

Formal Analysis of Vote Privacy using Computationally Complete Symbolic Attacker

Gergei Bana¹, Rohit Chadha², and Ajay Kumar Eeralla²

¹ University of Luxembourg

² University of Missouri-Columbia, USA

Abstract. We analyze the FOO electronic voting protocol in the provable security model using the technique of Computationally Complete Symbolic Attacker (CCSA). The protocol uses commitments, blind signatures and anonymous channels to achieve vote privacy. Unlike the Dolev-Yao analyses of the protocol, we assume neither perfect cryptography nor existence of perfectly anonymous channels. Our analysis reveals new attacks on vote privacy, including an attack that arises due to the inadequacy of the blindness property of blind signatures and not due to a specific implementation of anonymous channels. With additional assumptions and modifications, we were able to show that the protocol satisfies vote privacy. Our techniques demonstrate effectiveness of the CCSA technique for both attack detection and verification.

1 Introduction

The FOO protocol was introduced by Fujioka, Okamoto, and Ohta in [1]. It was one of the first protocols for large-scale secure electronic voting. The design was supposed to achieve fairness, eligibility, vote privacy and individual verifiability. Since its publication, it has been the subject of several attempts to formalize and verify its security properties. The focus of this paper is the formal analysis of vote privacy of FOO, namely, the property that the votes of honest voters cannot be linked to the voters.

As far as we are aware of, the FOO protocol was the first electronic voting protocol that was formally specified and analyzed [2]. In the seminal work of Kremer and Ryan [2], both the formalization and analysis of the FOO protocol was carried out in the Dolev-Yao (DY) attacker model using the applied pi-framework [3].

The vote privacy property was the most intricate property analyzed in [2], and Kremer and Ryan were only able to prove vote privacy by hand. Subsequent development in Dolev-Yao verification allowed the proof to be carried out automatically [4, 5]. As the Dolev-Yao attacker model makes the assumption of perfect cryptography, the question of whether the proof carries over in the provable security model (computational model) remained unanswered. In addition to the assumption of perfect cryptography, the Dolev-Yao analysis also makes the assumption of existence of anonymous channels.

In this paper, we formally verify vote privacy in the provable security framework, using symbolic verification techniques. In particular, we use the computationally complete symbolic attacker technique for indistinguishability properties introduced by Bana and Comon in [6] (a.k.a. CCSA framework). As far as we know, this is the first formal

analysis of FOO protocol in the provable security model. However, we emphasize that proving vote privacy for FOO is only one aim of this paper. The other is to further develop the library of axioms of the CCSA framework and to demonstrate its effectiveness in attack detection and verification.

CCSA technique was first introduced by Bana and Comon in [7] for reachability properties and then for indistinguishability properties in [6]. Since then it has been used to find new attacks to the Needham-Schroeder-Lowe protocol [8, 9]; to treat algebraic operations notoriously difficult to reason about in the Dolev-Yao framework, such as exponentiation along with it the decisional Diffie-Hellmann property and versions of the Diffie-Hellman key-exchange protocol [10]; to verify unlinkability of RFID protocols [11]; and to analyze key wrapping API's [12]. Automated tool is not yet available for the indistinguishability technique, but work is in progress. In the meantime, we continue developing the library of axioms, and verifying relatively simple protocols by hand. In this paper, we add the axioms for security assumptions for trapdoor commitments such as hiding property and for blind signatures such as the blindness property, which are necessary for the FOO verification.

FOO protocol assumes two election authorities: administrator and collector. The protocol proceeds in three phases: In the first phase, voters prepare their ballots in the form of a trapdoor commitment of their votes and obtain a blind signature on the ballots from the administrator indicating that they are eligible to vote. In the second, voters send their ballots to the collector using an anonymous channel who publishes them on a public bulletin board (BB). In the final phase, each voter verifies the presence of their ballots on the BB and then sends their trapdoor key along with the entry number of their commitment on bulletin board to the collector, again via an anonymous channel. The trapdoor keys are then also added to the bulletin board next to the commitments. After the votes finish, votes are tallied from BB. The creators of the FOO protocol did not specify how the anonymous channels are implemented. In our modeling, we model anonymous channel as a mix-net server [13] which, upon receiving a list of encrypted messages outputs their plaintexts in lexicographic order after decrypting them.

FOO protocol is designed to provide vote privacy even when the administrator and collector are corrupt and our analysis of vote privacy assumes this to be the case. As in [2, 14, 15], vote privacy is modeled as indistinguishability of two protocol executions: in one, honest voter A votes for candidate v_1 and honest voter B votes for candidate v_2 and in the other, honest voter A votes for candidate v_2 and honest voter B votes for candidate v_1 . Observe that, as stated above, FOO protocol does not satisfy vote privacy as the attacker may choose to forward only Alice's trapdoor key in the final phase of the protocol to the mix-net server. We argue that privacy of votes can never be guaranteed for any voting protocol if the attacker allows only one participant to complete the protocol. Hence, our formalization is carefully crafted to avoid these cases.

