output of the tally evaluation algorithm when the ballots are not computed by the cast algorithm.

We explain now that this is an issue. In our weakly verifiable scheme, the ballot verification algorithm accepts any ballot. Therefore, it would be possible to say that such a scheme is (fully) verifiable by just changing the tally verification algorithm so that it accept $y \stackrel{\Delta}{=} \perp$ only when all the ballots equal \perp . As can be seen, condition (1) in the definition of (full) verifiability (cf. Def. 3) is fulfilled because “if part” of the condition does not never holds.

However, intuitively, such a scheme is incorrect. Namely, if an honest authority that runs tally evaluation algorithm and gets $y = \perp$ (because some ballots where ill-formed), the tally verification algorithm should accept that result.

To address this issue, we add condition (2) to the definition of correctness (cf. Def. 3). This condition states that the tally verification algorithm must accept the output of the tally algorithm when run on input ballots that are accepted by the ballot verification algorithm (as opposed to ballots computed by the cast algorithm). We point out that in some works on definitional foundations (e.g., Bernhard *et al.* [BCG⁺15]) this issue has been overlooked.

3 Warm-up: Our Weakly Verifiable eVote (Sketch)

In this section, we sketch our construction for a weakly verifiable eVote, i.e. an eVote that fulfills the weak verifiability and weak privacy properties. A (fully)

verifiable eVote, which satisfies (full) verifiability and (full) privacy, is presented in Section 4. We stress that in practice such weakly verifiable eVote lacks fundamental security guarantees, but nonetheless it serves as a worthwhile warm-up to our fully verifiable eVote.

3.1 Intuition

Our weakly verifiable eVote uses 3 instances of a public key encryption (PKE) scheme in parallel. We require that the PKE scheme fulfills two properties: perfect correctness and unique secret key (see Def. 6). PKE schemes with those properties are known in the literature [DH76,BBS04] and can be constructed, e.g., from the Decision Linear assumption [BBS04]. The voter encrypts her vote 3 times using the PKE scheme *without* adding any proof of ciphertext well-formedness. Therefore, a ballot consists of three ciphertexts.

To compute the tally, the authority proceeds as follows. The authority decrypts the first ciphertext and the second ciphertext in a ballot. The authority replaces decrypted messages that do not belong to the message space by \perp .

The authority evaluates the tally function twice. First, the authority uses as input the messages encrypted in the first ciphertext of each ballot. Second, it uses the messages encrypted in the second ciphertext of each ballot. If both tallies are equal, the authority outputs the tally along with a proof of correctness, else the authority returns \perp to indicate an error.

The property of unique secret key guarantees that the decrypted message will be unique for each ciphertext. Without this property, it could be possible that a voter cast an invalid ciphertext Ct not belonging to the ciphertext space such that the decrypted message is different when using two well-formed secret keys Sk_1 and Sk_2 for the same public key. Note that this is not prevented by the correctness property, which only holds when the ciphertext is an output of the encryption algorithm.

Sketch of the construction. Let N be the number of voters and let F be a tally function with message space \mathcal{M} . The public key Pk of our eVote consists of the 3 PKs (Pk_1, \dots, Pk_3) of the underlying PKE. The secret key consists of the 3 corresponding SKs (Sk_1, \dots, Sk_3) of the PKE.

Our cast algorithm takes as input the public key (Pk_1, \dots, Pk_3) , the index j of the voter³, and for $j \in [N]$, a vote v . The cast algorithm outputs a ballot for the j -th voter. Our cast algorithm just encrypts the vote v with the 3 instances of the PKE to produce the ciphertexts Ct_1, \dots, Ct_3 . The ballot given as output is $Blt \triangleq (Ct_1, \dots, Ct_3)$.

³ The index is needed to associate a ballot with a unique voter. For instance, an eVote could require that each voter encrypts her ballot with a different PKE public key, adding a proof of well-formedness. The public key of the eVote would contain N PKE's public keys, one for each voter, and so the statement of the proof would have to contain the index of the voter in the set N .

The tally evaluation algorithm gives output y and works as follows. For all $j \in [N]$, if the corresponding voter cast her vote, for all $l \in [2]$, decrypt $\text{Ct}_{j,l}$ with Sk_l to get $m_{j,l}$. Then, for all $l \in [2]$ compute $y_l = F(m_{1,l}, \dots, m_{N,l})$, where for indices j such that either $m_{j,l} \notin \mathcal{M}$ or the j -th voter did not cast her vote, we set $m_{j,l} = \perp$. If the two y_l 's are equal to the *same* string y then return this as the tally, otherwise return an error $y = \perp$. Finally, compute a NIWI proof γ of the fact that $x = (\text{Bl}_1, \dots, \text{Bl}_N, \text{Pk}_1, \dots, \text{Pk}_3)$ satisfies the relation R^{dec} in Fig. 1 using as witness $(\text{Sk}_1, \text{Sk}_2, s_1, s_2)$. Another part of the witness is the two indices $i_1, i_2 \in [3]$, $i_1 < i_2$, which determine which two columns of ciphertexts that are used in the tally. In the real mode described above, we have $i_1 = 1, i_2 = 2$, but we can also have trapdoor modes with other index choices which will be essential for privacy.

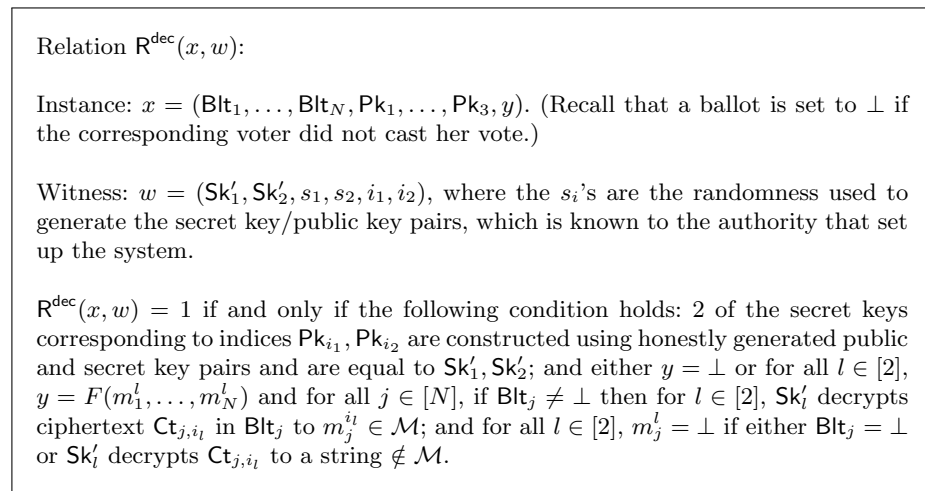


Fig. 1. Relation R^{dec} .

Note that the proof γ can be computed using as witness the randomness used to compute the public and secret key pairs. Finally, the algorithm outputs the pair (y, γ) . The tally verification algorithm verifies y by using the verification algorithm of the NIWI system.

3.2 Weak Verifiability of the Construction

Weak verifiability (cf. Def. 3) requires that, given a public key and a set of messages decrypted from the ballots, the authority cannot output a pair (y, γ) , $y \neq \perp$, such that y is an incorrect tally, but γ is accepted by the tally verification algorithm. However, the authority is able to claim that $y = \perp$ even if that is not the correct tally. The construction described above suffers from this problem.

We give a detailed proof that our construction fulfills the weak verifiability property in Theorem 1. In the following, we explain why our construction above fulfills the two conditions required by the weak verifiability property. First, we show that it fulfills the first condition. The first condition states that, if each ballot and the proof of correctness of the tally issued by the authority are verified successfully by the respective algorithms, then each ballot C_i (possibly computed on input a maliciously generated public key) must be associated with a unique message $m_i \in \mathcal{M} \cup \{\perp\}$, and the result y claimed by the authority equals $F(m_1, \dots, m_n)$.

We use a contradiction to show that our construction fulfills the first condition. Let us assume that there exist two results $y_0, y_1 \neq \perp$ such that $y_0 \neq y_1$, and two proofs γ_0, γ_1 that are accepted by the tally verification algorithm. By the unique secret key property, the decryption of the ciphertexts in the ballots produces a unique result. By the pigeon principle, there exists one index $i^* \in [3]$ used by both proofs. Therefore, it must be the case that either $y_1 = y_2 = \perp$ or y_1 and y_2 are equal to the evaluation of the tally function F on input the messages obtained by decrypting the ciphertexts. Consequently, $y_0, y_1 \neq \perp$ is a contradiction.

The second condition requires that, even when the adversary generates the public key, if honest voters cast a ballot that is accepted by the ballot verification algorithm, then the ballot has to be “counted”. We recall that an honest ballot for the j -th voter consists of three ciphertexts that encrypt the same message m . The perfect soundness of the NIWI ensures that the public key for the PKE scheme is honestly generated. The perfect correctness of the PKE scheme ensures that a ballot that encrypts m will be decrypted to m . Therefore, if the claimed tally y does not equal \perp , y has to be in the range of the function F restricted to m at index j .

3.3 Weak Privacy of the Construction

We explain how we prove that our construction fulfills the weak privacy property. The proof consists of a sequence of hybrid experiments [GM84], which are summarized in Table 3.3.

Table 1. Sequence of hybrid games to prove fulfillment of the weak privacy property.

Exp	$(Ct_{j,1}, Ct_{j,2}, Ct_{j,3})$	Sk index	γ	Security
H_1	$(m_{0,j}, m_{0,j}, m_{0,j})$	$(1,2,3)$	R	-
H_2	$(m_{0,j}, m_{0,j}, m_{1,j})$	$(1,2,3)$	R	IND-CPA
H_3	$(m_{0,j}, m_{0,j}, m_{1,j})$	$(1,2,3)$	T	WI
H_4	$(m_{0,j}, m_{1,j}, m_{1,j})$	$(1,2,3)$	T	IND-CPA
H_5	$(m_{0,j}, m_{1,j}, m_{1,j})$	$(1,2,3)$	T	WI
H_6	$(m_{1,j}, m_{1,j}, m_{1,j})$	$(1,2,3)$	T	IND-CPA
H_7	$(m_{1,j}, m_{1,j}, m_{1,j})$	$(1,2,3)$	R	WI

For simplicity, in this sketch we assume that the adversary submits a challenge that consists of two tuples $(m_{0,1}, \dots, m_{0,N})$ and $(m_{1,1}, \dots, m_{1,N})$, where each of the messages belongs to the message space \mathcal{M} . In the table, the first column shows the name of the hybrid experiment. The second column shows the three messages that are encrypted in the 3 ciphertexts $\text{Ct}_{j,1}, \dots, \text{Ct}_{j,3}$ contained in the challenge ballot $\text{Bl}_j = (\text{Ct}_{j,1}, \dots, \text{Ct}_{j,3})$ associated with voter j . The text in blue in the “Sk index” column denotes the indices used as witness in the proof γ . As mentioned above, if such blue indices correspond to the set $\{1, 2\}$ (resp. to a set different from $\{1, 2\}$) we say that the statement or proof is in real mode (resp. trapdoor mode), which we denote by R (resp. T) in the column γ . The text in red indicates the difference from the previous hybrid experiment.

The first four hybrid experiments and the last four are symmetrical. Therefore, it suffices to explain how we prove indistinguishability between the first four hybrid experiments.

Hybrid H_1 corresponds to the real experiment, except that the challenger sets the bit $b = 0$.

In hybrid H_2 , we switch the third message (in red) in any ballot to encrypt $m_{1,j}$. This is possible because the witness used to compute the proof γ does not contain the randomness used to compute the third secret key. Thanks to that, we can show indistinguishability between H_1 and H_2 by using the IND-CPA property of the PKE scheme.

In hybrid H_3 , the witness used to compute the proof γ contains the indices $\{1, 3\}$ instead of $\{1, 2\}$. Therefore, γ is in trapdoor mode. The witness-indistinguishability property of the NIWI allows us to show that H_3 cannot be distinguished from H_2 . Note that the result of the decryption does not change thanks to the constraint in the weak privacy definition that $F(m_{0,1}, \dots, m_{0,N}) = F(m_{1,1}, \dots, m_{1,N})$ must hold.

In hybrid H_4 , we switch the second message (in red) in any ballot to encrypt $m_{1,j}$. This is possible because the witness used to compute the proof γ does not contain the randomness used to compute the second secret key. Thanks to that, we can show indistinguishability between H_4 and H_3 by using the IND-CPA property of the PKE scheme.

We remark that, in order to switch the encrypted messages to $m_{1,j}$'s in every ballot, we use a simple property: the witness of the proof γ contains the randomness used to compute *two* of the secret keys. Thanks to that, we can show indistinguishability between H_1 and H_2 and between H_3 and H_4 by using the IND-CPA property of the PKE scheme whose randomness is not needed to compute γ . We point out that, to prove indistinguishability between those hybrid experiments, we need to use N “sub-hybrids”. In each “sub-hybrid”, we switch the message encrypted in just one ballot.

In Section 10, we present our weakly verifiable eVote in a more detailed manner.

4 Our Fully Verifiable eVote (Sketch)

The scheme sketched in Section 3 suffers from a severe problem: the authority can claim that the tally is \perp when it is not. That is, there can be two tallies $y_0 \neq \perp$ and $y_1 = \perp$ and two proofs γ_0 and γ_1 such that both proofs are accepted by the tally verification algorithm.

For instance, consider the following case. The ballots submitted by the voters are such that the tally y_1 obtained by evaluating the tally function on input the messages decrypted from the first ciphertext of each ballot equals the tally y_2 obtained when using the second ciphertext of each ballot, but differs from the tally y_3 obtained when using the third ciphertext. Then, by using the indices $(1, 2)$, the authority can prove successfully that the result of the election is $y_1 = y_2$, and by using indices $(1, 3)$, the authority can claim that the result of the election was \perp . The voters do not learn the indices that the authority used in the NIWI proof.

This also allows severe DoS attacks. For example, if just one voter submits a wrong ballot that makes the two tallies y_1 and y_2 be different from each other, then an honest authority has to output \perp . Furthermore, this scheme only fulfills the weak privacy property, which does not take into account corrupt voters.

Therefore, we propose a scheme that fulfills the (full) verifiability and the (full) privacy properties. This scheme solves the above-mentioned problems in an elegant way. Here we show a sketch of the scheme. In Section 11, we present our (fully) verifiable eVote in a more detailed manner.

4.1 Sketch of the Construction

In addition to the three public keys of the PKE scheme, the public key of the authority contains a perfectly binding commitment Z to the bit 1, i.e., the public key is $\text{Pk} = (\text{Pk}_1, \text{Pk}_2, \text{Pk}_3, Z)$, where $Z = \text{Com}(1)$.

A ballot consist of three ciphertexts, which are computed as in the weakly verifiable scheme, and of a proof that either the three ciphertexts encrypt the same message in the message space $\mathcal{M} \cup \{\perp\}$ or Z is a commitment to 0. Formally, the ballot contains a NIWI proof for the relation in Fig. 2.

The ballot verification algorithm runs the verification algorithm for the NIWI proof system for the relation $\mathbf{R}^{\text{enc,full}}$. (We recall that the ballot verification algorithm of our weakly verifiable eVote accepts any ballot.) The tally evaluation algorithm is the same as in our weakly verifiable eVote.

The tally verification algorithm also follows the one of the weakly verifiable eVote with the following modification. If either (1) not all inputs are \perp and $y = \perp$, or (2) all inputs are \perp and $y \neq \perp$, the tally verification algorithm outputs \perp .

We explain the reason for this modification. First, note that, in the (fully) verifiable scheme, the ballots that are rejected by the ballot verification algorithm are replaced by \perp as input to the tally evaluation algorithm. We recall that, in

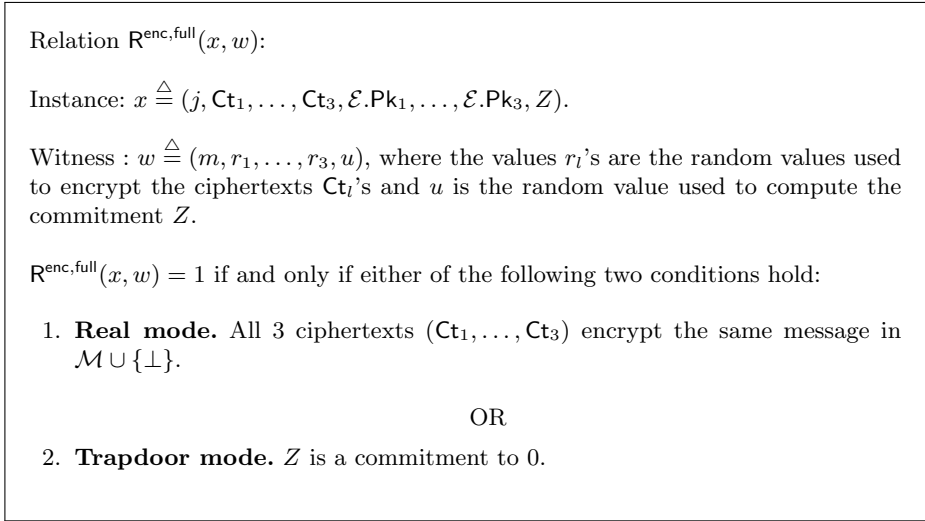


Fig. 2. Relation $\mathbf{R}^{\text{enc,full}}$.

the weakly verifiable scheme, the tally evaluation algorithm is run on input \perp only when voters do not send any ballot.

Our tally functions must fulfill a very natural property: $F(m_1, \dots, m_N) = \perp$ iff $m_1 = \perp, \dots, m_N = \perp$ (cf. Def. 2). That is, if at least one message is valid, then it has to be “counted”.

As we show below, if the public key is honestly generated, the tally evaluation algorithm never returns \perp on input a tuple of possibly dishonest ballots. Therefore, except for the case that all the ballots are invalid, a tally $y = \perp$ may only occur if the authority acted dishonestly and, consequently, the tally verification algorithm should not accept $y = \perp$.

4.2 (Full) Verifiability of the Construction

We show that our scheme fulfills the (full) verifiability property. This property consists of two conditions described in Section 2 and defined in Def. 3.

First, we show that our scheme fulfills the first condition. (A detailed proof is given in Theorem 4.) This condition requires that the authority cannot output two tallies y_1, y_2 such that $y_1 \neq y_2$ and two proofs γ_1 and γ_2 that are accepted by the tally verification algorithm.

Our tally verification algorithm only accepts a tally $y = \perp$ when all the ballots are invalid. Therefore, (1) the authority is not able to wrongly claim that a tally is \perp . Furthermore, as in our weakly verifiable eVote, (2) the authority cannot output two tallies y_1, y_2 such that $y_1 \neq y_2, y_1, y_2 \neq \perp$ along with proofs γ_1 and γ_2 that are accepted by the ballot verification algorithm. Therefore, (1) and (2) imply that the authority cannot output two tallies y_1, y_2 such that $y_1 \neq y_2$.

We show now that the second condition also holds. First, we note that the authority can only create a dishonest public key by setting the commitment dishonestly. The reason is that the authority has to prove that the public key of the PKE scheme is honestly generated. Therefore, the perfect correctness of the NIWI and of the PKE scheme guarantee that an honestly computed ballot⁴ for the j -th voter that encrypts message m will always be “counted”, i.e., for any (y, γ) pair that is accepted by the tally verification algorithm, y will be compatible with m at index j according to Def. 1.

4.3 (Full) Privacy of the Construction

We show now that our scheme fulfills the (full) privacy property. Here we summarize the proof. In Section 11.2, we describe the proof in detail. We stress that, for privacy to hold, the authority must be honest and thus the public key is honestly generated.

In the security proof, we consider a sequence of hybrid experiments. First, we define an experiment H^Z in which the commitment in the public key is a commitment to 0. We show that H^Z is indistinguishable from the real experiment under the computationally hiding property of the commitment.

Second, we define an event E^1 in experiment H^Z . In E^1 , the adversary submits a ballot that is accepted by the ballot verification algorithm but, when decrypting the three ciphertexts in the ballot, the three decrypted messages in $\mathcal{M} \cup \{\perp\}$ are not equal. We show that the probability of E^1 is negligible under the computationally hiding property of the commitment. More concretely, we show that, if E^1 occurs with non-negligible probability, then the adversary can be used to distinguish a commitment to 0 from a commitment to 1. We note that, if Z is a commitment to 1, the perfect soundness of the NIWI guarantees that the adversary can never submit an ill-formed ballot that is accepted by the ballot verification algorithm.