Our analysis revealed new attacks on the FOO protocol (See Section 3.2). The first attack occurs because of an inadequacy of blindness property of blind signatures. Intuitively, blindness [16] means that a dishonest signer who engages in two sessions (parallel or sequential) with an honest user on messages m_0 and m_1 cannot detect which session is for m_0 and which session is for m_1 if the user successfully outputs signatures in *both* sessions. The blindness property, however, does allow the possibility that

the attacker can distinguish between the sessions if the user is successful in only one session. In order to prevent this attack, we have to assume that the identities of the candidates are of equal length. This attack does not depend on the implementation of the anonymous channel. The second attack exploits the fact that encryption scheme used by the mix-net server may be length-revealing and hence the length of the encrypted messages to the mix-net server may reveal their senders. In order to prevent this, we have to assume that the signatures obtained on *equal length* messages by executing the blind signature interactive protocol with the *same* signer must be of equal length. The above two attacks lie outside the Dolev-Yao model and hence were not detected in previous works on formal analysis of FOO protocol. A third attack is a Dolev-Yao style replay attack in which messages from the Voting phase can be replayed in the Opening phase. This attack can be prevented by introducing phase numbers in the FOO protocol.

With these additional assumptions, we were able to establish vote privacy of the FOO protocol. The proof of vote privacy rests on the blindness property of the blind signature, the computational hiding property of the trapdoor commitments and of IND-CCA2 assumption on the encryption used in the anonymizing mechanism. The proof of vote privacy in the Dolev-Yao analysis, in contrast, relies only on the blindness property.³ The commitment hiding property of trapdoor commitments does not play a role in establishing vote privacy in the Dolev-Yao analysis. As Dolev-Yao analysis assumes anonymous communication as a primitive, IND-CCA2 justifiably does not play a role in the Dolev-Yao analysis.

Related work. There have been several attempts at formal analysis of FOO protocol in the Dolev-Yao model (see, for example, [2, 4, 5]). These analyses assume perfectly anonymous channels. In the computational model, there are several attempts at formalizing vote privacy in electronic voting such as in [14, 15, 17, 18]. Please see [18] for a comparison amongst these definitions. All these definitions apply to single phase voting protocols. Our definition is adapted to FOO-protocol which has three phases. The only other work at formally verifying vote privacy for electronic voting that we are aware of is the mechanized proof of vote privacy for the Helios family of single-phase protocols given in [19]. The inadequacy of the blindness axiom has also been pointed out in [20], who show how any blind signature scheme can be combined with trapdoor commitments to resolve this inadequacy.

2 FOO voting protocol and its computational modeling

We briefly recall the electronic voting protocol FOO introduced by Fujioka, Okamoto, and Ohta in [1]. We assume that the reader is familiar with the cryptographic primitives of public key encryption, trapdoor commitment schemes and digital signatures schemes. The FOO protocol also uses blind signature schemes. Informally, a blind signature scheme is an interactive protocol that allows a party \mathcal{U} to obtain the signature of a party \mathcal{S} on a message that is obfuscated by \mathcal{U} until \mathcal{S} completes the protocol and \mathcal{U} publishes the signed message. We assume as in [21, 22] that the interactive protocol

³ As evidence, we ran AKiSs [5] on a two-phase variant of the FOO protocol without commitments. The variant satisfies the vote privacy property in the Dolev-Yao model.

Authentication :

- 1 : V_i : computes $c_i := \xi(v_i, k_i)$, a commitment on her vote v_i with trapdoor k_i , $b_i := \chi(c_i, r_i)$, blinding the ballot c_i with blinding key r_i , and $s_i := \text{sign}_i(b_i)$, a signature on b_i
- 2 : $V_i \rightarrow A : \langle ID_i, b_i, s_i \rangle$
- 3 : A : verifies the signature s_i and eligibility (not yet applied for his signature) of V_i
- 4 : $A \rightarrow V_i : \text{bsign}_A(b_i)$, (blind) signature of A on the blinded ballot b_i

Voting :

- 5 : V_i : checks if the received message is A 's blind signature on b_i ; if the check passes, V_i unblinds the signature and it is denoted by $\sigma_A(c_i)$
- 6 : $V_i \dashrightarrow C : \langle c_i, \sigma_A(c_i) \rangle$
- 7 : C : verifies the signature, and adds the received message to Bulletin Board (BB) with label ℓ_i
- 8 : C : after a certain time, C publishes the BB

Opening :

- 9 : V_i : finds her ballot on the BB and records its label ℓ_i
- 10 : $V_i \dashrightarrow C : \langle \ell_i, k_i \rangle$
- 11 : C : opens the commitments at ℓ_i using the key k_i , and appends $\langle k_i, v_i \rangle$ to $\langle c_i, \sigma_A(c_i) \rangle$ in the BB, and after a certain time, publishes the new BB
- 12 : C : checks the validity of the votes, counts them and publishes the result

Fig. 1. FOO Voting Protocol

consists of three phases: blinding by \mathcal{U} , (blind) signing by \mathcal{S} and unblinding by \mathcal{U} . The primitives are described in detail in Section 2.2.

The FOO protocol has three roles: voters (V_i), administrator (A), and collector (C). It assumes the existence of anonymous channels. Following Kremer et al [2], we group the protocol in three phases, Authentication, Voting, and Opening, as described in Figure 1. In the Figure, \rightarrow means send, \dashrightarrow means send via anonymous channel.

2.1 Anonymous Channel

The creators of FOO did not specify how to implement anonymous channels. We now clarify our modeling of the anonymous channels. We model the anonymous communication using a mix-net server which, upon receiving a list of messages encrypted with its public key, checks if they are all different, decrypts each message in the list, outputs them in lexicographic order. Henceforth, we refer to mix-net server as the mixer. In particular, in the Voting phase (resp. Opening phase), the voter encrypts $\langle c_i, \sigma_A(c_i) \rangle$ (resp. $\langle \ell_i, k_i \rangle$) with the mixer M 's public key pk_M . The mixer waits until a certain time during which receives the messages from the voters, checks if all messages received until that time are distinct. If the check succeeds, the mixer decrypts the messages, shuffles the decrypted messages into lexicographic order and outputs them.