The next hybrid experiments are similar to the ones used in the security proof of the weakly verifiable scheme. Thanks to the hybrid experiment H^Z , we can still show indistinguishability between those hybrid experiments by using the IND-CPA property of the PKE scheme. The reason is that, thanks to H^Z , the NIWI proof in the ballots can be a proof that the commitment in the public key is a commitment to 0. Therefore, we avoid the computation of a proof that shows that the three ciphertexts encrypt the same message, which allows us to switch the message encrypted in one of the ciphertexts and prove indistinguishability by using the IND-CPA assumption.

⁴ Here, “honestly computed ballot” just means that it is computed by the voter using the Cast algorithm on input the public key of the authority, which could be honestly or dishonestly created. By design of our construction, an honestly generated ballot computed on input an honestly created public key has the same distribution of an honestly created ballot computed on input any possibly dishonest public key whenever the authority is able to compute proofs of tally correctness that are accepted by the tally verification algorithm.

There is one difference between the hybrid experiments in the weakly verifiable scheme and in the (full) verifiable scheme. Namely, in the (full) verifiable scheme, we have to handle possibly dishonest ballots. In particular, we have to guarantee that, when we switch the indices used as witness for the NIWI proof of tally correctness, the tally does not change. To illustrate this issue, suppose that, in an adversarial ballot, the first two ciphertexts encrypt the same message x but the third one encrypts a different message z . Then the tally computed by the secret keys for indices $\{1, 2\}$ could differ from the one computed with secret keys for indices $\{2, 3\}$. In that case, we cannot prove indistinguishability between a hybrid experiment where the NIWI witness comprises Sk_1, Sk_2 and a hybrid experiment where the NIWI witness comprises Sk_2, Sk_3 .

To solve this issue, we show that event E^1 occurs with negligible probability. Therefore, it is sufficient to analyze the advantage of the adversary in the hybrid experiments conditioned on the occurrence of \bar{E}^1 (i.e., the complement of E^1).

More concretely, the sequence of hybrid experiments after H^Z is as follows. We recall that the adversary sends two tuples $V_0 = (m_{0,1}, \dots, m_{0,N})$ and $V_1 = (m_{1,1}, \dots, m_{1,N})$, and a set $S \subset [N]$ that contains the indices of the strings of arbitrary ballots.

- Hybrid experiment H_1 is equal to the experiment H^Z , except that the challenger sets the bit $b \stackrel{\Delta}{=} 0$.
- Hybrid experiment H_2 switches the message encrypted in the third ciphertext in any ballot to encrypt $m_{1,j}$ instead of $m_{0,j}$. More in detail, for $k = 0$ to N , we define a sequence of hybrid experiments H_2^k . H_2^k is identical to H_1 , except that, for all $j = 1, \dots, k$ such that $j \notin S$, the challenger computes the third ciphertext of the ballot on input $m_{1,k}$. Therefore, H_2^0 is identical to H_1 , while H_2^N is identical to H_2 . We show that H_2^k and H_2^{k+1} are indistinguishable thanks to the IND-CPA property of the PKE scheme.
- Hybrid experiment H_3 is identical to experiment H_2 , except that the challenger computes the NIWI proof γ on input a witness that contains indices $(1, 3)$ and secret keys Sk_1, Sk_3 , instead of indices $(1, 2)$ and secret keys Sk_1, Sk_2 . We show that H_3 and H_2 are indistinguishable thanks to the witness-indistinguishability property of the NIWI proof. Because \bar{E}^1 occurs with overwhelming probability, any ballot in S is either replaced by \perp , if the ballot verification algorithm does not accept it, or decrypted to the same value in H_2 and H_3 . Therefore, the tally evaluation algorithm outputs the same tally in H_2 and H_3 .
- Hybrid experiment H_4 is identical to H_3 , except that the second ciphertext in any ballot encrypts $m_{1,j}$ instead of $m_{0,j}$. More in detail, for $k = 0$ to N , we define a sequence of hybrid experiments H_4^k . H_4^k is identical to H_3 , except that, for all $j = 1, \dots, k$ such that $j \notin S$, the challenger computes the second ciphertext of the ballot on input $m_{1,k}$. Therefore, H_4^0 is identical to H_3 , while H_4^N is identical to H_4 . We show that H_4^k and H_4^{k+1} are indistinguishable thanks to the IND-CPA property of the PKE scheme.

The remaining hybrid experiments are symmetrical to the ones described in above. In H_5 , the witness used to compute the NIWI proof contains the in-

dices $(2, 3)$ and secret keys Sk_2, Sk_3 , and indistinguishability between H_5 and H_4 follows from the witness-indistinguishability property of the NIWI proof. In H_6 , the first ciphertext of each ballot encrypts $m_{1,j}$ instead of $m_{0,j}$ and indistinguishability between H_6 and H_5 follows from the IND-CPA property of the PKE scheme. Finally, in H_7 the witness used to compute the NIWI proof contains the indices $(1, 2)$ and secret keys Sk_1, Sk_2 , and indistinguishability between H_7 and H_6 follows from the witness-indistinguishability property of the NIWI proof.

We would like to remark the subtle difference between ill-formed and invalid ballots. An ill-formed ballot is a ballot that is not in the range of the cast algorithm. However, an ill-formed ballot could be valid in the sense that, along with other (possibly ill- or well- formed) $N - 1$ ballots, the authority obtains a tally, i.e., the tally obtained when decrypting the first and the second ciphertext in the ballots is the same. An ill-formed ballot can be computed when the commitment in the public is computed dishonestly.

The event \bar{E}^1 may occur even if the adversary submits an ill-formed ballot that is accepted by the ballot verification algorithm. In fact, if a (non-honestly computed) ballot is formed by strings that are not in the ciphertext space of the encryption algorithm of the PKE, but decryption of those strings outputs the same message, such a ballot is not considered invalid.

Note also that the proof of well-formedness of the ballots states that the encrypted messages may be equal to \perp . Ballots that encrypt \perp are blank ballots. We consider tally functions in which the symbol \perp indicates a blank vote. For example, in case of an eVote for the sum function in which \perp is counted as 0, an adversary should not be able to distinguish three ballots that encrypt $(1, 1, \perp)$ from three ballots that encrypt $(1, \perp, 1)$.

5 eVote with Multiple Authorities and Threshold Privacy

In this section, we sketch how to generalize our (fully) verifiable construction to fit a model with multiple authorities. In this model, the tally evaluation algorithm is run by a set of authorities. The privacy property must hold if not all the authorities are corrupt. Our generalized scheme guarantees a statistical verifiability property (see below), which assumes that there is at least one honest voter.

First, we note that the multi-string model of Groth and Ostrovsky [GO14] does not provide a solution to this problem. The multi-string model assumes that the majority of the parties that set up the CRSs are honest. It does not guarantee soundness, which would provide verifiability in our application, when *all* those parties, which would be the authorities in our application, are corrupt. In the multi-string model, there is a trade-off between soundness and zero-knowledge. Namely, soundness could hold when all the authorities are corrupt, but then zero-knowledge does not hold. Zero-knowledge is guaranteed only when there is a majority of honest authorities. In contrast, our generalized scheme fulfills the privacy property when at least one authority is honest.

5.1 Sketch of the Construction

Our generalized construction works for tally functions that can be represented as polynomials. Such tally functions comprise many functions of interest for e-voting. For simplicity, henceforth we only consider the case of the sum function with a binary message space. The general case follows by using Lagrange’s polynomial interpolation.

Consider the sum function over a set of integers S^k , which we specify later. Consider m authorities. Each authority $k \in [m]$ publishes a public key that consists of the public key of our (fully) verifiable eVote and, in addition, a commitment com_j to a tuple of N 0’s for $j \in [m]$.

Each voter shares her vote v_j in m shares $v_{j,1}, \dots, v_{j,m}$ whose sum is v_j . (Later we describe how the shares are computed in order to preserve privacy.) For all $k \in [m]$, a voter encrypts $v_{j,k}$ under the public key of the k -th authority. As result, a voter obtains the ciphertexts $\text{Ct}_1, \dots, \text{Ct}_m$. In addition, the voter adds a proof that either (the real statement) for all $k \in [m]$, Ct_k encrypts a number in S^k such that the sum of the encrypted numbers is in $\{0, 1\}$ (for simplicity, here we do not consider messages equal to \perp) OR (the trapdoor statement) for all $k \in [m]$, com_k is a commitment to a tuple (z_1, \dots, z_N) such that z_j is equal to her ballot, i.e., to the tuple $(\text{Ct}_1, \dots, \text{Ct}_m)$.

For each $k \in [m]$, the k -th authority computes the tally y_k as in the (fully) verifiable scheme. The proof of correctness of the tally is a proof for the following modified relation: either (the real statement) the witness satisfies the relation $\text{R}^{\text{dec,full}}$ of the (fully) verifiable scheme and com_k is a commitment to 0 OR (the trapdoor statement) com_k is a commitment to a tuple (z_1, \dots, z_N) such that $z_j = \text{Bl}_j$ for all $j \in [N]$, where $\text{Bl}_1, \dots, \text{Bl}_N$ are the N ballots published on the public bulletin board.

Finally, the tally is computed by summing the tallies y_k ’s output by each of the authorities to obtain y . We give more details below.

To support functions represented as polynomials, the following modifications should be applied. To compute the shares $v_{j,1}, \dots, v_{j,m}$, the voter j chooses a polynomial p_j of degree $m - 1$ such that $p_j(0)$ equals her vote v_j . The shares are the evaluation of p_j on input $1, \dots, m$. The tally is computed by using Lagrange interpolation.

5.2 Verifiability of the Construction

We analyze now the verifiability of our generalized construction. If the commitments in the public key are computed honestly, we can show that the generalized construction fulfills the verifiability property by using the same arguments given for our construction with one authority.

Consider that w.l.o.g the k -th authority outputs a commitment com_k that does not commit to a tuple of 0’s. If at least one voter j is honest, the probability that this voter outputs a ballot Bl_j such that com_k is a commitment to a tuple (z_1, \dots, z_N) and $z_j = \text{Bl}_j$ is negligible over the random coins of the j -th voter. Therefore, assuming that there is at least one honest voter, the authorities

can compute proofs of tally correctness by using the witness for the “trapdoor statement” in the relation only with negligible probability. Similarly, assuming that there is at least one honest voter, the voters can compute proofs of ballot correctness by using the witness for the “trapdoor statement” only with negligible probability over the random coins of the honest voters. In conclusion, the generalized construction fulfills (a statistical variant of) the verifiability property thanks to the verifiability of the (fully) verifiable eVote in Section 11 and to the fact that, in real mode, the sum of the messages encrypted in the m ciphertexts of a ballot is equal to a number in $\{0, 1\}$.

5.3 Privacy of the Construction

We use a selectively-secure model [CHK04] for our definition of privacy. In the game between the challenger and the adversary, the adversary has to declare its challenge at the outset of the game before receiving the public keys of the authorities. The adversary is allowed to receive the secret keys of all except one authority.

We show that our generalized construction fulfills this definition of privacy. First, we define the sets S^k and a method for computing the shares $v_{j,1}, \dots, v_{j,m}$ for a vote v_j . This method must guarantee that any subset of $m - 1$ authorities does not get any information about v_j . For simplicity, consider that $m = 2$. Then, the sets $S^1 = S^2 = S$ are equal by definition to $\{-p, \dots, p\}$, where p is a number of size super-polynomial in the security parameter. The message space of the PKE scheme must comprise numbers between $-Np$ and Np . To encrypt 0 (resp. 1), the voter chooses a random number v_1 in S and sets v_2 to $-v_1$ (resp. $-v_1 + 1$). It is easy to see that, except when either v_1 or v_2 equal $-p$, any value of v_2 (resp. v_1) can correspond to $v_1 = -v_2$ (resp. $v_2 = -v_1$) if the voter cast a vote for 0 or to $v_1 = -v_2 + 1$ (resp. $v_2 = -v_1 + 1$) if the voter cast a vote for 1. The case in which either v_1 or v_2 equal $-p$ occurs with negligible probability, which is guaranteed by choosing p to be super-polynomial in the security parameter. Consequently, each authority does not get any information on the vote v_j . This method can be generalized to the case $m > 2$. We skip the details.

Because the adversary receives the public-keys after sending the challenge, in the security proof we can define a hybrid experiment where the commitments in the public key commit to ballots $(\text{Bl}_1 \dots, \text{Bl}_N)$ computed on input the challenge messages. Like in the reduction of Section 11.2, we prove that the probability that the adversary submits an ill-formed ballot that is accepted by the ballot verification algorithm is negligible by using the computationally hiding property of the commitment scheme.

In the next hybrid experiments, the NIWI proofs of ballot correctness and of tally correctness can be computed by using the witness for the trapdoor statement, i.e., the randomness used to compute the commitments. Thanks to that, we are able to compute ballots where not all the ciphertexts encrypt the same message. This allows us to switch the message encrypted in one of the ciphertexts of the ballots and prove indistinguishability between the experiments by using the IND-CPA property of the PKE scheme.

To prove that our scheme fulfills a definition for privacy in a non-selective (i.e., full) security model, one can use complexity leveraging arguments. Such arguments can profit from the fact that, in our formulation, we required the number of voters N and the size of the message space to be independent of the security parameter. This allows the challenger to just guess the challenge messages in advance with constant probability. This requirement can be weakened to the case of N and size of message space logarithmic in the security parameter. We leave open how to achieve full security without complexity leveraging.

Note that we do not require any interaction between the authorities. The public keys of the authorities are completely *independent* from each other. Moreover, the authorities do not need any *coordination* (e.g., to run sequentially), i.e., the tally can be computed and publicly verified from the output of each authority individually. Thus, our techniques completely diverge from previous approaches to the problem.

6 On The Reusability of the Public Parameters

Our definition of verifiability does not prevent the following undesirable case. Consider an ill-formed ballot Bl_1 . Consider other valid ballots $\text{Bl}_2, \dots, \text{Bl}_N$ that encrypt respectively v_2, \dots, v_N . The authority is able to compute a tally $y = F(v_1, \dots, v_N)$ and a valid proof of tally correctness. Consider now other valid ballots $\text{Bl}'_2, \dots, \text{Bl}'_N$ that encrypt v'_2, \dots, v'_N . The authority can possibly compute another tally $y' = F(v'_1, v'_2, \dots, v'_N)$ and another proof of tally correctness. The problem is that the ill-formed ballot Bl_1 can be decrypted to more than one message.

This does not contradict our definition because, for $\text{Bl}_1, \dots, \text{Bl}_N$, there still exist messages v_1, \dots, v_N that satisfy the statement of the definition, i.e., given Pk and $\text{Bl}_1, \dots, \text{Bl}_N$, the authority cannot output two different results and two valid proofs of tally correctness for each of them. However, it can occur that for $\text{Pk}, \text{Bl}_1, \text{Bl}'_2, \dots, \text{Bl}'_N$, there are different messages that satisfy the definition. We remark that the public key Pk does not change.

Let us present a concrete example. Consider two 0/1 elections with only 2 voters. A ballot could possibly be *reused* in the second election, i.e., if the public parameters of the system are reused, the same ballot can be cast again. Given an ill-formed ballot Bl_1 , there could exist two ballots Bl_2 and Bl'_2 such that, in an election with ballots Bl_1 and Bl_2 , the result is 2, and, in a election with ballots Bl_1 and Bl'_2 , the result is 0. This can only happen if the first ballot is “associated” with vote 1 in the first election and with vote 0 in the second election. Therefore, the first and the second elections are incoherent. More undesirable issues would emerge if different tally functions could be computed in different elections carried out with the same parameters and ballots.

A stronger definition could state that, for all Pk and all Bl_1 , there exists m_1 such that, for all $\text{Bl}_2, \dots, \text{Bl}_N$, there exist m_2, \dots, m_n such that the authority is only able to output a tally $y = F(m_1, \dots, m_N)$ along with a valid proof of tally

correctness. We note that this is a simplification because a general definition should take into account multiple dishonest voters.

Fortunately, in our e-voting model, as well as in other traditional models, the parameters cannot be reused through different elections. Therefore, the above-mentioned problem does not occur.

In a stronger model in which the parameters can be reused, our constructions would not be secure. Nevertheless, in our (fully) verifiable construction, the inconsistency of results through different elections would occur only in the case that a malicious authority sets the commitment in the public key dishonestly to 0, which allows the computation of ill-formed ballots.

This state of affairs could be paralleled to the case of garbled circuits, where the original one-time version [Yao86,LP09] can be based on the minimal assumption of existence of one-way functions, whereas the reusable variant [GKP⁺13] is known to be implementable only under stronger assumptions. Similarly, in functional encryption, the schemes with bounded security [SS10,GVW12] can be based just on public key encryption, whereas the unbounded secure variants are only known to be implementable under very strong assumptions [GGHZ16]. For instance, the scheme of Sahai and Seyalioglu [SS10] becomes completely insecure when the adversary can decrypt a ciphertext with two different secret keys, exactly as it occurs for our schemes.

7 Related Work

Our work is inspired by the work of Badrinarayanan *et al.* [BGJS16], which puts forward the concept of verifiable functional encryption. (We note that the committing IBE of [GH07] can be seen as a weaker variant of verifiable identity-based encryption.) Our work shares with BGJS the idea of “engineering” multiple witnesses, which are needed when using NIWI proofs, to enforce privacy in conjunction with verifiability.

Notwithstanding, the constructions are quite different, especially due to the different requirements of functional encryption and e-voting. For instance, in the security definition of functional encryption, the keys are handed to the adversary, so one needs a proof that each secret key and ciphertext is computed correctly. Instead, in our case, the adversary does not see the secret key. We can profit from this fact to just prove that the claimed tally equals the evaluation of the tally function over all ballots.

Such complications in functional encryption introduce a severe limitation: in the security reduction of BGJS, it is fundamental that the public key contain a commitment that in some hybrid experiment is set to the challenge ciphertext. Therefore, it is assumed that the adversary commits to the challenge before receiving the public key, i.e., security is proven in the *selective* model [CHK04]. On the contrary, our constructions are secure in the *full* (i.e., non-selective) model.

In other respects, in e-voting we face new challenges. In BGJS, the challenger computes the NIWI on input a witness that comprises *all* the secret keys and

proves the well-formedness of all the secret keys except one, but, in addition, proves that all the secret keys decrypt some challenge ciphertext correctly. This is sufficient to use the IND-CPA property of functional encryption to prove indistinguishability between two hybrid experiments where the message encrypted in one of the ciphertexts is switched from m_0 to m_1 . The reason is that the secret keys are supposed to be for the same function f such that $f(m_0) = f(m_1)$. (More concretely, in the IND-CPA property of functional encryption, the adversary is allowed to receive secret keys for a function f that evaluates both challenge messages to the same value.) Therefore, the secret keys do not allow to distinguish between the two ciphertexts. In our setting, we can *only* input to the NIWI all the secret keys except one. Otherwise we could not use the IND-CPA property to prove indistinguishability between two hybrid experiments where the encrypted message is switched from m_0 to m_1 .

Furthermore, in our (full) privacy definition, we have to handle challenge tuples that contain ill-formed ballots, whereas in verifiable multi-input functional encryption the challenge contains only honestly computed ciphertexts. Therefore, the differences between the two settings make the respective techniques utterly incomparable.

It is tempting to think that the construction of BGJS of multi-input verifiable functional encryption (which extends multi-input functional encryption of Goldwasser *et al.* [GGG⁺14]) can be directly used to construct a verifiable eVote. Though it seems plausible, we did not verify that. However, this would eventually result in a verifiable eVote based on indistinguishability obfuscation [GGH⁺13], a very strong assumption, and would only be secure in the selective model.

Needless to say, our techniques, as well as the ones of BGJS, owe a lot to the celebrated FLS' OR trick [FLS90]. They can be viewed as a generalization of it.

Kiayias *et al.* [KZZ15] (see also [CZZ⁺15] for a distributed implementation) put forth a verifiable eVote without trust assumptions that represents a breakthrough along this direction, but diverges from ours in several fundamental aspects:

- It requires interaction between the voters and the board, whereas all our algorithms are *non-interactive*.
- Receipt-freeness, accountability and degree of dependence on secure channels to distribute vote codes are undetermined issues. In our scheme, voters can verify the election if they just know the ballot they cast. In particular, voters do not need to store the randomness used to compute it, and the authority cannot cheat in the tally process.
- It does not achieve *universal verifiability*, whereas ours does.
- Its information-theoretical verifiability is parameterized and depends on the number of honest voters, whereas ours is *perfect*, i.e., the probability of a wrong tally being accepted by the verification algorithm is equal to *zero*.
- Its privacy can be reduced to group-based assumptions at the cost of using complexity leveraging and assuming sub-exponential security, whereas

ours only requires the standard version of Decision Linear assumption with polynomial security.⁵

- Its definition of verifiability requires an extraction property, whereas ours does not.