2.2 Computational modeling of the FOO protocol

As usual, we assume that all agents in the protocol execution, the voters, the administrator, the collector are interactive probabilistic polynomial-time Turing machines.

Furthermore, the network is controlled by an attacker, which is also an interactive probabilistic polynomial-time Turing machine. Each message goes through the attacker, except those that are “published”, in which case the message is written directly on each participant’s work tape synchronously. The Bulletin Board BB is simply a list of tuples. We further assume that the id’s ID_i are the same as the voters’ public keys used for verifying their digital signatures.

We assume that encryption used in the protocol is *indistinguishable against adaptive chosen cipher text attack*, i.e., satisfies the IND-CCA2 property. The interested reader is referred to [23]. Digital signatures are used in the protocol to ensure that only eligible voters vote and do not play a part in establishing vote privacy. We now present the computational modeling of the commitment and blind digital signature schemes, and their security properties we assume for the verification.

Trapdoor Commitments. A trapdoor commitment scheme allows one party, say V , to commit to a value obfuscating the value it committed to until it is revealed by V with the help of the trapdoor. In general, a commitment scheme satisfies two properties: *binding* and *hiding*. The former property states that the party V cannot change the message once V committed to it while the latter property says that it should be computationally infeasible for any other party to retrieve the message from the commitment unless the V reveals the trapdoor.

Formally, a commitment scheme [23] \mathcal{C} is a triple of algorithms $(KG, Commit, Open)$ such that KG is a PPT algorithm with input 1^η , while $Commit$ and $Open$ are PT algorithms such that there is a polynomial $p(\cdot)$ such that for each security parameter η and bit-string m of length at most $p(\eta)$, $Commit(m, k)$ computes the commitment on message m for $k \leftarrow KG(1^\eta)$, and $Open$ is a deterministic algorithm such that for all $k \leftarrow KG(1^\eta)$, $Open(Commit(m, k), m, k) = true$. We assume that the commitment schemes are length regular. That is if m_0 and m_1 are messages of equal length then the commitments of m_0 and m_1 are of equal length [24, 25].

The computational hiding property is formalized as a game between a (PPT) attacker and a (honest) challenger. Initially, the attacker generates two messages m_0 and m_1 of *equal length*, which it gives to the challenger. The challenger creates commitments, com_0 and com_1 for m_0 and m_1 with different (secret) commitment keys and sets a *secret* bit b with uniform probability. The challenger then gives the ordered pair (com_b, com_{1-b}) to the attacker who wins if it correctly guesses the value of b . The commitment scheme is said to satisfy the commitment hiding property [23] if the winning probability of the attacker cannot be non-negligibly different from $\frac{1}{2}$. Please see Appendix A for the formal definition.

Blind Digital Signatures. Blind signature schemes were introduced by Chaum [21]. They are small interactive protocols between a user and a signer. Informally, a blind signature scheme is an interactive protocol that allows a party \mathcal{U} to obtain the signature of a party \mathcal{S} on a message that is obfuscated by \mathcal{U} until \mathcal{S} completes the protocol and \mathcal{U} publishes the signed message.

Formally, a blind digital signature scheme [16, 22] B is a tuple $(Gen, \langle \mathcal{S}, \mathcal{U} \rangle, verify)$, a PPT key generation algorithm Gen , an interactive algorithm $\langle \mathcal{S}, \mathcal{U} \rangle$ of the PPT signer \mathcal{S} and the PPT user \mathcal{U} , and a PT verification algorithm $verify$, such that for each security parameter η : $Gen(1^\eta)$ outputs a public and private key pair (pk, sk) , the joint execution

of signer $\mathcal{S}(pk, sk)$ and the user $\mathcal{U}(pk, m)$ on the message $m \in \{0, 1\}^n$ produces an output σ for the user, and the deterministic algorithm $\text{verif}(m, \sigma, pk)$ outputs a bit. The algorithm $\mathcal{U}(pk, m)$ may fail to produce an output; in which case \mathcal{U} is said to abort and $\mathcal{U}(pk, m)$ said to be undefined. In that case, we write $\mathcal{U}(pk, m) = \perp$. A blind signature scheme is required to be *complete*, i.e., for all parameter η the following holds: for all messages $m \in \{0, 1\}^n$, if $(pk, sk) \leftarrow \text{Gen}(1^\eta)$ then the joint execution of $\mathcal{S}(pk, sk)$ and $\mathcal{U}(pk, m)$ must be defined and the produced output σ should be such that $\text{verif}(m, \sigma, pk) = \text{true}$.

Blind signatures should also satisfy two security properties, *blindness* and *unforgeability*. Intuitively, unforgeability says a malicious user who engages in ℓ sessions with an honest user should not be able to produce $\ell + 1$ valid message-signature pairs. In the FOO protocol, unforgeability is used to make sure only eligible voters vote. It does not play a role in voting privacy. For this reason, we omit the formalization of unforgeability. The interested reader is referred to [16].

The blindness property [16] on the other hand, is useful in establishing privacy of votes. Intuitively, blindness means that a dishonest signer cannot learn the contents of the message it signed. The blindness property is formalized as a game between a (PPT) attacker acting as the signer and a (honest) challenger. Initially, the attacker generates two messages m_0 and m_1 , which it gives to the challenger. The challenger sets a *secret* bit b with uniform probability. The challenger, acting as the user, then engages in two blind signature sessions with the attacker as the signer; the first is for obtaining signature on m_b and the second is for obtaining signature on m_{1-b} . If the user successfully completes both sessions and obtains signature σ_b, σ_{1-b} then it gives the ordered pair (σ_0, σ_1) to the attacker. Otherwise, it gives the pair (\perp, \perp) to the attacker. The attacker wins if it correctly guesses the value of b . Note that it is important that the challenger gives the pair (σ_0, σ_1) and not (σ_b, σ_{1-b}) (otherwise the attacker will win with probability 1). The blind signature scheme is said to satisfy the blindness property [16] if the winning probability of the attacker cannot be non-negligibly different from $\frac{1}{2}$. (Please see Appendix A for the formal definition).

We model the blind signatures using an interactive protocol between the user and the signer in three steps, blinding, blind signing, and unblinding and involving four algorithms, *blind*, *bsign*, *accept*, and *unblind* [21,22]. Before the protocol commences, the signer generates a public key and secret-key pair (pk, sk) and publishes pk . **Blinding:** User computes $b := \text{blind}(m, pk, r_b)$ for a message m , using the signer's public key pk , and a random seed r_b , and sends b to the signer. **Blind signing:** The signer computes $\rho := \text{bsign}(b, sk, r_s)$ using his own secret key sk and a random seed r_s , and sends it to the user. **Unblinding:** User checks validity of ρ by running $\text{accept}(m, pk, r_b, \rho)$. If it outputs false, the user quits the protocol without any output. If it outputs true, the user computes $\sigma = \text{unblind}(m, pk, r_b, \rho)$.

3 Vote privacy for the FOO protocol

We describe the targeted security notion of the voting, namely, that votes of honest voters remain secret. The fundamental idea of vote privacy in the literature is that permutation of votes cast by honest voters cannot be detected by the attacker. There are

numerous attempts to formalize this idea into rigorous terms, see for example [2, 18]. A major distinction between the protocols considered in [2, 18] and FOO protocol is that unlike the *single-phase* protocols considered therein, the Bulletin Board (BB) is built incrementally in FOO in different phases. Hence, we have to adapt the definition of privacy of votes to FOO protocol. We describe our formalization below. We consider only one session.

3.1 Formalization of Vote Privacy

We assume that the attacker controls the network and we allow the administrator and the collector to be corrupt. We require that the flipping of votes of two honest voters should result in computationally indistinguishable runs to any PPT attacker even if he controls all other voters. Under these assumptions however, there is an obvious unavoidable attack. Suppose A votes v_0 and B votes v_1 , while in the flipped voting scenario, A votes v_1 and B votes v_0 . If the attacker allows the ballots of both A and B to be posted in the Voting phase, but subsequently blocks B 's message in the opening phase, then the two situations will be distinguishable. This is because in the first scenario only v_0 vote will appear on the BB, while in the second, only v_1 . Therefore, our formalization requires that the scenario with the original votes and with the flipped votes be indistinguishable only when either both voters' votes appear on the BB or neither does.

Accordingly, we require the following two games to be computationally indistinguishable to a PPT attacker. For bit b , let Π_b be the following game between the Attacker and the Challenger simulating two honest voters of the FOO protocol and the mixer but aborting when only one of the commitment keys reach the mixer. The attacker simulates the corrupt administrator, the corrupt collector, and at the end has to guess the value of b :

1. Challenger publishes public keys of A , B , and the mixer, and a list of candidates.
2. Attacker publishes possibly dishonestly generated public keys of other voters, the administrator, and the collector. It then gives Challenger two valid votes v_0 and v_1 .
3. Challenger creates the ballots of v_b for A and v_{1-b} for B .
4. Attacker calls one of A or B to authorize vote with administrator.
5. Challenger prepares the corresponding blinded ballot. It sends the corresponding identity, the blinded ballot, and the corresponding signature to the attacker.
6. Attacker creates the possibly fake authorization, sends it back to the voter, and calls for the other voter to carry out the authorization.
7. Challenger prepares the appropriate blinded ballot and sends it to the attacker.
8. Attacker creates and sends again the authorization now for the other voter, and asks either A or B to proceed with sending the ballot.
9. If the corresponding blind signature was accepted, Challenger unblinds it, and sends the signed ballot encrypted with the mixer's public key. Otherwise skips.
10. Attacker makes some computations and sends a message back (which is possibly the unchanged ciphertext sent in the previous step).
11. If the other of A or B 's blind signature was accepted, Challenger now unblinds that, sends the signed ballot encrypted with the mixer's public key. Otherwise skips.
12. Attacker makes some computations and sends a message back (which is possibly the unchanged ciphertext sent in the previous step).