Moran and Naor [MN06] construct an universally verifiable e-voting protocol with very strong provable-security properties. However, it assumes either the availability of a “random beacon” that has to be sampled honestly or the soundness of the Fiat-Shamir’s heuristic. Therefore, verifiability does not hold unconditionally, i.e., without any assumption (both physical or computational).

We are not aware of other traditional e-voting schemes that achieve perfect verifiability without interaction and without trust assumptions. We refer to [CGK⁺16] and [BCG⁺15] for a survey.

We point out that our definition of verifiability is motivated by the guidelines of [CGK⁺16]. In its formalization, our definition is similar to the ones of [GIR16], the verifiability for multi-input functional encryption of BGJS and the uniqueness of tally of Bernhard *et al.* [BCG⁺15]. Anyhow, the latter is formulated to hold only against computationally bounded adversaries and both BGJS16 and Bernhard *et al.* do not take into account our condition (2) for verifiability.⁶ See also [KRS10] for symbolic approaches to verifiability.

Our privacy notion is inspired by the one of Benaloh [Ben87], which is called “PRIV” in [BCG⁺15]. We reformulate it by using modern terminology and we conveniently modify it to withstand the attacks shown in [BCG⁺15]. We refer to [BCG⁺15] for a survey on definitions of privacy for e-voting.

Perfect verifiability and perfect correctness seem incompatible with receipt-freeness [BT94,SK95,MH96,MN06,DKR06,CFG15]. Notwithstanding, we think that it should be possible to define a statistical variant of verifiability that could attain some form of receipt-freeness. Another possibility could be to resort to some voting server trusted for receipt-freeness but not for privacy, such as the server that re-randomizes the ballots in BeleniosRF of Cortier, Fuchsbaauer and Galindo [CFG15]. (We note that they also address the problem of authenticity that we neglect.) As it is out of the scope of this work, we deliberately omit receipt-freeness in our treatment.

Recently, Bellare, Fuchsbaauer and Scafuro [BFS16] started the study of security of NIZKs in face of public parameter subversion. They showed the impossibility of attaining subversion soundness while retaining zero-knowledge, thus justifying our need of sidestepping NIZKs.

⁵ At some point in the security reduction for our (fully) verifiable eVote, we make use of the fact that the number of voters N is a constant independent of the security parameter that could be viewed as a complexity leveraging trick or as problematic in the case that N be large. But we stress that this is done only for simplicity of exposition and we sketch how the reduction and our results can be generalized even to the case of $N(\cdot)$ function of the security parameter.

⁶ Needless to say, for many applications of multi-input functional encryption, the lack of condition (2) could not pose a threat.

8 Definitions

Notation. A *negligible* function $\text{negl}(k)$ is a function that is smaller than the inverse of any polynomial in k (from a certain point and on). We denote by $[n]$ the set of numbers $\{1, \dots, n\}$. If S is a finite set, we denote by $a \leftarrow S$ the process of setting a equal to a uniformly chosen element of S . With a slight abuse of notation, we assume the existence of a special symbol \perp that does not belong to $\{0, 1\}^*$.

If A is an algorithm, then $A(x_1, x_2, \dots)$ denotes the probability distribution of the output of A when A is run on input (x_1, x_2, \dots) and randomly chosen coin tosses. Instead, $A(x_1, x_2, \dots; r)$ denotes the output of A when run on input (x_1, x_2, \dots) and (sufficiently long) coin tosses r . All algorithms, unless explicitly noted, are probabilistic polynomial time (PPT) and all adversaries are modeled by non-uniform PPT algorithms.

If A is a PPT algorithm, we say that $y \in A(x)$ iff there exists a random value r such that $y = A(x; r)$; in that case, we say that y is in the range of $A(x)$. If E is an event in a probability space, \bar{E} denotes its complement.

The following definition is used in the definition of verifiability. Essentially, it states that a tally y is compatible with votes z_1, \dots, z_k if the latter values are in its pre-image.

Definition 1 Given a function $F(x_1, \dots, x_n) : A^n \rightarrow B$, we say that a value $y \in B$ is compatible with $z_1, \dots, z_k \in A$ at indices $i_1, \dots, i_k \in [n]$ if y is in the range of the restriction $F|_{C_{z_1, \dots, z_k, i_1, \dots, i_k}}$ of F to $C_{z_1, \dots, z_n, i_1, \dots, i_n} \triangleq \{(x_1, \dots, x_n) \mid \forall j \in [k], x_{i_j} = z_j\}$.

8.1 E-Voting Schemes

An e-voting scheme (eVote, in short) is associated with a natural number $N > 0$, the *number of voters*, a set \mathcal{M} , the *domain of valid votes*, a set Σ , the *range of possible results*, and a *tally function* $F : (\mathcal{M} \cup \{\perp\})^N \rightarrow \Sigma \cup \{\perp\}$, where \perp is a special symbol not in \mathcal{M} that indicates that the vote is invalid or has not been cast and so has not to be “counted“. We allow the tally function to take as input the special symbol \perp that denotes either an invalid ballot or an abstention and to output \perp to indicate an error, but we require that the tally outputs an error on input a sequence of string iff all strings are equal to \perp . Formally, the tally function is defined as follows.

Definition 2 [Tally function] A function F is a tally function if there exists a natural number $N > 1$, and sets $\mathcal{M}, \Sigma \subset \{0, 1\}^*$ such that the domain of F is $\mathcal{M} \cup \{\perp\}$, the range is $\Sigma \cup \{\perp\}$ and for all strings $m_1, \dots, m_N \in \mathcal{M} \cup \{\perp\}$, it holds that $F(m_1, \dots, m_N) = \perp$ iff $m_1 = \perp, \dots, m_N = \perp$.

Before presenting the formalization of an eVote, we explain how its algorithms have to be employed to carry out an election.

The voting ceremony. The voting ceremony occurs as follows.

- Setup phase. An authority (also called voting authority or election authority) uses a **Setup** algorithm to compute a pair of public-key Pk and secret-key Sk .
- Voting phase. Each of the N voters, identified by a number $j \in [N]$, holds her intended vote v and runs an algorithm **Cast** on input the identifier j , the public-key Pk and the vote $v \in \mathcal{M}$ to compute a ballot Bl and sends it to an append-only public bulletin board (PBB).
- Tallying phase. The well-formedness of each ballot Bl published in the PBB can be publicly verified by means of an algorithm **VerifyBallot**. If a voter did not cast her vote or her ballot turns out to be invalid it is replaced⁷ by the special symbol \perp in the next steps.

The authority runs an *evaluation tally* algorithm **EvalTally** on input the public- and secret-key, the N ballots resulting from the previous replacement, and outputs the result of the election (i.e, the tally), that can be eventually the special symbol \perp to indicate an error, and in addition a proof of correctness of the computation.

- Verification phase. There is a verification algorithm **VerifyTally** that takes as input the public-key, a tuple of N strings that represent either ballots or the special symbol \perp (that indicates abstention or invalid vote for the corresponding voter), the result of the election and the proof of correctness, and outputs a value in $\{\text{OK}, \perp\}$.

Each participant, not necessarily a voter, can verify the correctness of the result of the election as follows. First, verify whether the ballots cast by the voters are valid using the **VerifyBallot** algorithm and thus can verify whether the authority replaced with \perp only the invalid ballots and put \perp in correspondence of a voter who did not cast her vote.

After that, run the **VerifyTally** algorithm with the public-key, the N ballots resulting from the previous replacement (if any), the claimed result of the election and the proof output by the authority.

Definition 3 [E-voting Scheme] A $(N, \mathcal{M}, \Sigma, F)$ -*e-voting scheme* **EVOTE** for number of voters $N > 1$, domain of valid votes \mathcal{M} , range of possible results Σ and tally function $F : (\mathcal{M} \cup \{\perp\})^N \rightarrow \Sigma \cup \{\perp\}$ is a tuple

$\text{EVOTE} \triangleq (\text{Setup}, \text{Cast}, \text{VerifyBallot}, \text{EvalTally}, \text{VerifyTally})$ is a tuple of 5 PPT algorithms, with **VerifyBallot** and **VerifyTally** deterministic, following the following syntax:

1. **Setup**(1^λ): on input the security parameter in unary, outputs the *public-key* Pk and the *secret-key* Sk .
2. **Cast**(Pk, j, v): on input the public-key Pk , the identifier $j \in [N]$ of the voter, and a vote $v \in \mathcal{M}$, outputs a *ballot* Bl ;
3. **VerifyBallot**(Pk, j, Bl), on input the public-key Pk , the identifier $j \in [N]$ of the voter and a ballot Bl , output a value in $\{\text{OK}, \perp\}$;

⁷ Precisely, a new row in which the previous ciphertext is replaced by \perp is appended to the PBB and in the next steps only the new row will be used.

4. $\text{EvalTally}(\text{Pk}, \text{Sk}, \text{Bl}_1, \dots, \text{Bl}_N)$: on input the public-key Pk , and a tuple of N strings that can be either ballots cast by a voter or the special symbol \perp to indicate an abstention or an invalid vote, output the tally $y \in \Sigma \cup \{\perp\}$ and a proof γ of the correctness of the computation.
5. $\text{VerifyTally}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y, \gamma)$: on input the public-key Pk , a tuple of N strings that can be either ballots cast by a voter or the special symbol \perp to indicate an abstention or an invalid vote, a claimed tally $y \in \{0, 1\}^* \cup \{\perp\}$ and claimed proof γ of the correctness of the computation, output a value in $\{\text{OK}, \perp\}$.

In addition we require that an eVote satisfy correctness, verifiability, and privacy as defined next. We also define weak verifiability and weak privacy, and we require that a weakly verifiable eVote satisfy correctness, weak verifiability and weak privacy.

Correctness and verifiability.

- *(Perfect) Correctness.* We require the following conditions (1) and (2) to hold.
 1. Let Abst be a special symbol not in $\mathcal{M} \cup \{\perp\}$ used to denote that a voter did not cast her vote.⁸ For all $\text{Pk} \in \text{Setup}(1^\lambda)$, all $m_1, \dots, m_N \in \mathcal{M} \cup \{\perp, \text{Abst}\}$, all $(\text{Bl}_j)_{j=1}^N$ such that for all $j \in [N]$, $\text{Bl}_j = \perp$ if $m_j = \text{Abst}$, $\text{Bl}_j \in \text{Cast}(\text{Pk}, j, m_j)$ if $m_j \in \mathcal{M}$ and $\text{Bl}_j \in \text{Cast}(\text{Pk}, j, \perp)$ otherwise, the following two conditions (a) and (b) hold:
 - (a) For all $j \in [N]$, if $m_j \neq \perp$ then $\text{VerifyBallot}(\text{Pk}, j, \text{Bl}_j) = \text{OK}$.
 - (b) if $(y, \gamma) \stackrel{\Delta}{=} \text{EvalTally}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N)$ then it holds that:
 $y = F(m_1, \dots, m_N)$ and $\text{VerifyTally}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y, \gamma) = \text{OK}$.
 2. For all $\text{Pk} \in \text{Setup}(1^\lambda)$, $\text{Bl}_1, \dots, \text{Bl}_N \in \{0, 1\}^* \cup \{\perp\}$, if $S \stackrel{\Delta}{=} \{j \mid \text{Bl}_j \neq \perp \wedge \text{VerifyBallot}(\text{Pk}, j, \text{Bl}_j) = \perp\}$ and $\text{Bl}'_1, \dots, \text{Bl}'_N$ are such that for all $j \in [N]$, $\text{Bl}'_j = \text{Bl}_j$ if $j \notin S$ and $\text{Bl}_j = \perp$ otherwise, it holds that:
If $(y, \gamma) \stackrel{\Delta}{=} \text{EvalTally}(\text{Pk}, \text{Bl}'_1, \dots, \text{Bl}'_N)$ then $\text{VerifyTally}(\text{Pk}, \text{Bl}'_1, \dots, \text{Bl}'_N, y, \gamma) = \text{OK}$.
- *Weak verifiability.* We require the following conditions (1) and (2) to hold.
 1. For all $\text{Pk} \in \{0, 1\}^*$, $\text{Bl}_1, \dots, \text{Bl}_N \in \{0, 1\}^* \cup \{\perp\}$, there exist $m_1, \dots, m_N \in \mathcal{M} \cup \{\perp\}$ such that for all $y \neq \perp$ and γ in $\{0, 1\}^*$, if $S \stackrel{\Delta}{=} \{j \mid \text{Bl}_j \neq \perp \wedge \text{VerifyBallot}(\text{Pk}, j, \text{Bl}_j) = \perp\}$ and $\text{Ct}'_1, \dots, \text{Ct}'_N$ are such that for all $j \in [N]$, $\text{Bl}'_j = \text{Bl}_j$ if $j \notin S$ and $\text{Bl}_j = \perp$ otherwise, it holds that:
if $\text{VerifyTally}(\text{Pk}, \text{Bl}'_1, \dots, \text{Bl}'_N, y, \gamma) = \text{OK}$ then $y = F(m_1, \dots, m_N)$.

⁸ In the following definition, we need Abst to differentiate the case of a voter who did not cast a vote at all (Abst) from the case of a voter who casts \perp as her own vote but wishes to preserve the anonymity of her choice; however, in both cases the correctness has to guarantee that the result of the election equals the output of the tally function in which the inputs corresponding to the voters who, either cast \perp or did not cast any vote, are both counted as \perp .

2. For all $\text{Pk} \in \{0, 1\}^\lambda$, all $k \in [N]$, $i_1, \dots, i_k \in [N]$, all $m_{i_1}, \dots, m_{i_k} \in \mathcal{M} \cup \{\perp\}$, all $\text{Bl}_1, \dots, \text{Bl}_N \in \{0, 1\}^* \cup \{\perp\}$ such that for all $j \in [k]$, $\text{Bl}_j \in \text{Cast}(\text{Pk}, m_{i_j})$ and $\text{VerifyBallot}(\text{Pk}, \text{Bl}_j) = \text{OK}$, if $S \triangleq \{j \mid \text{Bl}_j \neq \perp \wedge \text{VerifyBallot}(\text{Pk}, j, \text{Bl}_j) = \perp\}$ and $\text{Ct}'_1, \dots, \text{Ct}'_N$ are such that for all $j \in [N]$, $\text{Bl}'_j = \text{Bl}_j$ if $j \notin S$ and $\text{Bl}_j = \perp$ otherwise, it holds that: if there exist $y \in \{0, 1\}^*$ and $\gamma \in \{0, 1\}^*$ such that $\text{VerifyTally}(\text{Pk}, \text{Bl}'_1, \dots, \text{Bl}'_N, y, \gamma) = \text{OK}$, then y is compatible with m_{i_1}, \dots, m_{i_k} at indices i_1, \dots, i_k .
- *Verifiability.* We require the following conditions (1) and (2) to hold.
1. For all $\text{Pk} \in \{0, 1\}^*$, $\text{Bl}_1, \dots, \text{Bl}_N \in \{0, 1\}^* \cup \{\perp\}$, there exist $m_1, \dots, m_N \in \mathcal{M} \cup \{\perp\}$ such that for all $y \in \{0, 1\}^* \cup \{\perp\}$ and γ in $\{0, 1\}^*$, if $S \triangleq \{j \mid \text{Bl}_j \neq \perp \wedge \text{VerifyBallot}(\text{Pk}, j, \text{Bl}_j) = \perp\}$ and $\text{Ct}'_1, \dots, \text{Ct}'_N$ are such that for all $j \in [N]$, $\text{Bl}'_j = \text{Bl}_j$ if $j \notin S$ and $\text{Bl}_j = \perp$ otherwise, it holds that: if $\text{VerifyTally}(\text{Pk}, \text{Bl}'_1, \dots, \text{Bl}'_N, y, \gamma) = \text{OK}$ then $y = F(m_1, \dots, m_N)$.
 2. For all $\text{Pk} \in \{0, 1\}^\lambda$, all $k \in [N]$, $i_1, \dots, i_k \in [N]$, all $m_{i_1}, \dots, m_{i_k} \in \mathcal{M} \cup \{\perp\}$, all $\text{Bl}_1, \dots, \text{Bl}_N \in \{0, 1\}^* \cup \{\perp\}$ such that for all $j \in [k]$, $\text{Bl}_j \in \text{Cast}(\text{Pk}, m_{i_j})$ and $\text{VerifyBallot}(\text{Pk}, \text{Bl}_j) = \text{OK}$, if $S \triangleq \{j \mid \text{Bl}_j \neq \perp \wedge \text{VerifyBallot}(\text{Pk}, j, \text{Bl}_j) = \perp\}$ and $\text{Ct}'_1, \dots, \text{Ct}'_N$ are such that for all $j \in [N]$, $\text{Bl}'_j = \text{Bl}_j$ if $j \notin S$ and $\text{Bl}_j = \perp$ otherwise, it holds that: if there exist $y \in \{0, 1\}^* \cup \{\perp\}$ and $\gamma \in \{0, 1\}^*$ such that $\text{VerifyTally}(\text{Pk}, \text{Bl}'_1, \dots, \text{Bl}'_N, y, \gamma) = \text{OK}$, then y is compatible with m_{i_1}, \dots, m_{i_k} at indices i_1, \dots, i_k .

(Note that the difference between condition (2) of verifiability and condition (2) of weak verifiability lies in the fact that in the latter y has to be $\neq \perp$, whereas in the former the statement has to hold even for $y = \perp$.)

In this work, we use the term verifiability and full verifiability interchangeably to differentiate it from weak verifiability.

Privacy. We formalize the notion of *privacy* in the style of indistinguishability-based security. The privacy for a $(n, \mathcal{M}, \Sigma, F)$ -eVote $\text{EVOTE} \triangleq (\text{Setup}, \text{Cast}, \text{VerifyBallot}, \text{EvalTally}, \text{VerifyTally})$ is formalized by means of the game in Fig. 3 $\text{Priv}_{\mathcal{A}}^{N, \mathcal{M}, \Sigma, F, \text{EVOTE}}$ between a stateful adversary \mathcal{A} and a *challenger* \mathcal{C} .

The advantage of adversary \mathcal{A} in the above game is defined as

$$\text{Adv}_{\mathcal{A}}^{\text{EVOTE}, \text{Priv}}(1^\lambda) \triangleq |\text{Prob}[\text{Priv}_{\mathcal{A}}^{N, \mathcal{M}, \Sigma, F, \text{EVOTE}}(1^\lambda) = 1] - 1/2|$$

Definition 4 We say that EVOTE for parameters $(N, \mathcal{M}, \Sigma, F)$ is *private* or IND-Secure if all PPT adversaries \mathcal{A} have at most advantage negligible in λ in the above game.

Definition 5 We say that EVOTE for parameters $(N, \mathcal{M}, \Sigma, F)$ is *weakly private* or wIND-Secure if all PPT adversaries \mathcal{A} have at most advantage negligible in λ in a game $\text{WeakPriv}_{\mathcal{A}}^{N, \mathcal{M}, \Sigma, F, \text{EVOTE}}(1^\lambda)$ that is identical to the one above except

$\text{Priv}_{\mathcal{A}}^{N,D,\Sigma,F,\text{EVOTE}}(1^\lambda)$

- Setup phase. \mathcal{C} generates $(\text{Pk}, \text{Sk}) \leftarrow \text{Setup}(1^\lambda)$, choose a random bit $b \leftarrow \{0, 1\}$ and runs \mathcal{A} on input Pk ;
- Query phase. The adversary \mathcal{A} outputs two tuples $M_0 \triangleq (m_{0,1}, \dots, m_{0,N})$ and $M_1 \triangleq (m_{1,1}, \dots, m_{1,N})$, and a set $S \subset [N]$.
(The set S indicates that the strings corresponding to the indices in S are possibly dishonest ballots and the others are supposed to be votes that have to be given as input to the **Cast** procedure.)
- Challenge phase. The challenger computes as follows. For all $j \in [N]$, if $j \in S$ then set $\text{Bl}_j \triangleq m_{0,j}$ and set $\text{Bl}_j \leftarrow \text{Cast}(\text{Pk}, j, m_{b,j})$ otherwise. For all $j \in S$, if $\text{VerifyBallot}(\text{Pk}, j, \text{Bl}_j) = \perp$ set $\text{Bl}_j \triangleq \perp$.
Compute $(y, \gamma) = \text{EvalTally}(\text{Pk}, \text{Sk}, \text{Bl}_1, \dots, \text{Bl}_N)$ and return $(\text{Bl}_1, \dots, \text{Bl}_N, y, \gamma)$ to the adversary.
- Output. At some point the adversary outputs its guess b' .
- Winning condition. The adversary wins the game if all the following conditions hold:
 1. $b' = b$.
 2. For all $j \in S$, $m_{0,j} = m_{1,j}$. (That is, if the adversary submits a dishonest ballot it has to be the same in both tuples.)
 3. For all $d_1, \dots, d_N \in \mathcal{M} \cup \{\perp\}$, for all $j \in [N]$, let $m'_{0,j} \triangleq m'_{1,j} \triangleq d_j$ if $j \in S$, and for all $b \in \{0, 1\}$ let $m'_{b,j} \triangleq m_{b,j}$ if $m_{b,j} \in \mathcal{M}$ and $m'_{b,j} \triangleq \perp$ if $m_{b,j} \notin \mathcal{M}$. Then, $F(m'_{0,1}, \dots, m'_{0,N}) = F(m'_{1,1}, \dots, m'_{1,N})$.
(That is, the tally function gives same output on the two tuples, even if the ballots corresponding to indices in S are replaced by arbitrary messages in $\mathcal{M} \cup \{\perp\}$.)

Fig. 3. Definition of privacy

that the adversary is required to output a set S that is empty (i.e., adversary cannot submit dishonest ballots).

Remark 1 We make some remarks on the previous definitions.

- Our definitions suppose that the algorithm `VerifyBallot` have to be run on each ballot before the `VerifyTally` algorithm replacing any ballot that either was not accepted by the `VerifyBallot` algorithm or was not cast by the corresponding voter with \perp . Another possibility would be to let the `VerifyTally` algorithm to make these tasks by itself.
- We required that `VerifyBallot` and `VerifyTally` be deterministic algorithms. Alternatively, they can be defined as PPT but definitions of weak verifiability and verifiability would have to be changed accordingly to hold with probability 1 over the random coins of the algorithms.
- Our definition is parameterized by the number of voters N . It is possible to define a more restricted eVote that may be possibly “unbounded“. Note that our definition is more general and for instance takes in account e-voting schemes in which the public-key has size proportional to the number of voters and has to depend on it.
- Both condition (2) of verifiability and condition (2) of weak verifiability lie in some sense between correctness and verifiability as they state a requirement about honest voters.
- In our weakly verifiable construction of Section 10, the `VerifyBallot` algorithm could be completely discarded as it accepts any ballot. Both for sake of generality (there could exist some weakly verifiable eVote that makes a non-trivial use of `VerifyBallot`) and for not overburdening the presentation, we preferred to not differentiate the syntax of weakly verifiable eVote schemes from verifiable eVote ones.
- For the necessity of condition (2) of correctness, we invite the reader to refer to the discussion in Section 2.3.
- As shown in [BCG⁺15] the definition of “Benaloh“ (recall that we restate it using modern terminology) is subject to attacks when instantiated for specific tally functions like the majority. Nonetheless, ours is strengthened to withstand such counter-examples. This is done adding the 3-rd winning condition.

9 Building Blocks

Our constructions use perfectly binding commitment schemes, (one-message) non-interactive witness-indistinguishable proof systems with perfect soundness for NP [GOS06] (see also [FLS90, DN00, DN00, BOV03, BP15]) and IND-CPA public key encryption with perfect correctness and unique secret key. In this section, we recall the definitions of those primitives.

Definition 6 [IND-CPA secure PKE with perfect correctness and unique secret key] An IND-CPA (or semantically) secure Public Key Encryption (PKE)

scheme consists of three PPT algorithms (Setup , Encrypt , Decrypt) defined as follows.

- $\text{Setup}(1^\lambda)$: On input 1^λ , it outputs public key Pk and decryption key Sk .
- $\text{Encrypt}(m, \text{Pk})$: On input message m and the public key, it outputs ciphertext Ct .
- $\text{Decrypt}(\text{Ct}, \text{Sk})$: On input ciphertext Ct and the decryption key, it outputs m .

The PKE scheme is said to be IND-CPA (or semantically) secure if for any PPT adversary \mathcal{A} , there exists a negligible function $\nu(\cdot)$ such that the following is satisfied for any two messages m_0, m_1 and for $b \in \{0, 1\}$:

$$|\Pr [\mathcal{A}(1^\lambda, \text{Encrypt}(m_0, \text{Pk})) = b] - \Pr [\mathcal{A}(1^\lambda, \text{Encrypt}(m_1, \text{Pk})) = b]| \leq \nu(\lambda).$$

Perfect correctness requires that, for all pairs $(\text{Pk}, \text{Sk}) \in \text{Setup}$, for all messages m and all ciphertexts Ct output by $\text{Encrypt}(\text{Pk}, m)$, $\text{Decrypt}(\text{Ct}, \text{Sk}) = m$ must hold. *Unique secret-key* requires that, for all Pk , there exist at most *one* Sk such that $(\text{Pk}, \text{Sk}) \in \text{Setup}(1^\lambda)$ (i.e., there exist some randomness $r \in \{0, 1\}^\lambda$ such that $(\text{Pk}, \text{Sk}) = \text{Setup}(1^\lambda; r)$).

The Decision Linear Encryption scheme [BBS04] fulfills those properties. It is secure under the Decision Linear Assumption [BBS04]. We recall them next.

Bilinear Groups. We assume the existence of a PPT algorithm $\mathcal{G}(1^\lambda)$, the *bilinear group generator*, that outputs a *pairing group setup* $(p, \mathbb{G}, \mathbb{G}_t, \mathbf{e}, g)$, where \mathbb{G} and \mathbb{G}_t are multiplicative groups of prime order p and $\mathbf{e} : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_t$ is a *bilinear map* satisfying the following three properties: (1) bilinearity, i.e., $\mathbf{e}(g^x, g^y) = \mathbf{e}(g, g)^{xy}$; (2) non-degeneracy, i.e., for all generators $g \in \mathbb{G}$, $\mathbf{e}(g, g)$ generates \mathbb{G}_t ; (3) efficiency, i.e., \mathbf{e} can be computed in polynomial time. Furthermore, we assume that the following problem be hard for \mathcal{G} .

Assumption 1 (Decision Linear Assumption for \mathcal{G} , [BBS04]) *Let $(p, \mathbb{G}, \mathbb{G}_t, \mathbf{e}, g)$ be a pairing group setup output by a \mathcal{G} defined as above, and let g_1, g_2 and g_3 be generators of \mathbb{G} . Given $(g_1, g_2, g_3, g_1^a, g_2^b, g_3^c)$, where a and b are picked randomly from \mathbb{Z}_p , the Decision Linear (DLIN) assumption is to decide whether $c = a + b \pmod p$. Precisely, the advantage of an adversary \mathcal{A} in solving the Decision Linear assumption is given by:*

$$\begin{aligned} & |\Pr [\mathcal{A}(\mathbb{G}, p, g_1, g_2, g_3, g_1^a, g_2^b, g_3^{a+b}) = 1 \mid (p, \mathbb{G}, \mathbb{G}_t, \mathbf{e}, g) \leftarrow \mathcal{G}(1^\lambda); \\ & (g_1, g_2, g_3) \leftarrow \mathbb{G}; (a, b) \leftarrow \mathbb{Z}_p] - \\ & \Pr [\mathcal{A}(\mathbb{G}, p, g_1, g_2, g_3, g_1^a, g_2^b, g_3^c) = 1 \mid (p, \mathbb{G}, \mathbb{G}_t, \mathbf{e}, g) \leftarrow \mathcal{G}(1^\lambda); \\ & (g_1, g_2, g_3) \leftarrow \mathbb{G}; (a, b, c) \leftarrow \mathbb{Z}_p]| \end{aligned}$$

The Decision Linear assumption states that the advantage of \mathcal{A} is negligible in λ . Boneh et al. [BBS04] provide a bilinear group generator \mathcal{G} for which such assumption is conjectured to hold.

Decision Linear Encryption Scheme. Consider the following PKE scheme described by a setup algorithm **Setup**, an encryption algorithm **Encrypt** and a decryption algorithm **Decrypt**.

Setup(1^λ): pick $(p, \mathbb{G}, \mathbb{G}_t, \mathbf{e}, g) \leftarrow \mathcal{G}(1^\lambda)$, pick randomly $g_3 \leftarrow \mathbb{G}$ and $(x, y) \leftarrow \mathbb{Z}_p$. Compute $g_1 = g_3^{1/x}$ and $g_2 = g_3^{1/y}$. Output the public key $\text{Pk} = (\mathbb{G}, p, g_1, g_2, g_3)$ and the secret key $\text{Sk} = (\text{Pk}, x, y)$.

Encrypt(Pk, m): on input a public key Pk and a message $m \in \mathbb{G}$, pick random $(a, b) \leftarrow \mathbb{Z}_p$. Output a ciphertext $\text{Ct} = (g_1^a, g_2^b, m \cdot g_3^{a+b})$.

Decrypt(Sk, Ct): on input a secret key Sk and a ciphertext $\text{Ct} = (c_1, c_2, c_3)$, output $m = c_3 / (c_1^x c_2^y)$.

This scheme fulfills the IND-CPA property under the Decision Linear assumption (see [BBS04] for details) and it is easy to verify that it fulfills the unique secret-key property.

Definition 7 [(Perfectly binding) Commitment Schemes] A commitment scheme Com is a PPT algorithm that takes as input a string x and randomness $r \in \{0, 1\}^k$ and outputs $\text{com} \leftarrow \text{Com}(x; r)$. A perfectly binding commitment scheme must satisfy the following properties:

- Perfectly Binding: This property states that two different strings cannot have the same commitment. More formally, $\forall x_1 \neq x_2$ and $r_1, r_2, \text{Com}(x_1; r_1) \neq \text{Com}(x_2; r_2)$.
- Computational Hiding: For all strings x_0 and x_1 (of the same length), for all PPT adversaries \mathcal{A} there exists a negligible function $\nu(\cdot)$ such that: $|\Pr_{r \in \{0, 1\}^k}[\mathcal{A}(\text{Com}(x_0; r)) = 1] - \Pr_{r \in \{0, 1\}^k}[\mathcal{A}(\text{Com}(x_1; r)) = 1]| \leq \nu(k)$.

NIWI proof systems. Next, we define (one-message) non-interactive witness indistinguishability (NIWI) proof systems [GOS06]. Groth *et al.* [GOS06] construct such NIWIs for all languages in NP, and in particular for **CircuitSat**.

Definition 8 [Non-interactive Proof System] A non-interactive proof system for a language L with a PPT relation R is a tuple of algorithms (**Prove**, **Verify**) such that **Verify** outputs a symbol in $\{\text{OK}, \perp\}$ and the following properties hold:

- Perfect Completeness: For every $(x, w) \in R$, it holds that $\Pr[\text{Verify}(x, \text{Prove}(x, w)) = 1] = \text{OK}$, where the probability is taken over the coins of **Prove** and **Verify**.
- Perfect Soundness: For every adversary \mathcal{A} it holds that:

$$\Pr \left[\begin{array}{l} \text{Verify}(x, \pi) = \text{OK} \wedge x \notin L : \\ (x, \pi) \leftarrow \mathcal{A}(1^k) \end{array} \right] = 0.$$

Definition 9 [NIWI] A non-interactive proof system $\text{NIWI} = (\text{Prove}, \text{Verify})$ for a language L with a PPT relation R is witness-indistinguishable (WI, in short) if for any triplet (x, w_0, w_1) such that $(x, w_0) \in R$ and $(x, w_1) \in R$, the distributions $\{\text{Prove}(x, w_0)\}$ and $\{\text{Prove}(x, w_1)\}$ are computationally indistinguishable.

10 Our Weakly Verifiable eVote

In this Section, we present our eVote scheme EVOTE that is wIND-Secure and weakly verifiable.

Definition 10 [EVOTE] Let $\mathcal{E} = (\mathcal{E}.\text{Setup}, \mathcal{E}.\text{Encrypt}, \mathcal{E}.\text{Decrypt})$ be a PKE scheme with perfectly correctness and unique secret-key (see Def. 6), and let $\text{NIWI}^{\text{dec}} \triangleq (\text{Prove}^{\text{dec}}, \text{Verify}^{\text{dec,full}})$ be a NIWI proof systems for the relation \mathbf{R}^{dec} to be specified later.

We define a $(N, \mathcal{M}, \Sigma, F)$ -eVote

EVOTE $^{N, \mathcal{M}, \Sigma, F, \mathcal{E}, \text{NIWI}^{\text{dec}}} = (\text{Setup}, \text{Cast}, \text{VerifyBallot}, \text{EvalTally}, \text{VerifyTally})$ as follows.

- **Setup**(1^λ): on input the security parameter in unary, compute as follows.
 1. For all $l \in [3]$, run $(\mathcal{E}.\text{Pk}_l, \mathcal{E}.\text{Sk}_l) = \mathcal{E}.\text{Setup}(1^\lambda; s_l)$ with randomness s_l .
 2. Output $\text{Pk} \triangleq (\mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3)$ and $\text{Sk} \triangleq (\mathcal{E}.\text{Sk}_1, \mathcal{E}.\text{Sk}_2, s_1, s_2)$.
(Actually, as the randomness for the setup of a PKE uniquely determines the secret-key, it would be sufficient to just include the s_l 's in Sk .)
- **Cast**(Pk, j, v): on input the public-key $\text{Pk} \triangleq (\mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3)$, the index $j \in [N]$ of the voter, a vote v , compute as follows.
 1. For all $l \in [3]$, compute $\text{Ct}_{j,l} \leftarrow \mathcal{E}.\text{Encrypt}(\mathcal{E}.\text{Pk}_l, v)$.
 2. Output $\text{Blt}_j \triangleq (\text{Ct}_{j,1}, \dots, \text{Ct}_{j,3})$.
- **VerifyBallot**(Pk, j, Blt): on input the public-key $\text{Pk} \triangleq (\mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3)$, the index $j \in [N]$ of the voter and a ballot $\text{Blt} \triangleq (\text{Ct}_1, \dots, \text{Ct}_3, \pi)$, just output OK (i.e., accept any ballot, even invalid ones).
- **EvalTally**($\text{Pk}, \text{Sk}, \text{Blt}_1, \dots, \text{Blt}_N$): on input the public-key $\text{Pk} \triangleq (\mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3)$, the secret-key $\text{Sk} \triangleq (\mathcal{E}.\text{Sk}_1, \mathcal{E}.\text{Sk}_2, s_1, s_2)$ and a tuple of N strings $(\text{Blt}_1, \dots, \text{Blt}_N)$ made up of elements that can be either ballots cast by a voter or the special symbol \perp to indicate an abstention, the evaluation tally procedure computes what follows.
 1. For all $j \in [N], l \in [2]$,

$$m_j^l = \begin{cases} \perp & \text{if } \text{Blt}_j = \perp, \\ \perp & \text{if } \text{Blt}_j \neq \perp \wedge \mathcal{E}.\text{Decrypt}(\text{Ct}_{j,l}, \mathcal{E}.\text{Sk}_l) \notin \mathcal{M}, \\ \mathcal{E}.\text{Decrypt}(\text{Ct}_{j,l}, \mathcal{E}.\text{Sk}_l) & \text{otherwise.} \end{cases}$$

2. For all $l \in [2]$, compute $y_l = F(m_{1,l}, \dots, m_{N,l})$.
3. If $y_1 = y_2$ then set $y = y'$; otherwise set $y = \perp$.
4. Consider the following relation \mathbf{R}^{dec} in Fig. 4.

Henceforth, if a statement or proof for the relation \mathbf{R}^{dec} is satisfied with indices $i_1 = 1, i_2 = 2$ (resp. $i_1 \neq 1$ or $i_2 \neq 2$) we will say that the statement or the proof is in real mode (resp. trapdoor mode).

Run $\text{Prove}^{\text{dec}}$ on input $x \triangleq (\text{Blt}_1, \dots, \text{Blt}_N, \mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3, y)$ and the witness $(\mathcal{E}.\text{Sk}_1, \mathcal{E}.\text{Sk}_2, s_1, s_2, i_1 \triangleq 1, i_2 \triangleq 2)$ to compute a proof γ .

Relation $R^{\text{dec}}(x, w)$:

Instance: $x \triangleq (\text{Bl}_1, \dots, \text{Bl}_N, \mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3, y)$. (Recall that a ballot is set to \perp if the corresponding voter did not cast her vote.)

Witness: $w \triangleq (\mathcal{E}.\text{Sk}'_1, \mathcal{E}.\text{Sk}'_2, s_1, s_2, i_1, i_2)$, where the s_l 's are the randomness used to generate the secret- and public- keys pairs (that are known to the authority who set-up the system).

$R^{\text{dec}}(x, w) = 1$ if and only if the following condition holds.

The 2 secret-keys corresponding to the public-keys $\mathcal{E}.\text{Pk}_{i_1}, \mathcal{E}.\text{Pk}_{i_2}$ are constructed using honestly generated public- and secret-key pairs and are equal to $\mathcal{E}.\text{Sk}'_1, \mathcal{E}.\text{Sk}'_2$; and either $y = \perp$ or for all $l \in [2]$, $y = F(m_1^l, \dots, m_N^l)$ and for all $j \in [N]$, if $\text{Bl}_j \neq \perp$ then for all $l \in [2]$ the ciphertext Ct_{j, i_l} in Bl_j decrypts with the corresponding secret-key $\mathcal{E}.\text{Sk}_{i_l}$ to $m_j^{i_l} \in \mathcal{M}$; and for all $l \in [2]$, $m_j^l \triangleq \perp$ if $\text{Bl}_j = \perp$.

Precisely, $R^{\text{dec}}(x, w) = 1$ if and only if the following conditions hold. In the following, items (a) and (c) are not actually conditions that have to be checked but are steps needed to define the variables (note the use of " \triangleq ") $\mathcal{E}.\text{Pk}_{i_l}$'s, $\mathcal{E}.\text{Sk}_{i_l}$'s and $m_j^{i_l}$'s that are used in the checks (b) and (d).

- (a) For all $l \in [2]$, $(\mathcal{E}.\text{Pk}_{i_l}, \mathcal{E}.\text{Sk}_{i_l}) \triangleq \mathcal{E}.\text{Setup}(1^\lambda; s_l)$.
- (b) For all $l \in [2]$, $\mathcal{E}.\text{Sk}'_l = \mathcal{E}.\text{Sk}_{i_l}$.
- (c) For all $j \in [N], l \in [2]$,

$$m_j^{i_l} \triangleq \begin{cases} \perp & \text{if } \text{Bl}_j = \perp, \\ \perp & \text{if } \text{Bl}_j \neq \perp \wedge \mathcal{E}.\text{Decrypt}(\text{Ct}_{j, i_l}, \mathcal{E}.\text{Sk}_{i_l}) \notin \mathcal{M}, \\ \mathcal{E}.\text{Decrypt}(\text{Ct}_{j, i_l}, \mathcal{E}.\text{Sk}_{i_l}) & \text{otherwise.} \end{cases}$$

- (d) $(y = \perp) \vee$ (for all $l \in [2]$, $y = F(m_1^{i_l}, \dots, m_N^{i_l})$).

(Note that $\mathcal{E}.\text{Sk}'_1$ and $\mathcal{E}.\text{Sk}'_2$ do not have to necessarily correspond to the first two secret-keys.)

Fig. 4. Relation R^{dec}

5. Output (y, γ) .
- $\text{VerifyTally}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y, \gamma)$: on input the public-key Pk , a tuple of N strings that can be either ballots cast by a voter or the special symbol \perp to indicate an abstention, a claimed tally y and a claimed proof γ of the correctness of the computation:
 1. if $\gamma = \perp$ output \perp , otherwise output the decision of $\text{Verify}^{\text{dec,full}}((\text{Bl}_1, \dots, \text{Bl}_N, \mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3, y), \gamma)$.

Henceforth, for simplicity of notation we will omit the parameters of the scheme and we will just write **EVOTE**.

10.1 Correctness and Weak Verifiability of the Construction

Correctness. The (perfect) correctness of **EVOTE** follows from the perfect correctness of \mathcal{E} and the perfect completeness of NIWI^{dec} .

Weak verifiability.

Theorem 1 For all $N > 0$, all sets $\mathcal{M}, \Sigma \subset \{0, 1\}^*$, all tally functions $F : (\mathcal{M} \cup \{\perp\})^N \rightarrow \Sigma \cup \{\perp\}$, if \mathcal{E} is a perfectly correct PKE with unique secret-key (cf. Def. 6 and NIWI^{dec} is a (one-message) NIWI (cf. Def. 9) for the relation R^{dec} , then $\text{EVOTE}^{N, \mathcal{M}, \Sigma, F, \mathcal{E}, \text{NIWI}^{\text{dec}}}$ satisfies weak verifiability (cf. Def. 3).