13. Attacker also sends the Challenger further possible encrypted signed ballots of the other, possibly corrupted voters.
14. Challenger waits for all the signed encrypted ballots to arrive. It decrypts the signed ballots with the secret key of the mixer, and puts them through the shuffle. Gives the result of the shuffle to the attacker.
15. Attacker creates the bulletin board BB and gives it to the Challenger. Note that the ballots of A and B may or may not appear on the BB. Attacker also specifies which of A or B should move to the next step first.
16. If A (resp. B) was asked to move next, A (resp. B) accepted the blind signature in phase 2 and A 's (resp. B 's) ballot appears on the bulletin board, then the Challenger sends A 's (resp. B 's) $\langle l_i, k_i \rangle$ to the attacker encrypted with the public key of the mixer. Otherwise, the challenger skips this step.
17. Attacker makes some computations and sends a message back (which is possibly the unchanged cipher sent in the previous step).
18. If the other agent accepted the blind signature in phase 2 and its ballot appeared on the bulletin board, then the Challenger sends the other agent's $\langle l_i, k_i \rangle$ to the attacker encrypted with the public key of the mixer. Otherwise, it skips this step.
19. Attacker makes some computations and sends a message back (which is possibly the unchanged cipher sent in the previous step).
20. Attacker may also send other messages of the form $\langle l_i, k_i \rangle$ encrypted.
21. Challenger puts the decrypted messages through the shuffle. If only one of A 's or B 's commitment key appears, the Challenger aborts. If none of them appear or both of them appear, the Challenger gives the result of the shuffle to the attacker.
22. The attacker outputs 0 or 1.

If in the situation above the probability that the attacker outputs 1 playing game Π_0 is non-negligibly different from outputting 1 playing game Π_1 , then the attacker wins the game. Our aim is to show that no PPT attacker can win the above game.

3.2 Attacks on the FOO Protocol

We describe three attacks on the vote privacy for FOO protocol, which we caught with the CCSA technique. We will comment in the next section on how attacks are found. The first two attacks cannot be captured in the Dolev-Yao attacker model as they cannot be discovered under the assumption of perfect cryptography. The third attack can be captured by a Dolev-Yao attacker in nature but does not appear in previous Dolev-Yao analyses of the protocol as those analyses assume perfectly anonymous channels.

The first attack exploits an insufficiency in the blindness property. This insufficiency has also been pointed out in [20]. Recall that in the blindness game, a (potentially dishonest) signer engages in two sessions (parallel or sequential) with an honest user on messages m_0 and m_1 . If one session successfully completes and the other aborts then the information of which session aborted is not revealed to the attacker. Thus, the blindness property does *not rule out* the possibility that the signer is able to deduce which session corresponded to m_0 and which corresponded to m_1 if one session aborted and the signer knows which session aborted. For example, in the three-step blind signature scheme described above, replace *accept* by *accept'* where *accept'*(m, pk, r_b, ρ) always

returns false if m is of even length and returns the value of $\text{accept}(m, pk, r_b, \rho)$ otherwise. The resulting blind signature scheme continues to satisfy the blindness property if the original one does. Now, if m_0 is of even length and m_1 is of odd length, then the session corresponding to message m_0 always aborts and thus the dishonest signer be able to associate the aborted session to m_0 .

The first attack on FOO protocol proceeds as follows. Assume that we use the modified blind signature scheme in the FOO protocol and that the candidate choice v_0 is of even length and the choice v_1 is of odd length. Since commitments may reveal the length, it is possible that the commitment of v_0 is also even while of v_1 is odd. Assume that A chooses v_b and Bob chooses v_{1-b} . Vote privacy requires that the attacker not be able to deduce the value of b . In authentication phase, the attacker, acting as the administrator, correctly follows the protocol. The blinding still hides the commitments even though their lengths are different, but according to our modified accept algorithm, the signature on v_0 will be rejected while the one on v_1 will be accepted. Then during the voting phase, only one voter will send its ballot to the administrator. If this voter is A then b must be 1, otherwise b must be 0. Note, this attack is not a feature of our specific implementation of the anonymous channel. Note also that even though the commitments are revealed on the BB, there they appear in the same order on the two sides, so the attack works not because the commitments reveal to the attacker the vote inside, but because one accept may pass, the other may not. Observe also that the attack does not manifest if v_0 and v_1 are of equal length. Indeed, if v_0 and v_1 are of equal length then the property of commitment hiding ensures that an abort by a voter does not reveal information about its intended vote.

The second attack relies on the fact that there is no guarantee on the lengths of the (blind) signatures and that an encryption scheme used may reveal the lengths of the underlying plaintexts. If a length revealing encryption scheme (such as IND-CCA2) is used then the encrypted ballots in the voting phase are different, which allows the attacker to deduce how A and B voted. To *rule out* this attack, we have to assume that the (blind) signatures on equal length messages are of equal length.

Finally, we point out the Dolev-Yao style attack on the FOO protocol. The attack replays a message in the voting phase of the protocol in the opening phase. Assume that there is a third voter C in addition to A and B . The attacker remembers A 's encrypted ballot to the mixer in the voting phase. In the opening phase of the protocol, the attacker forwards the encrypted commitment keys of A and B to the mixer, but replaces C 's commitment key by A 's encrypted ballot from the voting phase. The mixer checks that the three encryptions are different, decrypts them and outputs them after shuffling them into lexicographic order. The only ballot that appears in the output is A 's. The commitment keys of A and B are also part of the output. The attacker can then deduce how A and B voted. In order to prevent this attack, we modify the FOO protocol to *also include phase numbers* inside the encryptions in the voting and opening phases of the FOO protocol. The mixer *must* check that the phase numbers correspond to the current phase in the protocol.

4 Computationally complete symbolic attacker (CCSA)

The CCSA model was introduced by Bana and Comon in [7] with the aim of establishing computational guarantees for symbolic analysis by making the symbolic attacker as powerful as the computational one. We limit our attention in this paper to their indistinguishability framework given in [6]. The CCSA framework is similar to the Dolev-Yao framework in the sense that explicit symbolic adversarial messages are created. The Dolev-Yao framework specifies *all* rules that the attacker can use to create new messages from what he has seen, and the protocol agents use pattern matching to check whether messages coming from the attacker have the correct form. On the other hand, in the CCSA framework, each message from the attacker is modeled by a function f_i applied to the sequence of messages that the attacker has seen thus far. Pattern matching is replaced by applying function symbols on the term coming from the attacker. Limitations on attacker capabilities originating from computational assumptions on the primitives are specified as a set of axioms A in first-order logic based on a single indistinguishability predicate (representing computational indistinguishability of sequences of messages). Computationally, a security property of a protocol Π is formulated as the computational indistinguishability of two protocols Π_1 and Π_2 constructed (depending on the security property) from the original Π .

Verification. In the CCSA framework, the security translates to the validity of the formula obtained by applying the indistinguishability predicate on the list of terms produced by the symbolic execution of Π_1 and Π_2 . Hence the security formula obtained this way must be derived as a logical consequence of the axioms A using first-order inference. If the formula is derived and the axioms in A shown to be computationally sound then the protocol Π is computationally secure according to Theorem 1 in [6].

Attack finding. If the security formula cannot be derived then the negation of the security formula is consistent with A and a symbolic attacker model is then obtained. In practice, for the security proof, a proof tree is being built, and some branch cannot be reduced to axioms. For example, often what happens is that it cannot be shown that an attacker message $f_i(\dots)$ cannot equal a term, that is, the attacker may produce a message that passes some check when it should not. Then from this branch, an Herbrand model can be built on the terms, resulting in a symbolic attack. This symbolic attack then has to be checked to see if it corresponds to a computational attacker or if it originates from having too few axioms in A , as the set of axioms is not complete.

4.1 Syntax

Terms: Let S be a finite set of *sorts* that includes at least the sorts `bool` and `msg`. \mathcal{X} is an infinite set of *variable symbols*, each coming with a sort $s \in S$. Terms are built on a set of function symbols \mathcal{F} representing honest computation of primitives, a set of function symbols \mathcal{G} representing attacker's computation, a set of zero-arity function symbols (names) \mathcal{N} representing honest generation of randomness, and a set of variables \mathcal{X} . The set \mathcal{F} contains the basic symbols such as `0`, `true`, `false`, `L`, `<-`, `->`, `EQ(-, -)`, `if _ then _ else _`, π_1 and π_2 with the typing rules as follows:

- `0` : `msg` representing the empty message.

- $\mathit{true}, \mathit{false} : \mathit{bool}$.
- $L(_) : \mathit{msg} \rightarrow \mathit{msg}$. $L(x)$ represents the length of x in unary.
- $\langle _, _ \rangle : \mathit{msg} \times \mathit{msg} \rightarrow \mathit{msg}$ representing pairs.
- Polymorphic equality test $\mathit{EQ}(_, _) : \mathit{msg} \times \mathit{msg} \rightarrow \mathit{bool}$
 $\mathit{bool} \times \mathit{bool} \rightarrow \mathit{bool}$.
- Polymorphic conditional branching $\mathit{if_then_else_} : \mathit{bool} \times \mathit{msg} \times \mathit{msg} \rightarrow \mathit{msg}$
 $\mathit{bool} \times \mathit{bool} \times \mathit{bool} \rightarrow \mathit{bool}$.
- $\pi_1(_), \pi_2(_) : \mathit{msg} \rightarrow \mathit{msg}$. π_i represents the i -th projection of a pair

Formulas: As presented in [10], we have for every sequence of sorts s_1, \dots, s_n a *predicate symbol* that takes $2 \times n$ arguments of sort $(s_1 \times \dots \times s_n)^2$, which we write as $t_1, \dots, t_n \sim u_1, \dots, u_n$. The predicate $t_1, \dots, t_n \sim u_1, \dots, u_n$ represents *computational indistinguishability* of the two sequences of terms t_1, \dots, t_n and u_1, \dots, u_n .

The first-order formulas are built from the above atomic formulas combining the Boolean connectives \neg, \wedge, \vee , and \rightarrow , and the quantifiers \forall and \exists . The formulas are used to represent both axioms (assumptions) and the security properties of the protocols.

Equational Theory: We use binary relation symbol $=_E$ to indicate equations the function symbols have to satisfy. Note that $=_E$ is not part of the first-order signature. The equations specified by $=_E$ will result in axioms for \sim . We assume that the following hold: $\forall x_1, x_2. \pi_i(\langle x_1, x_2 \rangle) =_E x_i$, for $i = 1, 2$; $L(L(x)) =_E L(x)$.

Abbreviations: In order to reduce the size of the terms, we shall use the following abbreviations: $\mathit{not}(b) \stackrel{\text{def}}{=} \text{if } b \text{ then } \mathit{false} \text{ else } \mathit{true}$, $b_1 \ \& \ b_2 \stackrel{\text{def}}{=} \text{if } b_1 \text{ then } b_2 \text{ else } \mathit{false}$, $b_1 \ \text{or} \ b_2 \stackrel{\text{def}}{=} \text{if } b_1 \text{ then } \mathit{true} \text{ else } b_2$ and $x = y \stackrel{\text{def}}{=} \mathit{EQ}(x, y) \sim \mathit{true}$. The abbreviation $\mathit{EQL}(x_1, x_2) \stackrel{\text{def}}{=} \mathit{EQ}(L(x_1), (L(x_2)))$ modeling length equality test.