Proof. We first prove that condition (1) of verifiability is satisfied.

Since VerifyBallot accepts any ballot, even invalid ones, we have to prove that for all $\text{Pk} \in \{0, 1\}^*$, all $\text{Bl}_1, \dots, \text{Bl}_N \in \{0, 1\}^* \cup \{\perp\}$, there exist $m_1, \dots, m_N \in \mathcal{M} \cup \{\perp\}$ such that for all $y \neq \perp$ and all γ in $\{0, 1\}^*$, it holds that: if $\text{VerifyTally}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y, \gamma) = 1$ then $y = F(m_1, \dots, m_N)$.

Henceforth, w.l.o.g, we let Pk and $\text{Bl}_1, \dots, \text{Bl}_N$ be arbitrary strings. First, we prove the following claim.

Claim 1 Given Pk and $(\text{Bl}_1, \dots, \text{Bl}_N)$, for every two pairs (y_0, γ_0) and (y_1, γ_1) such that $y_0, y_1 \neq \perp$, if $\text{VerifyTally}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y_0, \gamma_0) = \text{VerifyTally}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y_1, \gamma_1) = \text{OK}$ then $y_1 = y_2$.

Let $y_0, \gamma_0, y_1, \gamma_1$ be arbitrary strings in $\{0, 1\}^* \cup \{\perp\}$ such that $y_0, y_1 \neq \perp$. Suppose that $\text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y_0, \gamma_0) = \text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y_1, \gamma_1) = \text{OK}$, otherwise the claim is proven. The perfect soundness of NIWI^{dec} implies that, for all $b \in \{0, 1\}$, the proof γ_b is verified by some witness $w_b \triangleq (\mathcal{E}.\text{Sk}_1^b, \mathcal{E}.\text{Sk}_2^b, s_1^b, s_2^b, i_1^b, i_2^b)$.

Now, by the pigeon principle, there exists an index i^* such that one of the following cases holds.

1. $i^* = i_1^0 = i_2^1$. For all $b \in \{0, 1\}$, let $(m_1^{i^*, b}, \dots, m_N^{i^*, b})$ be the messages guaranteed by condition (iii) of relation R^{dec} for proof γ_b . Condition (i) for proof γ_0 (resp. γ_1) implies that the secret-key $\text{Sk}_1^{i^0}$ (resp. $\text{Sk}_2^{i^1}$) is honestly computed and thus, the unique secret-key property and the fact that it corresponds to $\mathcal{E}.\text{Pk}_{i_1^0} = \mathcal{E}.\text{Pk}_{i^*}$ (resp. $\mathcal{E}.\text{Pk}_{i_2^1} = \mathcal{E}.\text{Pk}_{i^*}$ imply that for all $j \in [N]$, $\mathcal{E}.\text{Decrypt}(\text{Ct}_{j, i^*}, \mathcal{E}.\text{Sk}_1^{i^0}) = \mathcal{E}.\text{Decrypt}(\text{Ct}_{j, i^*}, \mathcal{E}.\text{Sk}_2^{i^1})$.

Furthermore, condition (ii) and (iii) for proof γ_0 (resp. γ_1) imply that for all $j \in [N]$, either for all $b \in \{0, 1\}$, $m_j^{i^*,b} = \perp$ or $m_j^{i^*,0} = \mathcal{E}.\text{Decrypt}(\text{Ct}_{j,i^*}, \mathcal{E}.\text{Sk}_1^0) \in \mathcal{M}$ (resp. $m_j^{i^*,1} = \mathcal{E}.\text{Decrypt}(\text{Ct}_{j,i^*}, \mathcal{E}.\text{Sk}_2^1) \in \mathcal{M}$).

Hence, for all $j \in [N]$, $m_j^{i^*,0} = m_j^{i^*,1} \in \mathcal{M} \cup \{\perp\}$. Now, condition (iv) for proof γ_0 (resp. γ_1) implies that either $y_0 = F(m_1^{i_1^0,0}, \dots, m_N^{i_N^0,0})$ or $y_0 = \perp$ (resp. either $y_1 = F(m_1^{i_1^1,1}, \dots, m_N^{i_N^1,1})$ or $y_1 = \perp$) and, as by hypothesis $y_0, y_1 \neq \perp$, it holds that $y_0 = y_1$.

(Here, it is where the “weakness“ arises. In fact, EVOTE cannot be proven (fully) verifiable since it could occur that $y_0 \neq y_1, y_0 = \perp$.)

2. $i^* = i_2^0 = i_1^1$. This case is symmetrical to the first one, having care of replacing i_1^0 with i_2^0 and i_2^1 with i_1^1 .
3. $i^* = i_1^0 = i_1^1$. This case is symmetrical to the first one, having care of replacing i_2^1 with i_1^1 .
4. $i^* = i_2^0 = i_2^1$. This case is symmetrical to the first one, having care of replacing i_1^0 with i_2^0 .

In all cases, we have that if $\text{VerifyTally}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y_0, \gamma_0) = \text{VerifyTally}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y_1, \gamma_1) = \text{OK}$ then $y_0 = y_1$ and the claim is proved.

Now, from the previous claim it follows that there exists a *unique* value y^* such that for all (y, γ) such that $y \neq \perp$, if $\text{VerifyTally}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y, \gamma) = \text{OK}$ then $y = y^*$ (1).

Moreover, it is easy to see that for all (y, γ) , if $\text{VerifyTally}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y, \gamma) = \text{OK}$, there exist messages $m_1, \dots, m_N \in \mathcal{M} \cup \{\perp\}$ such that $y = F(m_1, \dots, m_N)$ (2).

Now, we have two mutually exclusive cases.

- For all (y, γ) such that $y \neq \perp$, $\text{VerifyTally}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y, \gamma) = \perp$. Then, letting m_1, \dots, m_N in the statement of the theorem be arbitrary messages in $\mathcal{M} \cup \{\perp\}$, the statement is verified with respect to Pk and $\text{Bl}_1, \dots, \text{Bl}_N$.
- There exists (y', γ) such that $y' \neq \perp$ and $\text{VerifyTally}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y', \gamma) = \text{OK}$. In this case, (2) implies that there exist $m'_1, \dots, m'_N \in \mathcal{M} \cup \{\perp\}$ such that $y' = F(m'_1, \dots, m'_N)$ (3). Hence, (1) and (3) together imply that $y^* = F(m'_1, \dots, m'_N)$ (4).

Therefore, for all (y, γ) such that $y \neq \perp$, if $\text{VerifyTally}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y, \gamma) = \text{OK}$ then (by (1)) $y = y^* =$ (by (4)) $= F(m'_1, \dots, m'_N)$.

Then, for $m_1 \triangleq m'_1, \dots, m_N \triangleq m'_N$, the statement of condition (1) of weak verifiability is verified with respect to Pk and $\text{Bl}_1, \dots, \text{Bl}_N$.

In both cases, for $m_1 \triangleq m'_1, \dots, m_N \triangleq m'_N$, the statement of condition (1) of weak verifiability is verified with respect to Pk and $\text{Bl}_1, \dots, \text{Bl}_N$.

As Pk and $\text{Bl}_1, \dots, \text{Bl}_N$ are arbitrary strings, the statement of condition (1) of weak verifiability is proven.

It is also easy to check that condition (2) of weak verifiability is satisfied; this follows straightforward from the perfect soundness of $\text{NIWI}^{\text{dec,full}}$ observing the following. The authority always proves that the the public-key of the PKE

is honestly generated and so, by the perfect correctness of the PKE, an honestly computed ballot for message m for the j -th voter will be decrypted to m (since an honestly computed ballot, by definition, is constituted by three ciphertexts that encrypt the same message), and thus if the claimed tally y is different from \perp (i.e., if the evaluation of the tally function is equal for all indices), then y has to be compatible with m at index j (cf. Def. 1).

10.2 Weak Privacy of the Construction

Theorem 2 For all $N > 0$, all sets $\mathcal{M}, \Sigma \subset \{0, 1\}^*$, all tally functions $F : (\mathcal{M} \cup \{\perp\})^N \rightarrow \Sigma \cup \{\perp\}$, if \mathcal{E} is a perfectly correct PKE with unique secret-key (cf. Def. 6 and NIWI^{dec} is a (one-message) NIWI (cf. Def. 9) for the relation R^{dec} , then $\text{EVOTE}^{N, \mathcal{M}, \Sigma, F, \mathcal{E}, \text{NIWI}^{\text{dec}}}$ is wIND-Secure (cf. Def. 5).

Proof. Let \mathcal{A} be an arbitrary adversary against the wIND-Security of EVOTE (henceforth, for simplicity, we will omit its parameters). We prove that $\text{Adv}_{\mathcal{A}}^{\text{EVOTE}, \text{WeakPriv}}(1^\lambda) \leq \nu(1^\lambda)$ for some negligible function $\nu(\lambda)$.

We prove that by means of a series of hybrid experiments. We invite the reader to refer to Table 3.3 for a pictorial explanation of the experiments as explained in Section 3.2. In the table, for simplicity, we omitted the indices k for the experiments H_2^k 's, H_4^k 's, H_6^k 's presented below, and so hybrid H_2 (resp. H_4, H_6) in the table would correspond to hybrid H_2^N (resp. H_4^N, H_6^N) presented below.

Hybrid H_1 . Experiment H_1 is equal to the experiment $\text{WeakPriv}_{\mathcal{A}}^{N, \mathcal{M}, \Sigma, F, \text{EVOTE}}$ except that the challenger sets $b \triangleq 0$.

Hybrid H_2^k , for $k = 0, \dots, N$. For all $k = 0, \dots, N$, experiment H_2^k is identical to experiment H_1 except that for all $j = 1, \dots, k$, the challenger computes $\text{Ct}_{k,3}$ to be encryption of $m_{1,k}$. Note that H_2^0 is identical to H_1 .

Claim 2 For all $k = 1, \dots, N$ the advantage of \mathcal{A} in distinguishing H_2^{k-1} from H_2^k is negligible.

Proof. Suppose toward a contradiction that \mathcal{A} have instead non-negligible advantage $\epsilon(\lambda)$. We construct an adversary \mathcal{B} that has advantage at most $\epsilon(\lambda)$ against the IND-CPA security of \mathcal{E} .

\mathcal{B} receives from the challenger of IND-CPA a public-key pk and sets $\text{Pk}_3 \triangleq \text{pk}$. For $l \in [2]$, \mathcal{B} runs $\mathcal{E}.\text{Setup}$ to compute $(\mathcal{E}.\text{Pk}_l, \mathcal{E}.\text{Sk}_l)$ and runs \mathcal{A} on input $\text{Pk} \triangleq (\mathcal{E}.\text{Pk}_1, \mathcal{E}.\text{Pk}_2, \mathcal{E}.\text{Pk}_3)$.

\mathcal{A} outputs two tuples $(m_{0,1}, \dots, m_{0,N})$ and $(m_{1,1}, \dots, m_{1,N})$ (and a set S empty for the wIND-Security game).

\mathcal{B} returns $(m_{0,k}, m_{1,k})$ as its pair of challenge messages to the IND-CPA challenger that in turn returns to \mathcal{B} the challenge ciphertext ct^* .

\mathcal{B} computes $\text{Bl}_k \triangleq (\text{Ct}_{k,1}, \text{Ct}_{k,2}, \text{ct}^*)$ by encrypting $m_{0,j}$ in $\text{Ct}_{k,1}$ and $\text{Ct}_{k,2}$. \mathcal{B} can compute the ballots Bl_j for all $j \in [N], j \neq k$ exactly as the challenger in both experiments would do.

\mathcal{B} computes y as in the previous experiment (i.e., running EvalTally over $(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N)$) and uses the 2 secret-keys $\mathcal{E}.\text{Sk}_1, \mathcal{E}.\text{Sk}_2$ to compute a proof γ exactly as the challenger in both experiments would do.

\mathcal{B} restarts \mathcal{A} on input the so computed ballots along with (y, γ) and returns the output of \mathcal{A} .

It is easy to see that if ct^* is an encryption of $m_{0,k}$ then \mathcal{B} simulates experiment H_2^{k-1} and if ct^* is an encryption of $m_{1,k}$, then \mathcal{B} simulates experiment H_2^k .

Therefore, \mathcal{B} has probability at least $\epsilon(\lambda)$ of winning the IND-CPA game, a contradiction.

Hybrid H_3 . Experiment H_3 is identical to experiment H_2^N except that the challenger computes the proof γ with indices $(1, 3)$ and secret-keys Sk_1, Sk_3 (precisely, with the randomness used to compute them but henceforth, for simplicity, we omit this detail).

Claim 3 The advantage of \mathcal{A} in distinguishing H_2^N from H_3 is negligible.

Proof. This follows straight-forward from the WI of NIWI^{dec} observing that both the randomness used to compute Sk_1, Sk_2 and the randomness used to compute Sk_1, Sk_3 constitute both valid witnesses for $(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N)$.

Hybrid H_4^k , for $k = 0, \dots, N$. For all $k = 0, \dots, N$, experiment H_4^k is identical to experiment H_3 except that for all $j = 1, \dots, k$, the challenger computes $\text{Ct}_{k,2}$ to be encryption of $m_{1,k}$. Note that H_4^0 is identical to H_3 .

Claim 4 For all $k = 1, \dots, N$ the advantage of \mathcal{A} in distinguishing H_4^{k-1} from H_4^k is negligible.

Proof. The proof is exactly symmetrical to the one for Claim 2 except that the third index is swapped with the second index.

Hybrid H_5 . Experiment H_5 is identical to experiment H_4^N except that the challenger computes the proof γ with indices $(2, 3)$ and secret-keys Sk_2, Sk_3 .

Claim 5 The advantage of \mathcal{A} in distinguishing H_4^N from H_5 is negligible.

Proof. This follows straight-forward from the WI of NIWI^{dec} observing that both the randomness used to compute Sk_1, Sk_3 and the randomness used to compute Sk_2, Sk_3 constitute both valid witnesses for $(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N)$.

Hybrid H_6^k , for $k = 0, \dots, N$. For all $k = 0, \dots, N$, experiment H_6^k is identical to experiment H_5 except that for all $j = 1, \dots, k$, the challenger computes $\text{Ct}_{k,1}$ to be encryption of $m_{1,k}$. Note that H_6^0 is identical to H_5 .

Claim 6 For all $k = 1, \dots, N$ the advantage of \mathcal{A} in distinguishing H_6^{k-1} from H_6^k is negligible.

Proof. The proof is exactly symmetrical to the one for Claim 2 except that the third index is swapped with the first index.

Hybrid H_7 ,. Experiment H_7 is identical to experiment H_6^N except that the challenger computes the proof γ with indices $(1, 2)$ and secret-keys Sk_1, Sk_2 .

Claim 7 The advantage of \mathcal{A} in distinguishing H_6^N from H_7 is negligible.

Proof. This follows straight-forward from the WI of NIWI^{dec} observing that both the randomness used to compute Sk_1, Sk_2 and the randomness used to compute Sk_2, Sk_3 constitute both valid witnesses for $(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N)$.

Now, by observing that experiment H_1 (resp. H_7) is identical to the experiment $\text{WeakPriv}_{\mathcal{A}}^{N, \mathcal{M}, \Sigma, F, \text{EVOTE}}$ except that the challenger sets $b = 0$ (resp. $b = 1$), it follows that $\text{WeakPriv}_{\mathcal{A}}^{N, \mathcal{M}, \Sigma, F, \text{EVOTE}}$ equals at most the sum of the advantages of \mathcal{A} in distinguishing the previous hybrids and, since N is a constant, such advantage is negligible and the theorem is proven.

Corollary 3 If the Decision Linear assumption (see Section 9) holds, then there exists a weakly verifiable eVote.

Proof. Boneh *et al.* [BBS04] show the existence of a PKE with perfect correctness and unique secret-key from Decision Linear, and Groth *et al.* [GOS06] show the existence of (one-message) NIWI (with perfect soundness) for all languages in NP from the Decision Linear assumption.

Then, combining theorems 1 and 2, the corollary follows.

11 Our (Fully) Verifiable eVote

In this Section, we present our eVote scheme $\text{EVOTE}_{\text{full}}$ that is IND-Secure and (fully) verifiable.

Definition 11 [$\text{EVOTE}_{\text{full}}$] Let $\mathcal{E} = (\mathcal{E}.\text{Setup}, \mathcal{E}.\text{Encrypt}, \mathcal{E}.\text{Decrypt})$ be a PKE scheme with perfectly correctness and unique secret-key (see Def. 6), Com a perfectly binding commitment, and let $\text{NIWI}^{\text{dec}, \text{full}} \triangleq (\text{Prove}^{\text{dec}, \text{full}}, \text{Verify}^{\text{dec}, \text{full}})$ and $\text{NIWI}^{\text{enc}, \text{full}} \triangleq (\text{Prove}^{\text{enc}, \text{full}}, \text{Verify}^{\text{enc}, \text{full}})$ be two NIWI proof systems, respectively, for the relations $\text{R}^{\text{dec}, \text{full}}$ and $\text{R}^{\text{enc}, \text{full}}$ to be specified later.

We define a $(N, \mathcal{M}, \Sigma, F)$ -eVote

$\text{EVOTE}_{\text{full}}^{N, \mathcal{M}, \Sigma, F, \mathcal{E}, \text{Com}, \text{NIWI}^{\text{dec}, \text{full}}} = (\text{Setup}_{\text{full}}, \text{Cast}_{\text{full}}, \text{VerifyBallot}_{\text{full}}, \text{EvalTally}_{\text{full}}, \text{VerifyTally}_{\text{full}})$ as follows.

- $\text{Setup}_{\text{full}}(1^\lambda)$: on input the security parameter in unary, compute as follows.
 1. Choose randomness $r \leftarrow \{0, 1\}^\lambda$ and for all $l \in [2]$, randomness $s_l \leftarrow \{0, 1\}^\lambda$.
 2. For all $l \in [3]$, run $(\mathcal{E}.\text{Pk}_l, \mathcal{E}.\text{Sk}_l) = \mathcal{E}.\text{Setup}(1^\lambda; s_l)$ with randomness s_l , and set $Z = \text{Com}(1; r)$.
 3. Output $\text{Pk} \triangleq (\mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3, Z)$ and $\text{Sk} \triangleq (\mathcal{E}.\text{Sk}_1, \mathcal{E}.\text{Sk}_2, s_1, s_2, r)$.
 (Actually, as the randomness for the setup of a PKE uniquely determines the secret-key, it would be sufficient to just include the s_l 's in Sk .)
- $\text{Cast}_{\text{full}}(\text{Pk}, j, v)$: on input the public-key $\text{Pk} \triangleq (\mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3)$, the index $j \in [N]$ of the voter, a vote v , compute as follows.
 1. For all $l \in [3]$, choose randomness $r_l \leftarrow \{0, 1\}^\lambda$.
 2. For all $l \in [3]$, compute $\text{Ct}_{j,l} = \mathcal{E}.\text{Encrypt}(\mathcal{E}.\text{Pk}_l, v; r_l)$.
 3. Consider the following relation $\text{R}^{\text{enc,full}}$ in Fig. 5.

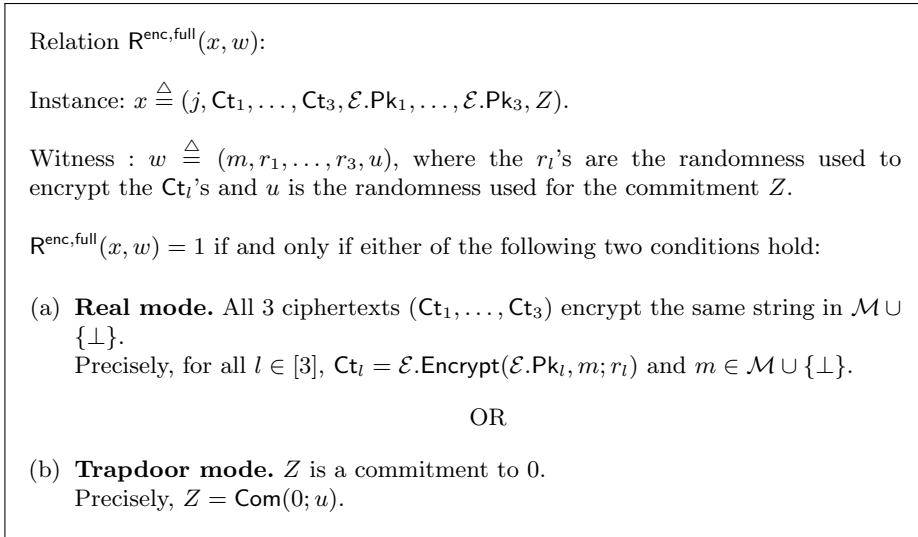


Fig. 5. Relation $\text{R}^{\text{enc,full}}$

- Run $\text{Prove}^{\text{enc,full}}$ on input $x \triangleq (j, \text{Ct}_1, \dots, \text{Ct}_3, \mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3, Z)$ and the witness (r_1, \dots, r_3) to compute a proof π .
- Output $\text{Bl}_j \triangleq (\text{Ct}_{j,1}, \dots, \text{Ct}_{j,3}, \pi_j)$.
- $\text{VerifyBallot}_{\text{full}}(\text{Pk}, j, \text{Bl}_j)$: on input the public-key $\text{Pk} \triangleq (\mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3)$, the index $j \in [N]$ of the voter and a ballot $\text{Bl}_j \triangleq (\text{Ct}_1, \dots, \text{Ct}_3, \pi)$, output the decision of $\text{Verify}^{\text{enc,full}}((\text{Pk}, j, \text{Bl}_j), \pi)$.