4.2 Semantics

As Bana and Comon defined it in [6], the logic is interpreted over a *computational model*. A computational model \mathcal{M}^c is a particular first-order model in which the domain consists of probabilistic polynomial-time algorithms taking the input 1^η together with two infinitely long random tapes (ρ_1, ρ_2) , where ρ_1 is for honestly generated names, and ρ_2 for adversarial use. We use $\llbracket _ \rrbracket$ to denote the semantics of syntactic objects. Once $\llbracket _ \rrbracket$ is given on function symbols, and a valuation σ of variables of term t in the domain, the interpretation of t in the domain is defined the usual way and denoted as $\llbracket t \rrbracket^\sigma$. Let $\llbracket t \rrbracket_{\eta, \rho}^\sigma := \llbracket t \rrbracket^\sigma(1^\eta; \rho_1, \rho_2)$. In this model, function symbols in \mathcal{F} representing the cryptographic primitives are interpreted as polynomial-time algorithms and they act on the outputs of the PPT algorithms in their arguments. For example, $\llbracket \mathit{EQ}(t_1, t_2) \rrbracket_{\eta, \rho}^\sigma := 1$ if $\llbracket t_1 \rrbracket_{\eta, \rho}^\sigma = \llbracket t_2 \rrbracket_{\eta, \rho}^\sigma$, otherwise it is 0, and $\llbracket \text{if } t_1 \text{ then } t_2 \text{ else } t_3 \rrbracket_{\eta, \rho}^\sigma := \llbracket t_2 \rrbracket_{\eta, \rho}^\sigma$ if $\llbracket t_1 \rrbracket_{\eta, \rho}^\sigma = 1$, and $\llbracket \text{if } t_1 \text{ then } t_2 \text{ else } t_3 \rrbracket_{\eta, \rho}^\sigma := \llbracket t_3 \rrbracket_{\eta, \rho}^\sigma$ if $\llbracket t_1 \rrbracket_{\eta, \rho}^\sigma = 0$. Function symbols in \mathcal{G} representing adversarial computation are interpreted as probabilistic polynomial-time algorithms and they also act on the outputs of the PPT algorithms in their arguments, but they can also use randomness from ρ_2 . Each name $n \in \mathcal{N}$ is interpreted as a machine $\llbracket n \rrbracket$ which extracts a random word of length $p(\eta)$ (where p is a polynomial globally fixed by \mathcal{M}^c) from ρ_1 . Different names draw from disjoint parts of the tape ρ_1 . The predicate symbol \sim is interpreted as computational indistinguishability of the outputs of the PPT algorithms on the two sides of \sim . Interested reader can consult either [6]

Commitments	Encryptions	Blind Signatures	Signatures	Other Symbols
$k_c(-): \text{msg}$		$\perp: \text{msg}$		$ph_2, ph_3: \text{msg}$
$com(-, -): \text{msg}$	$k_e(-): \text{msg}$	$k_b(-): \text{msg}$		$A, B, M: \text{msg}$
$open(-, -, -): \text{msg}$	$r_e(-): \text{msg}$	$r_b(-): \text{msg}$	$k_s(-): \text{msg}$	$C_1, C_2, C_3: \text{msg}$
	$\{-\}_-: \text{msg}$	$bsign(-, -, -): \text{msg}$	$r_s(-): \text{msg}$	$to(-): \text{msg}$
Shuffling	$dec(-, -): \text{msg}$	$b(-, -, -): \text{msg}$	$sign(-, -, -): \text{msg}$	$V_0(-): \text{msg}$
$shuff(-, -, -): \text{msg}$		$ub(-, -, -, -): \text{msg}$	$ver(-, -, -): \text{bool}$	$V_1(-): \text{msg}$
		$acc(-, -, -, -): \text{bool}$		$pubkey(-): \text{msg}$
		$bver(-, -, -): \text{bool}$		

Table 1: FOO function symbols with their co-domain sorts (each argument sort is msg)

or [10] for complete definition of this semantics. A formula is computationally valid if it is satisfied in all computational interpretations.

4.3 Computationally Valid Axioms

A core set of valid axioms of the CCSA have been given in [10] and are reproduced in Appendix B. Essentially, these core axioms can be divided into four sets. The first set of axioms reflect the properties of the indistinguishability predicate \sim , namely, that it is an equivalence relation and is preserved under projection, permutation and function application. The second set of axioms says that the abbreviation $=$ is a congruence and preserves the equational theory $=_E$. The third set of axioms reflect the properties of the function symbol `if _ then _ else _` such as `if true then x else y = x`. Finally, we have a couple of axioms that reflect the *truly* random nature of names such as replacing any name n_1 by a fresh name n_2 yields indistinguishable sequence of terms. The interested reader is referred to [10] for additional details, where several small examples of usage of the axioms are also given.

5 Vote Privacy for the FOO Protocol in CCSA Framework

We now explain how vote privacy for modified FOO protocol can be modeled in CCSA framework. We assume that the FOO protocol is modified to ensure that the length candidate identities are equal and that the phase identifiers are used in Voting and Opening phase. We will drop the adjective modified in the rest of the section. In order to make things simpler, we assume that other than voters A and B , there is only a single other (possibly dishonest) voter. The proof of vote privacy carries over to n voters for any n .

5.1 Function Symbols and Axioms

In order to formalize the FOO voting protocol in the CCSA framework, we shall include the following (see Table 1) function symbols in \mathcal{F} . Please note that each argument in these symbols has sort `msg`. Each key-generation algorithm and each random seed generation will be represented by a corresponding unary function symbol. When the function symbol is applied to a name n (recall that names represent truly random objects), then the resulting term shall represent a correctly generated key or seed.