- $\text{EvalTally}_{\text{full}}(\text{Pk}, \text{Sk}, \text{Bl}_1, \dots, \text{Bl}_N)$: on input the public-key $\text{Pk} \triangleq (\mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3)$, the secret-key $\text{Sk} \triangleq (\mathcal{E}.\text{Sk}_1, \mathcal{E}.\text{Sk}_2, s_1, s_2)$ and a tuple of N strings $(\text{Bl}_1, \dots, \text{Bl}_N)$ made up of elements that can be either ballots cast by a voter or the special symbol \perp to indicate an abstention, the evaluation tally procedure computes what follows.

1. For all $j \in [N]$, it checks whether $\text{VerifyBallot}(\text{Pk}, \text{Bl}_j) = \perp$ and in this case it sets $\text{Bl}_j = \perp$ (i.e., it changes its previous value).

If the check failed for all $j \in [N]$, then output $(y \triangleq \perp, \gamma \triangleq \perp)$ without proceeding to next step.

2. For all $j \in [N], l \in [2]$,

$$m_j^l = \begin{cases} \perp & \text{if } \text{Bl}_j = \perp, \\ \perp & \text{if } \text{Bl}_j \neq \perp \wedge \mathcal{E}.\text{Decrypt}(\text{Ct}_{j,l}, \mathcal{E}.\text{Sk}_l) \notin \mathcal{M}, \\ \mathcal{E}.\text{Decrypt}(\text{Ct}_{j,l}, \mathcal{E}.\text{Sk}_l) & \text{otherwise.} \end{cases}$$

3. For all $l \in [2]$, compute $y_l = F(m_{1,l}, \dots, m_{N,l})$.
4. If $y_1 = y_2$ then set $y = y'$.
5. Consider the following relation $\text{R}^{\text{dec}, \text{full}}$ in Fig. 6.

(The reader may have noticed that the value Z in the instance of the statement may be completely discarded (as it is never used) and that the relation $\text{R}^{\text{dec}, \text{full}}$ is equivalent to the relation R^{dec} used for our weakly verifiable eVote but for exigency of exposition we prefer to distinguish among them.)

Henceforth, if a statement or proof for the relation $\text{R}^{\text{dec}, \text{full}}$ is satisfied with indices $i_1 = 1, i_2 = 2$ (resp. $i_1 \neq 1$ or $i_2 \neq 2$) we will say that the statement or the proof is in real mode (resp. trapdoor mode).

Run $\text{Prove}^{\text{dec}, \text{full}}$ on input $x \triangleq (\text{Bl}_1, \dots, \text{Bl}_N, \mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3, y)$ and the witness $(\mathcal{E}.\text{Sk}_1, \mathcal{E}.\text{Sk}_2, s_1, s_2, i_1 \triangleq 1, i_2 \triangleq 2)$ to compute a proof γ .

6. Output (y, γ) .

- $\text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y, \gamma)$: on input the public-key Pk , a tuple of N strings that can be either ballots cast by a voter or the special symbol \perp to indicate an abstention, a claimed tally y and a claimed proof γ of the correctness of the computation: If $y = \perp$, then checks whether all Bl_j 's are equal to \perp , otherwise output the decision of $\text{Verify}^{\text{dec}, \text{full}}((\text{Bl}_1, \dots, \text{Bl}_N, (\mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3), y), \gamma)$, after having replaced the Bl_j 's with \perp when $\text{VerifyBallot}(\text{Pk}, \text{Bl}_j) = \perp$.

Precisely, the algorithm computes as follows:

1. For all $j \in [N]$, if $\text{VerifyBallot}(\text{Pk}, \text{Bl}_j) = \perp$, set $\text{Bl}_j = \perp$.
2. If $y \neq \perp$, then output the decision of $\text{Verify}^{\text{dec}, \text{full}}((\text{Bl}_1, \dots, \text{Bl}_N, \mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3, y), \gamma)$.
3. If $y = \perp$, then:

If for all $j \in [N], \text{Bl}_j = \perp$ output OK, otherwise output \perp .

Henceforth, for simplicity of notation we will omit the parameters of the scheme and we will just write $\text{EVOTE}_{\text{full}}$.

Relation $R^{\text{dec,full}}(x, w)$:

Instance: $x \triangleq (\text{Bl}_1, \dots, \text{Bl}_N, \mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3, Z, y)$. (Recall that a ballot is set to \perp if the corresponding voter either did not cast her vote or her ballot is not accepted by the verification ballot algorithm.)

Witness: $w \triangleq (\mathcal{E}.\text{Sk}'_1, \mathcal{E}.\text{Sk}'_2, s_1, s_2, i_1, i_2, r)$, where the s_l 's are the randomness used to generate the secret- and public- keys pairs and r is the randomness used to compute the commitment Z (that are all known to the authority who set-up the system).

$R^{\text{dec}}(x, w) = 1$ if and only if the following condition holds:

The 2 secret-keys corresponding to the public-keys $\mathcal{E}.\text{Pk}_{i_1}, \mathcal{E}.\text{Pk}_{i_2}$ are constructed using honestly generated public- and secret-key pairs and are equal to $\mathcal{E}.\text{Sk}'_1, \mathcal{E}.\text{Sk}'_2$; and either $y = \perp$ or for all $l \in [2]$, $y = F(m_1^l, \dots, m_N^l)$ and for all $j \in [N]$, if $\text{Bl}_j \neq \perp$ then for all $l \in [2]$ the ciphertext Ct_{j,i_l} in Bl_j decrypts with the corresponding secret-key $\mathcal{E}.\text{Sk}_{i_l}$ to $m_j^{i_l} \in \mathcal{M}$; and for all $l \in [2]$, $m_j^l \triangleq \perp$ if $\text{Bl}_j = \perp$.

Precisely, $R^{\text{dec,full}}(x, w) = 1$ if and only if the following conditions hold. In the following, items (a) and (c) are not actually conditions that have to be checked but are steps needed to define the variables (note the use of " \triangleq ") $\mathcal{E}.\text{Pk}_{i_l}$'s, $\mathcal{E}.\text{Sk}_{i_l}$'s and $m_j^{i_l}$'s that are used in the checks (b) and (d).

- (a) For all $l \in [2]$, $(\mathcal{E}.\text{Pk}_{i_l}, \mathcal{E}.\text{Sk}_{i_l}) \triangleq \mathcal{E}.\text{Setup}(1^\lambda; s_l)$.
- (b) For all $l \in [2]$, $\mathcal{E}.\text{Sk}'_l = \mathcal{E}.\text{Sk}_{i_l}$.
- (c) For all $j \in [N], l \in [2]$,

$$m_j^{i_l} \triangleq \begin{cases} \perp & \text{if } \text{Bl}_j = \perp, \\ \perp & \text{if } \text{Bl}_j \neq \perp \wedge \mathcal{E}.\text{Decrypt}(\text{Ct}_{j,i_l}, \mathcal{E}.\text{Sk}_{i_l}) \notin \mathcal{M}, \\ \mathcal{E}.\text{Decrypt}(\text{Ct}_{j,i_l}, \mathcal{E}.\text{Sk}_{i_l}) & \text{otherwise.} \end{cases}$$

- (d) For all $l \in [2]$, $y = F(m_1^{i_l}, \dots, m_N^{i_l})$.

(Note that $\mathcal{E}.\text{Sk}'_1$ and $\mathcal{E}.\text{Sk}'_2$ do not have to necessarily correspond to the first two secret-keys.)

Fig. 6. Relation $R^{\text{dec,full}}$

11.1 Correctness and (Full) Verifiability of the Construction

Correctness. Condition (1) of (perfect) correctness of $\text{EVOTE}_{\text{full}}$ follows from the perfect correctness of \mathcal{E} and the perfect completeness of $\text{NIWI}^{\text{dec,full}}$ and $\text{NIWI}^{\text{enc,full}}$. Condition (2) follows analogously observing the following. For all Pk honestly computed it holds that $\text{Pk} = (\text{Pk}_1, \dots, \text{Pk}_2, \text{Pk}_3, Z)$ for some $\text{Pk}_1, \text{Pk}_2, \text{Pk}_3$ and Z and Z is a commitment to 1. Thus, relation $\text{R}^{\text{enc,full}}$ implies that if there exists a proof π and a statement $x \triangleq (j, \text{Ct}_1, \dots, \text{Ct}_3, \text{Pk}_1, \dots, \text{Pk}_3, Z)$ such that VerifyBallot accepts (x, π) , then it must be the case that $\text{Ct}_1, \dots, \text{Ct}_3$ encrypt the same string in $\mathcal{M} \cup \{\perp\}$.

Therefore, for all $\text{Bl}_1, \dots, \text{Bl}_N$, if for all $j \in [N]$, Bl'_j is equal to Bl_j if Bl_j is accepted by $\text{VerifyBallot}_{\text{full}}$ or is replaced by \perp otherwise, then it holds that if $(y, \gamma) \triangleq \text{EvalTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N)$, then $y = F(m_1, \dots, m_N)$, where for all $j \in [N]$, m_j is the string encrypted in the first two ciphertexts of Bl_j if Bl_j is accepted by $\text{VerifyBallot}_{\text{full}}$ or \perp if it is refused.

Then, it is easy to see that $\text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y, \gamma) = \text{OK}$.

(Full) verifiability.

Theorem 4 For all $N > 0$, all sets $\mathcal{M}, \Sigma \subset \{0, 1\}^*$, all tally functions $F : (\mathcal{M} \cup \{\perp\})^N \rightarrow \Sigma \cup \{\perp\}$, if \mathcal{E} is a perfectly correct PKE with unique secret-key (cf. Def. 6, Com is a PPT algorithm, and $\text{NIWI}^{\text{dec,full}}$ and $\text{NIWI}^{\text{enc,full}}$ are (one-message) NIWIs (cf. Def. 9), respectively, for the relations $\text{R}^{\text{dec,full}}$ and $\text{R}^{\text{enc,full}}$, then $\text{EVOTE}_{\text{full}}^{N, \mathcal{M}, \Sigma, F, \mathcal{E}, \text{Com}, \text{NIWI}^{\text{dec,full}}}$ satisfies (fully) verifiability (cf. Def. 3).

Proof. We first prove that condition (1) of verifiability is satisfied. We have to prove that for all $\text{Pk} \in \{0, 1\}^*$, all $\text{Bl}_1, \dots, \text{Bl}_N \in \{0, 1\}^* \cup \{\perp\}$ such that for all $j \in [N]$ either $\text{Bl}_j = \perp$ or $\text{VerifyBallot}_{\text{full}}(\text{Pk}, j, \text{Bl}_j) = \text{OK}$, then there exist $m_1, \dots, m_N \in \mathcal{M} \cup \{\perp\}$ such that for all $y, \gamma \in \{0, 1\}^*$, it holds that: if $\text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y, \gamma) = 1$ then $y = F(m_1, \dots, m_N)$.

Henceforth, w.l.o.g, we let Pk and $\text{Bl}_1, \dots, \text{Bl}_N$ be arbitrary strings such that for all $j \in [N]$ either $\text{Bl}_j = \perp$ or $\text{VerifyBallot}_{\text{full}}(\text{Pk}, j, \text{Bl}_j) = \text{OK}$.

First, we prove the following claim.

Claim 8 Given Pk and $(\text{Bl}_1, \dots, \text{Bl}_N)$, for every two pairs (y_0, γ_0) and (y_1, γ_1) , if $\text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y_0, \gamma_0) = \text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y_1, \gamma_1) = \text{OK}$ then $y_1 = y_2$.

For every (y_0, γ_0) and (y_1, γ_1) , we have two cases.

1. Either $y_0 = \perp$ and $y_1 \neq \perp$ or $y_1 = \perp$ and $y_0 \neq \perp$. Suppose w.l.o.g. that $y_0 = \perp$ and $y_1 \neq \perp$. The other case (i.e., $y_1 = \perp$ and $y_0 \neq \perp$) is symmetrical. By construction, for all y, γ , it holds that (A) if $\text{Bl}_1 = \dots = \text{Bl}_N = \perp$, then $\text{VerifyBallot}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y, \gamma) = \text{OK}$ if and only if $y = \perp$ and (B) if for some $j \in [N]$, $\text{Bl}_j \neq \perp$, then $\text{VerifyBallot}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, \perp, \gamma) = \perp$. We now have two cases.

- (a) $\text{Bl}_1 = \dots = \text{Bl}_N = \perp$. Suppose that $\text{VerifyBallot}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, \mathbf{y}_1, \gamma_1) = \text{VerifyBallot}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, \mathbf{y}_0, \gamma_0)$. Then, $\text{VerifyBallot}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, \mathbf{y}_1, \gamma_1) = \text{OK}$ and by (A) $y_1 = \perp$, a contradiction
- (b) It is not the case that $\text{Bl}_1 = \dots = \text{Bl}_N = \perp$. Then, by (B) $\text{VerifyBallot}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, \perp, \gamma) = \perp$, and the statement of the claim is verified with respect to (y_0, γ_0) and (y_1, γ_1) .
2. $y_0, y_1 \neq \perp$. Suppose that $\text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y_0, \gamma_0) = \text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y_1, \gamma_1) = \text{OK}$, otherwise the claim is proven. The perfect soundness of $\text{NIWI}_{\text{dec,full}}^{\text{dec,full}}$ implies that, for all $b \in \{0, 1\}$, the proof γ_b is verified by some witness $w_b \triangleq (\mathcal{E}.\text{Sk}_1^b, \mathcal{E}.\text{Sk}_2^b, s_1^b, s_2^b, i_1^b, i_2^b)$. Now, by the pigeon principle, there exists an index i^* such that one of the following cases holds.
- (a) $i^* = i_1^0 = i_2^1$. For all $b \in \{0, 1\}$, let $(m_1^{i^*,b}, \dots, m_N^{i^*,b})$ be the messages guaranteed by condition (iii) of relation $\text{R}^{\text{dec,full}}$ for proof γ_b . Condition (i) for proof γ_0 (resp. γ_1) implies that the secret-key Sk_1^0 (resp. Sk_2^1) is honestly computed and thus, the unique secret-key property and the fact that it corresponds to $\mathcal{E}.\text{Pk}_{i_1^0} = \mathcal{E}.\text{Pk}_{i^*}$ (resp. $\mathcal{E}.\text{Pk}_{i_2^1} = \mathcal{E}.\text{Pk}_{i^*}$) imply that for all $j \in [N]$, $\mathcal{E}.\text{Decrypt}(\text{Ct}_{j,i^*}, \mathcal{E}.\text{Sk}_1^0) = \mathcal{E}.\text{Decrypt}(\text{Ct}_{j,i^*}, \mathcal{E}.\text{Sk}_2^1)$. Furthermore, condition (ii) and (iii) for proof γ_0 (resp. γ_1) imply that for all $j \in [N]$, either for all $b \in \{0, 1\}$, $m_j^{i^*,b} = \perp$ or $m_j^{i^*,0} = \mathcal{E}.\text{Decrypt}(\text{Ct}_{j,i^*}, \mathcal{E}.\text{Sk}_1^0) \in \mathcal{M}$ (resp. $m_j^{i^*,1} = \mathcal{E}.\text{Decrypt}(\text{Ct}_{j,i^*}, \mathcal{E}.\text{Sk}_2^1) \in \mathcal{M}$). Hence, for all $j \in [N]$, $m_j^{i^*,0} = m_j^{i^*,1} \in \mathcal{M} \cup \{\perp\}$. Now, condition (iv) for proof γ_0 (resp. γ_1) implies that either $y_0 = F(m_1^{i^*,0}, \dots, m_N^{i^*,0})$ or $y_0 = \perp$ (resp. either $y_1 = F(m_1^{i^*,1}, \dots, m_N^{i^*,1})$ or $y_1 = \perp$) and, as by hypothesis $y_0, y_1 \neq \perp$, it holds that $y_0 = y_1$.
- (b) $i^* = i_2^0 = i_1^1$. This case is symmetrical to the first one, having care of replacing i_1^0 with i_2^0 and i_2^1 with i_1^1 .
- (c) $i^* = i_1^0 = i_1^1$. This case is symmetrical to the first one, having care of replacing i_2^1 with i_1^1 .
- (d) $i^* = i_2^0 = i_2^1$. This case is symmetrical to the first one, having care of replacing i_1^0 with i_2^0 .

In all cases, we have that if $\text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y_0, \gamma_0) = \text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N, y_1, \gamma_1) = \text{OK}$ then $y_0 = y_1$ and the claim is proved.

Now, from the previous claim it follows that there exists a *unique* value y^* such that for all (y, γ) , if $\text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y, \gamma) = \text{OK}$ then $y = y^*$ (1).

Moreover, it is easy to see that for all (y, γ) , if $\text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y, \gamma) = \text{OK}$, there exist messages $m_1, \dots, m_N \in \mathcal{M} \cup \{\perp\}$ such that $y = F(m_1, \dots, m_N)$ (2).

Now, we have two mutually exclusive cases.

- For all (y, γ) , $\text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y, \gamma) = \perp$. Then, letting m_1, \dots, m_N in the statement of the theorem be arbitrary messages in $\mathcal{M} \cup \{\perp\}$, the statement is verified with respect to Pk and $\text{Bl}_1, \dots, \text{Bl}_N$.

- There exists (y', γ) such that $\text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y', \gamma) = \text{OK}$. In this case, (2) implies that there exist $m'_1, \dots, m'_N \in \mathcal{M} \cup \{\perp\}$ such that $y' = F(m'_1, \dots, m'_N)$ (3). Hence, (1) and (3) together imply that $y^* = F(m'_1, \dots, m'_N)$ (4).
Therefore, for all (y, γ) , if $\text{VerifyTally}_{\text{full}}(\text{Pk}, \text{Bl}_1, \text{Bl}_N, y, \gamma) = \text{OK}$ then (by (1)) $y = y^* =$ (by (4)) $= F(m'_1, \dots, m'_N)$.
Then, for $m_1 \triangleq m'_1, \dots, m_N \triangleq m'_N$, the statement of condition (1) of verifiability is verified with respect to Pk and $\text{Bl}_1, \dots, \text{Bl}_N$.

In both cases, for $m_1 \triangleq m'_1, \dots, m_N \triangleq m'_N$, the statement of condition (1) of verifiability is verified with respect to Pk and $\text{Bl}_1, \dots, \text{Bl}_N$.

As Pk and $\text{Bl}_1, \dots, \text{Bl}_N$ are arbitrary strings, the condition (1) of verifiability is proven.

It is also easy to check that condition (2) of verifiability is satisfied; this follows straightforward from the perfect soundness of $\text{NIWI}^{\text{dec,full}}$ observing the following. The authority always proves that the the public-key of the PKE is honestly generated and so, by the perfect correctness of the PKE, an honestly computed ballot for message m for the j -th voter will be decrypted to m (this holds independently from the value committed in Z since an honestly computed ballot, by definition, is constituted by three ciphertexts that encrypt the same message), and thus the claimed tally y has to be compatible with m at index j (cf. Def. 1).

(In essence, the condition (2) is verified because the degree of freedom of the authority in creating a dishonest public-key is only in setting the commitment dishonestly, but this will not affect how the honest ballots will be decrypted and “counted“.)

Note that for the proof of the theorem above, the security of the commitment scheme Com is not needed (i.e., the theorem holds for any PPT algorithm Com , even if insecure).

11.2 Privacy of the Construction

Theorem 5 For all $N > 0$, all sets $\mathcal{M}, \Sigma \subset \{0, 1\}^*$, all tally functions $F : (\mathcal{M} \cup \{\perp\})^N \rightarrow \Sigma \cup \{\perp\}$, if \mathcal{E} is a perfectly correct PKE with unique secret-key (cf. Def. 6, Com is a perfectly hiding scheme (cf. Def. 7), and $\text{NIWI}^{\text{dec,full}}$ and $\text{NIWI}^{\text{enc,full}}$ are (one-message) NIWIs (cf. Def. 9), respectively, for the relations $\mathbf{R}^{\text{dec,full}}$ and $\mathbf{R}^{\text{enc,full}}$, then $\text{EVOTE}_{\text{full}}^{N, \mathcal{M}, \Sigma, F, \mathcal{E}, \text{Com}, \text{NIWI}^{\text{dec,full}}}$ is IND-Secure (cf. Def. 4).

Proof. Consider the following experiment $H_{\mathcal{A}}^Z(1^\lambda)$ between a challenger and \mathcal{A} (henceforth, often we omit the parameters).

Experiment H^Z . Experiment H^Z is equal to the experiment $\text{Priv}_{\mathcal{A}}^{N, \mathcal{M}, \Sigma, F, \text{EVOTE}_{\text{full}}}$ except that the challenger sets the commitment Z in the public-key to be a commitment to 0 instead of 1. Note that the output of the experiment is defined to

be a bit that is 1 if and only if all winning conditions are satisfied. Then, we have the following.

Claim 9 The probability P^0 that \mathcal{A} wins in the experiment $\text{Priv}_{\mathcal{A}}^{N, \mathcal{M}, \Sigma, F, \text{EVOTE}_{\text{full}}}$ is negligibly different from the probability P^1 that \mathcal{A} wins in game H^Z .

Proof. Suppose towards a contradiction the the difference between P_0 and P_1 is some non-negligible function $\epsilon(\lambda)$. We construct an adversary \mathcal{B} that breaks the computational hiding property of Com with probability $\geq \epsilon(\lambda)$.

\mathcal{B} receives as input a commitment com that is either a commitment for 0 or for 1. For $l \in [3]$, \mathcal{B} runs $\mathcal{E}.\text{Setup}(1^\lambda)$ to compute $(\mathcal{E}.\text{Pk}_l, \mathcal{E}.\text{Sk}_l)$ and forms the public-key for the eVote by setting $\text{Pk} \triangleq (\mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3, Z \triangleq \text{com})$.

\mathcal{B} can emulate the rest of the experiments (note that from this point onwards both experiments coincide) using the first two secret-keys and gets the output b' of \mathcal{A} . \mathcal{B} outputs 1 if and only if all winning conditions are satisfied.

By hypothesis, if com is a commitment to 0, the probability that \mathcal{B} outputs 1 equals the probability that \mathcal{A} wins in $\text{Priv}_{\mathcal{A}}^{N, \mathcal{M}, \Sigma, F, \text{EVOTE}_{\text{full}}}$, and if com is a commitment to 1, the probability that \mathcal{B} outputs 1 equals the probability that \mathcal{A} wins in H^Z .

Thus, the advantage of \mathcal{B} in breaking the computational hiding property of Com is $\epsilon(\lambda)$, a contradiction.

(Before continuing the proof we would like to remark a subtle point. In the previous claim, we implicitly supposed that the adversary \mathcal{B} be able to make all checks of the winning conditions efficiently.

This is possible if \mathcal{M} is efficiently enumerable and its cardinality, as well as the number of voters N , are *constant* in the security parameter.

This could seem like resorting to “complexity leveraging“ arguments and in fact one could ask if our proof would break down in the case that N and \mathcal{M} may depend on the security parameter. However, the whole proof can be generalized to the case of N and $|\mathcal{M}|$ polynomial in the security parameter using the following observation.

Let A be the event that \mathcal{A} submits challenges that satisfy the winning condition. Then, if the probability that \mathcal{A} wins in the $\text{Priv}_{\mathcal{A}}^{N, \mathcal{M}, \Sigma, F, \text{EVOTE}_{\text{full}}}$ is non-negligible, then it must the case that the event A occurs with non-negligible probability and, conditioned on it, \mathcal{A} wins with non-negligible probability as well.

Therefore, the rest of the proof would follow analyzing the probability that \mathcal{A} win in the next hybrid experiments conditioned under the occurrence of the event that, in such experiments, \mathcal{A} submit challenges satisfying the winning condition. As we will see now, a similar “conditioning“argument will be anyhow necessary for the rest of the proof.)

Let E^1 be the event that in experiment H^Z , \mathcal{A} submits as challenge two tuples $M_0 \triangleq (m_{0,1}, \dots, m_{0,N})$ and $M_1 \triangleq (m_{1,1}, \dots, m_{1,N})$, and a set $S \subset [N]$ such that there exists $j \in S$ such that $m_{0,j} = m_{1,j}$ and, letting $B \triangleq$

$m_{0,j} \triangleq (\text{Ct}_1, \dots, \text{Ct}_3)$ (supposing that can be parsed in a such way), it holds that $\text{VerifyBallot}(\text{Pk}, B) = \text{OK}$ but there exist $i_1, i_2 \in [3], i_1 \neq i_2$ such that $\mathcal{E}.\text{Decrypt}(\text{Ct}_{i_1}, \text{Sk}_{i_1}) \neq \mathcal{E}.\text{Decrypt}(\text{Ct}_{i_2}, \text{Sk}_{i_2})$.

(Note that E^0 and E^1 are events in different probability spaces but we can still study the difference of such probabilities, as we are going to do next.)

Claim 10 The probability that E^1 occurs is negligible.

Proof. Suppose towards a contradiction the the probability of occurrence of E^1 be some non-negligible function $\epsilon(\lambda)$. We construct an adversary \mathcal{B} that breaks the computational hiding property of Com with probability $\geq \epsilon(\lambda)$.

\mathcal{B} receives as input a commitment com that is either a commitment to 0 or to 1. For $l \in [3]$, \mathcal{B} runs $\mathcal{E}.\text{Setup}(1^\lambda)$ to compute $(\mathcal{E}.\text{Pk}_l, \mathcal{E}.\text{Sk}_l)$ and forms the public-key for the eVote by setting $\text{Pk} \triangleq (\mathcal{E}.\text{Pk}_1, \dots, \mathcal{E}.\text{Pk}_3, Z \triangleq \text{com})$.

\mathcal{B} can emulate the rest of the experiments (note that from this point onwards both experiments coincide) using the first two secret-keys and get the two tuples $M_0 \triangleq (m_{0,1}, \dots, m_{0,N})$ and $M_1 \triangleq (m_{1,1}, \dots, m_{1,N})$, and a set $S \subset [N]$.

For all $j \in S$, \mathcal{B} checks whether the following conditions are all satisfied: $m_{0,j} = m_{1,j}$ and, setting $B \triangleq m_{0,j}$, B can be parsed as $(\text{Ct}_1, \dots, \text{Ct}_3)$ and it holds that $\text{VerifyBallot}(\text{Pk}, B) = \text{OK}$ but there exist $i_1, i_2 \in [3], i_1 \neq i_2$ such that $\mathcal{E}.\text{Decrypt}(\text{Ct}_{i_1}, \text{Sk}_{i_1}) \neq \mathcal{E}.\text{Decrypt}(\text{Ct}_{i_2}, \text{Sk}_{i_2})$. If for some $j \in S$ the conditions are satisfied \mathcal{B} outputs 0, otherwise it outputs 1.

If com is a commitment to 1, by the perfect soundness of $\text{NIWI}^{\text{enc,full}}$ and the definition of relation $\text{R}^{\text{enc,full}}$ it cannot ever happen that the conditions above are satisfied for some $j \in S$ and so \mathcal{B} outputs 1 with probability 1.

On the other hand, if com is a commitment to 0, that probability that the conditions are satisfied for some $j \in [S]$ is exactly the probability of occurrence of E^1 and thus with probability ϵ , \mathcal{B} outputs 0 and with probability $1 - \epsilon$ outputs 1.

Thus, the advantage of \mathcal{B} in breaking the computational hiding property of Com is $\epsilon(\lambda)$, a contradiction.

From Claim 9 and 10 we now know that for some negligible function $\text{negl}(\cdot)$ the following equations hold:

$$|\Pr[\text{Priv} = 1] - \Pr[H^Z = 1]| \leq \text{negl}(\lambda), \quad (1)$$

$$\Pr[E^1] \leq \text{negl}(\lambda), \quad (2)$$

$$\Pr[H^Z = 1] = \Pr[H^Z = 1|E^1]\Pr[E^1] + \Pr[H^Z = 1|\bar{E}^1]\Pr[\bar{E}^1] \leq \text{negl} + \Pr[H^Z = 1|\bar{E}^1](1 - \text{negl}). \quad (3)$$

(Here and henceforth, we omit the parameters but it is meant that the experiments are parameterized by λ as well as $\text{negl}(\cdot)$.)

Thus, to show that $\Pr[\text{Priv} = 1]$ equals $1/2$ plus a negligible quantity, it is sufficient to show that $\Pr[H^Z = 1|\bar{E}^1]$ equals $1/2$ plus negligible as well.

We prove that by means of a series of hybrid experiments. The reader could still refer to Table 3.3 for a pictorial explanation but the reader is warned that the experiments in the table, though very similar conceptually to the next ones, correspond to the security reduction for the weakly verifiable eVote and moreover, in the following we will analyze the behavior of the adversary conditioned on the occurrence of the event \bar{E}^1 .

Hybrid H_1 . Experiment H_1 is equal to the experiment H^Z except that the challenger sets $b \triangleq 0$.

Hybrid H_2^k , for $k = 0, \dots, N$. For all $k = 0, \dots, N$, experiment H_2^k is identical to experiment H_1 except that for all $j = 1, \dots, k$ such that $j \notin S$, the challenger computes $\text{Ct}_{k,3}$ to be encryption of $m_{1,k}$. Note that H_2^0 is identical to H_1 .

Claim 11 For all $k = 1, \dots, N$, $|\Pr[H_2^{k-1} = 1|\bar{E}^1] - \Pr[H_2^k = 1|\bar{E}^1]|$ is negligible.

Proof. Suppose toward a contradiction that the difference in such probabilities is non-negligible function $\epsilon(\lambda)$. We construct an adversary \mathcal{B} that has advantage at most $\epsilon(\lambda)$ against the IND-CPA security of \mathcal{E} .

\mathcal{B} receives from the challenger of IND-CPA a public-key pk and sets $\text{Pk}_3 \triangleq \text{pk}$. For $l \in [2]$, \mathcal{B} runs $\mathcal{E}.\text{Setup}$ to compute $(\mathcal{E}.\text{Pk}_l, \mathcal{E}.\text{Sk}_l)$, computes $Z \leftarrow \text{Com}(0)$ and runs \mathcal{A} on input $\text{Pk} \triangleq (\mathcal{E}.\text{Pk}_1, \mathcal{E}.\text{Pk}_2, \mathcal{E}.\text{Pk}_3, Z)$.

\mathcal{A} outputs two tuples $(m_{0,1}, \dots, m_{0,N})$ and $(m_{1,1}, \dots, m_{1,N})$ (and a set S empty for the wIND-Security game).

If $k \in S$, \mathcal{B} returns $(0, 0)$ as its pair of challenge messages to the IND-CPA challenger that in turn returns to \mathcal{B} the challenge ciphertext ct^* . If $k \notin S$, \mathcal{B} returns $(m_{0,k}, m_{1,k})$ as its pair of challenge messages to the IND-CPA challenger that in turn returns to \mathcal{B} the challenge ciphertext ct^* .

If $k \in S$, \mathcal{B} sets Bl_j as the challenger in the real experiment would do; otherwise \mathcal{B} computes $\text{Bl}_k \triangleq (\text{Ct}_{k,1}, \text{Ct}_{k,2}, \text{ct}^*)$ by encrypting $m_{0,j}$ in $\text{Ct}_{k,1}$ and $\text{Ct}_{k,2}$. \mathcal{B} can compute the ballots Bl_j for all $j \in [N], j \neq k$ exactly as the challenger in both experiments would do.

\mathcal{B} computes y using EvalTally and uses the 2 secret-keys $\mathcal{E}.\text{Sk}_1, \mathcal{E}.\text{Sk}_2$ to compute a proof γ exactly as the challenger in both experiments would do.

\mathcal{B} restarts \mathcal{A} on input the so computed ballots along with (y, γ) and returns the output of \mathcal{A} .

It is easy to see that if ct^* is an encryption of $m_{0,k}$ and if $k \notin S$ then \mathcal{B} simulates experiment H_2^{k-1} and if ct^* is an encryption of $m_{1,k}$ and $k \notin S$, then \mathcal{B} simulates experiment H_2^k . Instead, if $k \in S$ the advantage of \mathcal{A} is clearly 0.

Therefore, \mathcal{B} has probability at least $\epsilon(\lambda)$ of winning the IND-CPA game, a contradiction.

Hybrid H_3 . Experiment H_3 is identical to experiment H_2^N except that the challenger computes the proof γ with indices $(1, 3)$ and secret-keys Sk_1, Sk_3 (precisely, with the randomness used to compute them but henceforth, for simplicity, we omit this detail).

Claim 12 $|\Pr [H_2^N = 1 | \bar{E}^1] - \Pr [H_3 = 1 | \bar{E}^1]|$ is negligible.

Proof. This follows from the WI of NIWI^{dec} observing that both the randomness used to compute Sk_1, Sk_2 and the randomness used to compute Sk_1, Sk_3 constitute both valid witnesses for $(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N)$ and by observing that if event \bar{E}^1 occurs any ballot in the set S either is replaced by \perp in both experiments if $\text{VerifyBallot}_{\text{full}}$ refuses it or it is decrypted in both experiments to the same values and thus the tally is computed identically in both experiments.

Hybrid H_4^k , for $k = 0, \dots, N$. For all $k = 0, \dots, N$, experiment H_4^k is identical to experiment H_3 except that for all $j = 1, \dots, k$ such that $j \notin S$, the challenger computes $\text{Ct}_{k,2}$ to be encryption of $m_{1,k}$. Note that H_4^0 is identical to H_3 .

Claim 13 For all $k = 1, \dots, N$, $|\Pr [H_4^{k-1} = 1 | \bar{E}^1] - \Pr [H_4^k = 1 | \bar{E}^1]|$ is negligible.

Proof. The proof is exactly symmetrical to the one for Claim 11 except that the third index is swapped with the second index.

Hybrid H_5 . Experiment H_5 is identical to experiment H_4^N except that the challenger computes the proof γ with indices $(2, 3)$ and secret-keys Sk_2, Sk_3 .

Claim 14 $|\Pr [H_4^N = 1 | \bar{E}^1] - \Pr [H_5 = 1 | \bar{E}^1]|$ is negligible.

Proof. This follows straight-forward from the WI of NIWI^{dec} observing that both the randomness used to compute Sk_1, Sk_3 and the randomness used to compute Sk_2, Sk_3 constitute both valid witnesses for $(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N)$.

Hybrid H_6^k , for $k = 0, \dots, N$. For all $k = 0, \dots, N$, experiment H_6^k is identical to experiment H_5 except that for all $j = 1, \dots, k$ such that $j \notin S$, the challenger computes $\text{Ct}_{k,1}$ to be encryption of $m_{1,k}$. Note that H_6^0 is identical to H_5 .

Claim 15 For all $k = 1, \dots, N$, $|\Pr [H_6^{k-1} = 1 | \bar{E}^1] - \Pr [H_6^k = 1 | \bar{E}^1]|$ is negligible.

Proof. The proof is exactly symmetrical to the one for Claim 11 except that the third index is swapped with the first index.

Hybrid H_7 . Experiment H_7 is identical to experiment H_6^N except that the challenger sets $b = 1$ (so that the winning condition will be computed differently) and computes the proof γ with indices $(1, 2)$ and secret-keys Sk_1, Sk_2 .

Claim 16 $|\Pr [H_6^N = 1 | \bar{E}^1] - \Pr [H_7 = 0 | \bar{E}^1]|$ is negligible.

Proof. This follows straight-forward from the WI of NIWI^{dec} observing that both the randomness used to compute Sk_1, Sk_2 and the randomness used to compute Sk_2, Sk_3 constitute both valid witnesses for $(\text{Pk}, \text{Bl}_1, \dots, \text{Bl}_N)$.

Note that according to the proof received, an adversary against NIWI can emulate experiment H^6 or H^7 , and return the output of \mathcal{A} . In the first case, the probability that \mathcal{A} output 0 is exactly $\Pr[H_6^N = 1 | \bar{E}^1]$ because the winning condition is computed with respect to $b = 0$, whereas in the second case it is $\Pr[H_7 = 0 | \bar{E}^1]$ because the winning condition is computed with respect to $b = 1$.

Now, observe that:

$$\begin{aligned}
& \Pr [H^Z = 1 | \bar{E}^1] = \\
& \Pr [H^Z = 1 | \bar{E}^1 \wedge b = 0] \Pr [b = 0] + \Pr [H^Z = 1 | \bar{E}^1 \wedge b = 1] \Pr [b = 1] = \\
& = 1/2 \cdot (\Pr [H^Z = 1 | \bar{E}^1 \wedge b = 0] + \Pr [H^Z = 1 | \bar{E}^1 \wedge b = 1]) = \\
& \text{(since } H_1 \text{ is identically distributed to } H^Z \text{ with bit } b = 0 \text{ and } H_7 \text{ to } H^Z \text{ with } b = 1) \\
& = 1/2 \cdot (\Pr [H_1 = 1 | \bar{E}^1] + \Pr [H_7 = 1 | \bar{E}^1]) = \\
& = 1/2 + 1/2 \cdot (\Pr [H_1 = 1 | \bar{E}^1] - \Pr [H_7 = 0 | \bar{E}^1]) = \\
& \text{(since } H_1 \text{ (resp. } H_3, H_5) \text{ is identically distributed to } H_2^0 \text{ (resp. } H_4^0, H_6^0)) \\
& = 1/2 + 1/2 \cdot \left(\sum_{k=0}^{N-1} (\Pr [H_2^k = 1 | \bar{E}^1] - \Pr [H_2^{k+1} = 1 | \bar{E}^1]) + (\Pr [H_2^N = 1 | \bar{E}^1] - \Pr [H_4^0 = 1 | \bar{E}^1]) \right) + \\
& \quad \sum_{k=0}^{N-1} (\Pr [H_4^k = 1 | \bar{E}^1] - \Pr [H_4^{k+1} = 1 | \bar{E}^1]) (\Pr [H_4^N = 1 | \bar{E}^1] - \Pr [H_6^0 = 1 | \bar{E}^1]) + \\
& \quad \sum_{k=0}^{N-1} (\Pr [H_6^k = 1 | \bar{E}^1] - \Pr [H_6^{k+1} = 1 | \bar{E}^1]) (\Pr [H_6^N = 1 | \bar{E}^1] - \Pr [H_7 = 0 | \bar{E}^1]) \leq \\
1 & \leq 1/2 + 1/2 \cdot \left| \left(\sum_{k=0}^{N-1} (\Pr [H_2^k = 1 | \bar{E}^1] - \Pr [H_2^{k+1} = 1 | \bar{E}^1]) + (\Pr [H_2^N = 1 | \bar{E}^1] - \Pr [H_4^0 = 1 | \bar{E}^1]) \right) + \right. \\
& \quad \left. \sum_{k=0}^{N-1} (\Pr [H_4^k = 1 | \bar{E}^1] - \Pr [H_4^{k+1} = 1 | \bar{E}^1]) (\Pr [H_4^N = 1 | \bar{E}^1] - \Pr [H_6^0 = 1 | \bar{E}^1]) + \right. \\
& \quad \left. \sum_{k=0}^{N-1} (\Pr [H_6^k = 1 | \bar{E}^1] - \Pr [H_6^{k+1} = 1 | \bar{E}^1]) (\Pr [H_6^N = 1 | \bar{E}^1] - \Pr [H_7 = 0 | \bar{E}^1]) \right| \leq \\
& \hspace{15em} \text{(by the triangle inequality)} \\
& \leq 1/2 + 1/2 \cdot \left(\sum_{k=0}^{N-1} |\Pr [H_2^k = 1 | \bar{E}^1] - \Pr [H_2^{k+1} = 1 | \bar{E}^1]| + |(\Pr [H_2^N = 1 | \bar{E}^1] - \Pr [H_4^0 = 1 | \bar{E}^1])| + \right. \\
& \quad \sum_{k=0}^{N-1} |\Pr [H_4^k = 1 | \bar{E}^1] - \Pr [H_4^{k+1} = 1 | \bar{E}^1]| |\Pr [H_4^N = 1 | \bar{E}^1] - \Pr [H_6^0 = 1 | \bar{E}^1]| + \\
& \quad \left. \sum_{k=0}^{N-1} |\Pr [H_6^k = 1 | \bar{E}^1] - \Pr [H_6^{k+1} = 1 | \bar{E}^1]| |\Pr [H_6^N = 1 | \bar{E}^1] - \Pr [H_7 = 0 | \bar{E}^1]| \right) \leq \\
& \hspace{15em} \text{(by Claims 11 - 16)} \\
& \hspace{15em} \leq 3k \cdot \text{negl}, \\
& \hspace{15em} (4)
\end{aligned}$$

where negl is the sum of the negligible functions guaranteed by Claims 11 - 16.

Finally, Claim 9 and Equations 2,3 and 4 imply that $\Pr[\text{Priv} = 1] \leq \nu$ for some negligible function ν and the theorem is proven.

Corollary 6 If the Decision Linear assumption (see Section 9) holds, then there exists a (fully) verifiable eVote.

Proof. Boneh *et al.* [BBS04] show the existence of a PKE with perfect correctness and unique secret-key from Decision Linear, and Groth *et al.* [GOS06] show the existence of (one-message) NIWI (with perfect soundness) for all languages in NP and of statistically binding commitments, both from the Decision Linear Assumption.

Then, combining Theorems 4 and 5, the corollary follows.

References

- Adi08. Ben Adida. Helios: Web-based open-audit voting. In *USENIX Security Symposium*, volume 17, pages 335–348, 2008.
- BBS04. Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In Matthew Franklin, editor, *Advances in Cryptology – CRYPTO 2004*, volume 3152 of *Lecture Notes in Computer Science*, pages 41–55. Springer, August 2004.
- BCG⁺15. David Bernhard, Véronique Cortier, David Galindo, Olivier Pereira, and Bogdan Warinschi. Sok: A comprehensive analysis of game-based ballot privacy definitions. In *2015 IEEE Symposium on Security and Privacy*, pages 499–516. IEEE, 2015.
- BDPA11. G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche. The KECCAK reference, 2011. <http://keccak.noekeon.org/>.
- BDSG⁺13. Nir Bitansky, Dana Dachman-Soled, Sanjam Garg, Abhishek Jain, Yael Tauman Kalai, Adriana López-Alt, and Daniel Wichs. Why “fiat-shamir for proofs” lacks a proof. In *Theory of Cryptography: 10th Theory of Cryptography Conference, TCC 2013, Tokyo, Japan, March 3-6, 2013.*, pages 182–201. Springer, 2013.
- Ben87. J. Benaloh. *Verifiable secret-ballot elections*. PhD thesis, Yale University, 1987.
- BF03. Dan Boneh and Matthew K. Franklin. Identity based encryption from the Weil pairing. *SIAM Journal on Computing*, 32(3):586–615, 2003.
- BFM88. Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *20th Annual ACM Symposium on Theory of Computing*, pages 103–112. ACM Press, May 1988.
- BFS16. Mihir Bellare, Georg Fuchsbauer, and Alessandra Scafuro. NIZKs with an untrusted CRS: Security in the face of parameter subversion. *Cryptology ePrint Archive*, Report 2016/372, 2016. <http://eprint.iacr.org/2016/372>. To appear in ASIACRYPT 2016.
- BGJS16. Sai Krishna Badrinarayanan, Vipul Goyal, Aayush Jain, and Amit Sahai. Verifiable functional encryption. *Cryptology ePrint Archive*, Report 2016/629, 2016. <http://eprint.iacr.org/2016/629>. To appear in ASIACRYPT 2016.

- Bit14. Nir Bitansky. *Getting inside the Adversary's Head: New Directions in Non-Black-Box Knowledge Extraction*. PhD thesis, Tel Aviv University, 2014.
- BOV03. Boaz Barak, Shien Jin Ong, and Salil P. Vadhan. Derandomization in cryptography. In Dan Boneh, editor, *Advances in Cryptology – CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 299–315. Springer, August 2003.
- BP15. Nir Bitansky and Omer Paneth. Zaps and non-interactive witness indistinguishability from indistinguishability obfuscation. In *Theory of Cryptography Conference*, pages 401–427. Springer, 2015.
- BR93. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In V. Ashby, editor, *ACM CCS 93: 1st Conference on Computer and Communications Security*, pages 62–73. ACM Press, November 1993.
- BSW11. Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In Yuval Ishai, editor, *TCC 2011: 8th Theory of Cryptography Conference*, volume 6597 of *Lecture Notes in Computer Science*, pages 253–273. Springer, March 2011.
- BT94. Josh Cohen Benaloh and Dwight Tuinstra. Receipt-free secret-ballot elections (extended abstract). In *26th Annual ACM Symposium on Theory of Computing*, pages 544–553. ACM Press, May 1994.
- CCC⁺09. David Chaum, Richard Carback, Jeremy Clark, Aleksander Essex, Stefan Popoveniuc, Ronald L. Rivest, Peter Y. A. Ryan, Emily Shen, Alan T. Sherman, and Poorvi L. Vora. Scantegrity II: end-to-end verifiability by voters of optical scan elections through confirmation codes. *IEEE Trans. Information Forensics and Security*, 4(4):611–627, 2009.
- CFG15. Véronique Cortier, Georg Fuchsbauer, and David Galindo. Beleniosrf: A strongly receipt-free electronic voting scheme. Cryptology ePrint Archive, Report 2015/629, 2015. <http://eprint.iacr.org/2015/629>.
- CG15. Pyrros Chaidos and Jens Groth. Making sigma-protocols non-interactive without random oracles. In *Public-Key Cryptography - PKC 2015 - 18th IACR International Conference on Practice and Theory in Public-Key Cryptography, Gaithersburg, MD, USA, March 30 - April 1, 2015, Proceedings*, pages 650–670, 2015.
- CGGI14. Véronique Cortier, David Galindo, Stéphane Glondu, and Malika Izabachène. Election verifiability for helios under weaker trust assumptions. In Mirosław Kutylowski and Jaideep Vaidya, editors, *ESORICS 2014: 19th European Symposium on Research in Computer Security, Part II*, volume 8713 of *Lecture Notes in Computer Science*, pages 327–344. Springer, September 2014.
- CGH98. Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In *30th Annual ACM Symposium on Theory of Computing*, pages 209–218. ACM Press, May 1998.
- CGK⁺16. Véronique Cortier, David Galindo, Ralf Kuesters, Johannes Mueller, and Tomasz Truderung. Verifiability notions for e-voting protocols. Cryptology ePrint Archive, Report 2016/287, 2016. <http://eprint.iacr.org/2016/287>.
- CGS97. Ronald Cramer, Rosario Gennaro, and Berry Schoenmakers. A secure and optimally efficient multi-authority election scheme. In Walter Fumy, editor, *Advances in Cryptology – EUROCRYPT'97*, volume 1233 of *Lecture Notes in Computer Science*, pages 103–118. Springer, May 1997.

- Cha81. David L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–90, 1981.
- CHK04. Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology – EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 207–222. Springer, May 2004.
- CKLM13. Melissa Chase, Markulf Kohlweiss, Anna Lysyanskaya, and Sarah Meiklejohn. Verifiable elections that scale for free. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013: 16th International Workshop on Theory and Practice in Public Key Cryptography*, volume 7778 of *Lecture Notes in Computer Science*, pages 479–496. Springer, February / March 2013.
- CPSV16. Michele Ciampi, Giuseppe Persiano, Luisa Siniscalchi, and Ivan Visconti. A transform for NIZK almost as efficient and general as the fiat-shamir transform without programmable random oracles. In *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part II*, pages 83–111, 2016.
- CS03a. Jan Camenisch and Victor Shoup. Practical verifiable encryption and decryption of discrete logarithms. In Dan Boneh, editor, *Advances in Cryptology – CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 126–144. Springer, August 2003.
- CS03b. Ronald Cramer and Victor Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM Journal on Computing*, 33(1):167–226, 2003.
- CS10. Veronique Cortier and Ben Smyth. Attacking and fixing helios: An analysis of ballot secrecy. Cryptology ePrint Archive, Report 2010/625, 2010. <http://eprint.iacr.org/2010/625>.
- CZZ⁺15. Nikos Chondros, Bingsheng Zhang, Thomas Zacharias, Panos Diamantopoulos, Stathis Maneas, Christos Patsonakis, Alex Delis, Aggelos Kiyayas, and Mema Roussopoulos. A distributed, end-to-end verifiable, internet voting system. *CoRR*, abs/1507.06812, 2015.
- DDO⁺01. Alfredo De Santis, Giovanni Di Crescenzo, Rafail Ostrovsky, Giuseppe Persiano, and Amit Sahai. Robust non-interactive zero knowledge. In *Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings*, pages 566–598, 2001.
- DFN06. Ivan Damgård, Nelly Fazio, and Antonio Nicolosi. Non-interactive zero-knowledge from homomorphic encryption. In Shai Halevi and Tal Rabin, editors, *TCC 2006: 3rd Theory of Cryptography Conference*, volume 3876 of *Lecture Notes in Computer Science*, pages 41–59. Springer, March 2006.
- DH76. Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.
- DJ01. Ivan Damgård and Mats Jurik. A generalisation, a simplification and some applications of Paillier’s probabilistic public-key system. In Kwangjo Kim, editor, *PKC 2001: 4th International Workshop on Theory and Practice in Public Key Cryptography*, volume 1992 of *Lecture Notes in Computer Science*, pages 119–136. Springer, February 2001.
- DJ03. Ivan Damgård and Mads Jurik. A length-flexible threshold cryptosystem with applications. In Reihaneh Safavi-Naini and Jennifer Seberry, editors,

- ACISP 03: 8th Australasian Conference on Information Security and Privacy*, volume 2727 of *Lecture Notes in Computer Science*, pages 350–364. Springer, July 2003.
- DKR06. Stéphanie Delaune, Steve Kremer, and Mark D. Ryan. Coercion-resistance and receipt-freeness in electronic voting, July 2006.
- DMP88. Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Non-interactive zero-knowledge proof systems. In Carl Pomerance, editor, *Advances in Cryptology – CRYPTO’87*, volume 293 of *Lecture Notes in Computer Science*, pages 52–72. Springer, August 1988.
- DN00. Cynthia Dwork and Moni Naor. Zaps and their applications. In *41st Annual Symposium on Foundations of Computer Science*, pages 283–293. IEEE Computer Society Press, November 2000.
- Fed12. Federal Agency on Technical Regulation and Metrology. Gost r 34.11-2012: Streebog hash function, 2012. <https://www.streebog.net>.
- FLS90. Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *31st Annual Symposium on Foundations of Computer Science*, pages 308–317. IEEE Computer Society Press, October 1990.
- FS87. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology – CRYPTO’86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer, August 1987.
- GGG⁺14. Shafi Goldwasser, S. Dov Gordon, Vipul Goyal, Abhishek Jain, Jonathan Katz, Feng-Hao Liu, Amit Sahai, Elaine Shi, and Hong-Sheng Zhou. Multi-input functional encryption. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 578–602. Springer, May 2014.
- GGH⁺13. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual Symposium on Foundations of Computer Science*, pages 40–49. IEEE Computer Society Press, October 2013.
- GGHZ16. Sanjam Garg, Craig Gentry, Shai Halevi, and Mark Zhandry. Functional encryption without obfuscation. In Eyal Kushilevitz and Tal Malkin, editors, *Theory of Cryptography: 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part II*, pages 480–511. Springer, 2016.
- GH07. Matthew Green and Susan Hohenberger. Blind identity-based encryption and simulatable oblivious transfer. In Kaoru Kurosawa, editor, *Advances in Cryptology – ASIACRYPT 2007*, volume 4833 of *Lecture Notes in Computer Science*, pages 265–282. Springer, December 2007.
- GIR16. Rosario Giustolisi, Vincenzo Iovino, and Peter Rønne. On the possibility of non-interactive voting in the public-key setting. In *Financial Cryptography and Data Security - FC 2016 International Workshops, BITCOIN, VOTING, and WAHC, Christ Church, Barbados, February 26, 2016, Revised Selected Papers*, 2016.
- GK03. Shafi Goldwasser and Yael Tauman Kalai. On the (in)security of the Fiat-Shamir paradigm. In *44th Annual Symposium on Foundations of Computer Science*, pages 102–115. IEEE Computer Society Press, October 2003.

- GKP⁺13. Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th Annual ACM Symposium on Theory of Computing*, pages 555–564. ACM Press, June 2013.
- GM84. Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.
- GO14. Jens Groth and Rafail Ostrovsky. Cryptography in the multi-string model. *Journal of Cryptology*, 27(3):506–543, July 2014.
- Gol01. Oded Goldreich. *Foundations of Cryptography: Basic Techniques*, volume 1. Cambridge University Press, Cambridge, UK, 2001.
- GOS06. Jens Groth, Rafail Ostrovsky, and Amit Sahai. Non-interactive zaps and new techniques for NIZK. In Cynthia Dwork, editor, *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 97–111. Springer, August 2006.
- Gro04. Jens Groth. Efficient maximal privacy in boardroom voting and anonymous broadcast. In *International Conference on Financial Cryptography*, pages 90–104. Springer, 2004.
- GS08. Jens Groth and Amit Sahai. Efficient non-interactive proof systems for bilinear groups. In Nigel P. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 415–432. Springer, April 2008.
- GVW12. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions via multi-party computation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 162–179. Springer, August 2012.
- HRZ10. Feng Hao, Peter Y. A. Ryan, and Piotr Zielinski. Anonymous voting by two-round public discussion. *IET Information Security*, 4(2):62–67, 2010.
- JCJ10. Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-resistant electronic elections. In *Towards Trustworthy Elections*, pages 37–63. Springer, 2010.
- Jou04. Antoine Joux. A one round protocol for tripartite Diffie-Hellman. *Journal of Cryptology*, 17(4):263–276, September 2004.
- Kal06. Yael Tauman Kalai. *Attacks on the Fiat-Shamir paradigm and program obfuscation*. PhD thesis, Massachusetts Institute of Technology, 2006.
- KRS10. Steve Kremer, Mark Ryan, and Ben Smyth. Election verifiability in electronic voting protocols. In *European Symposium on Research in Computer Security*, pages 389–404. Springer, 2010.
- KSRH12. Dalia Khader, Ben Smyth, Peter Y. A. Ryan, and Feng Hao. A fair and robust voting system by broadcast. In *5th International Conference on Electronic Voting 2012, (EVOTE 2012), Co-organized by the Council of Europe, Gesellschaft für Informatik and E-Voting.CC, July 11-14, 2012, Castle Hofen, Bregenz, Austria*, pages 285–299, 2012.
- KY02. Aggelos Kiayias and Moti Yung. Self-tallying elections and perfect ballot secrecy. In David Naccache and Pascal Paillier, editors, *PKC 2002: 5th International Workshop on Theory and Practice in Public Key Cryptography*, volume 2274 of *Lecture Notes in Computer Science*, pages 141–158. Springer, February 2002.

- KZZ15. Aggelos Kiayias, Thomas Zacharias, and Bingsheng Zhang. End-to-end verifiable elections in the standard model. In *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II*, pages 468–498, 2015.
- Lin15. Yehuda Lindell. An efficient transform from sigma protocols to NIZK with a CRS and non-programmable random oracle. In *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part I*, pages 93–109, 2015.
- Lip05. Helger Lipmaa. Secure electronic voting protocols. In Hossein Bidgoli, editor, *Handbook of Information Security, Volume 2, Information Warfare, Social, Legal, and International Issues and Security Foundations*, pages 647–657. John Wiley & Sons, Inc., 2005. Electronic edition available at <http://kodu.ut.ee/~lipmaa/papers/voting4hb.pdf>.
- LP09. Yehuda Lindell and Benny Pinkas. A proof of security of Yao’s protocol for two-party computation. *Journal of Cryptology*, 22(2):161–188, April 2009.
- MH96. Markus Michels and Patrick Horster. Some remarks on a receipt-free and universally verifiable mix-type voting scheme. In Kwangjo Kim and Tsutomu Matsumoto, editors, *Advances in Cryptology - ASIACRYPT’96*, volume 1163 of *Lecture Notes in Computer Science*, pages 125–132. Springer, November 1996.
- MN06. Tal Moran and Moni Naor. Receipt-free universally-verifiable voting with everlasting privacy. In Cynthia Dwork, editor, *Advances in Cryptology - CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 373–392. Springer, August 2006.
- Nao03. Moni Naor. On cryptographic assumptions and challenges (invited talk). In Dan Boneh, editor, *Advances in Cryptology - CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 96–109. Springer, August 2003.
- NY90. Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *22nd Annual ACM Symposium on Theory of Computing*, pages 427–437. ACM Press, May 1990.
- Riv06. Ronald L. Rivest. The threeballot voting system, 2006.
- RR06. Brian Randell and Peter Y. A. Ryan. Voting technologies and trust. In *IEEE Security and Privacy*, pages 50–56, 2006.
- RRI16. Peter Y. A. Ryan, Peter B. Rønne, and Vincenzo Iovino. Selene: Voting with transparent verifiability and coercion-mitigation. In *Financial Cryptography and Data Security - FC 2016 International Workshops, BITCOIN, VOTING, and WAHC, Christ Church, Barbados, February 26, 2016, Revised Selected Papers*, pages 176–192, 2016.
- RS92. Charles Rackoff and Daniel R. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In Joan Feigenbaum, editor, *Advances in Cryptology - CRYPTO’91*, volume 576 of *Lecture Notes in Computer Science*, pages 433–444. Springer, August 1992.
- RS06. P.Y.A. Ryan and S. A. Schneider. Prêt à voter with re-encryption mixes. Technical Report CS-TR-956, University of Newcastle, 2006.
- RT09. Peter Y. A. Ryan and Vanessa Teague. Pretty good democracy. In *IN: WORKSHOP ON SECURITY PROTOCOLS*, 2009.
- SK95. Kazue Sako and Joe Kilian. Receipt-free mix-type voting scheme - a practical solution to the implementation of a voting booth. In Louis C. Guillou

- and Jean-Jacques Quisquater, editors, *Advances in Cryptology – EURO-CRYPT’95*, volume 921 of *Lecture Notes in Computer Science*, pages 393–403. Springer, May 1995.
- SS10. Amit Sahai and Hakan Seyalioglu. Worry-free encryption: functional encryption with public keys. In Ehab Al-Shaer, Angelos D. Keromytis, and Vitaly Shmatikov, editors, *ACM CCS 10: 17th Conference on Computer and Communications Security*, pages 463–472. ACM Press, October 2010.
- Yao86. Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th Annual Symposium on Foundations of Computer Science*, pages 162–167. IEEE Computer Society Press, October 1986.

12 Future Directions

Our work opens up new directions in e-voting and generally in cryptography. We discuss some of them.

- **Efficiency.** In our work, in order to compute the NIWI proofs of Groth *et al.* [GOS06] for CircuitSat, we need to represent the computation as a Boolean circuit and, though this can be done in polynomial time, it can be inefficient in practice. An important objective is to sidestep the reduction to circuits by employing a more direct approach. A possibility would be to explore the achievability of our results from variants of Groth-Sahai NIWIs [GS08]. The NIWI of Groth-Sahai, as it stands, is formulated in the CRS model but it is worthy to study in which settings it can be instantiated without CRS. Another important direction is to improve the efficiency of the verification. It would be desirable that the verifiers make a work sub-linear in the number of voters. The verifiability guarantees attained would then be computational but hopefully it could be possible to avoid trust assumptions. A possibility would be to employ variants of succinct arguments (see [Bit14] for a survey).
- **Receipt-freeness.** Perfect verifiability and perfect correctness seem incompatible with receipt-freeness [BT94,SK95,MH96,MN06,DKR06,CFG15], but we think that it should be possible to define a statistical variant of verifiability that could coexist with some form of receipt-freeness. Another possibility could be to resort to some voting server trusted for receipt-freeness but not for privacy that re-randomizes the ballots, as done in BeleniosRF of Cortier, Fuchsbauer and Galindo [CFG15].
- **Other applications of our techniques.** We think that our techniques could be of wide applicability to other settings. For instance, Camenisch and Shoup [CS03a] put forth the concept of verifiable encryption (that in some sense could be also viewed as a special case of verifiable functional encryption [BGJS16]) and present numerous applications of it, such as key escrow, optimistic fair exchange, publicly verifiable secret and signature sharing, universally composable commitments, group signatures, and confirmer signatures. We believe that our techniques can be employed profitably to improve their results with the aim of removing the need of trust assumptions.