Commitments. We include a symbol $k_c(\cdot)$ in \mathcal{F} to represent the key generation algorithm for the commitment schemes defined in Section 2.2. $\text{com}(x, y)$ is the commitment of x using the key y whereas $\text{open}(u, z)$ is the opening of the commitment u using the key z . We assume that, for all x, y , $\text{open}(\text{com}(x, k_c(y)), k_c(y)) =_E x$.

Encryptions. The symbol $k_e(\cdot)$ denotes the public-key secret-key pair generation algorithm of encryptions. We define $\text{pk}_e(x) \stackrel{\text{def}}{=} \pi_1(k_e(x))$ and $\text{sk}_e(x) \stackrel{\text{def}}{=} \pi_2(k_e(x))$, the public encryption key and secret decryption key parts of $k_e(x)$ respectively. In order to formalize the random seed of encryptions, we introduce a symbol $r_e(\cdot)$ (see Table 1). Encryption of x with random seed z and public key y is denoted as $\{x\}_y^z$ and decryption of x with secret key z is denoted as $\text{dec}(x, z)$. We assume that for all x, y, z , $\text{dec}(\{x\}_{\text{pk}_e(y)}^z, \text{sk}_e(y)) =_E x$.

Blind signatures. We introduce a symbol $k_b(\cdot)$ to represent a public-verification-key secret-signing-key pair generation algorithm of blind signatures. Once again, $\text{pk}_b(x) \stackrel{\text{def}}{=} \pi_1(k_b(x))$ and $\text{sk}_b(x) \stackrel{\text{def}}{=} \pi_2(k_b(x))$ represent the public verification key and secret signing key parts of the key $k_b(x)$. We also introduce a symbol $r_b(\cdot)$ to represent blinding key generation algorithm. The interpretations of the function symbols b , bsign , bver , ub and acc are *blind*, *blindsign*, *verify*, *unblind* and *accept* as defined in Section 2.2. $b(x, y, z)$ is the blinding of the message x using the public verification key y and the blinding key z . $\text{bsign}(x, y, r)$ is the (blinded) message x blind-signed with secret blind signing key y and random seed r . $\text{acc}(x, y, z, w)$ is the acceptance check of the blind-sign w created for the blinded message $b(x, y, z)$ and $\text{ub}(x, y, z, w)$ its unblinding. \perp represents undefined that is used to model blindness game. Finally, $\text{bver}(x, s, y)$ is the verification of unblinded signature s on the message x with the public verification key y . The co-domains of the symbol acc and bver is **bool**.

Digital signatures. In addition to blind signatures, we also need plain digital signatures. As in [10], let $k_s(\cdot)$ represent generation of a key pair, the public verification key and the secret signing key. $\text{vk}(x) \stackrel{\text{def}}{=} \pi_1(k_s(x))$ and $\text{ssk}(x) \stackrel{\text{def}}{=} \pi_2(k_s(x))$ represent the public verification key and secret signing key parts of the key $k_s(x)$. In order to randomize signatures, we introduce a symbol $r_s(\cdot)$ for random seed generation. $\text{sign}(x, y, r)$ is the message x signed with secret signing key y with a random seed r and $\text{ver}(z, u, y)$ is the verification of signature u on the message z with the public verification key y .

Mixer. We introduce a symbol $\text{shuf}(\cdot, \cdot, \cdot)$ to model the shuffling functionality of mixnets. $\text{shuf}(x_1, x_2, x_3)$ is the lexicographic ordering of the messages x_1, x_2 , and x_3 . As we model only three voters, shuf has only three arguments.

Other function symbols. We introduce constants ph_2 (resp. ph_3) to represent the Voting (resp. Opening) Phase of the FOO protocol. These symbols are necessary to avoid *replay attacks* in which an attacker can learn honest vote value by forwarding one of the honest voter's message from the Voting Phase to the Opening Phase (see 3.2). Constants A and B model the identity of two honest voters and constant M models the identity of the mixer. C_1, C_2, C_3 model the identities of the candidates. (For simplicity, we assume only three candidates.) Finally, the function symbols $\text{to}(\cdot), \text{V}_0(\cdot), \text{V}_1(\cdot), \text{pubkey}(\cdot)$ are used to model the security game and explained in Section 5.2.

Axioms. We present axioms (see Table 2) that formalize the security assumptions on the primitives. The axioms use the notion of *freshness* [6]: For any list of names \vec{n} , and

a (possibly empty) list of terms \vec{v} , $\text{fresh}(\vec{n}; \vec{v})$ is the constraint that the names in \vec{n} are pairwise distinct and that none of the names in \vec{n} occur in \vec{v} .

The axiom `COMPHID` models the computational hiding property of the commitments (see Section 2.2). Intuitively, the axiom says that if the attacker is presented with commitments on equal length messages but the order of these messages is hidden, then the attacker cannot determine the order if it does not have access to commitment keys. `COMPHID` can be shown to be computationally valid if and only if $\llbracket \text{k}_c(-) \rrbracket$, $\llbracket \text{com}(-, -) \rrbracket$, $\llbracket \text{open}(-, -) \rrbracket$ satisfy the computational hiding property.

The axiom `BLINDNESS` models the blindness property of the blind signatures (see Section 2.2). Intuitively, in the blindness game, the dishonest signer attacker is asked to sign two blinded messages (could be of unequal length), but the order of these messages is hidden. Thus, the blinded messages on the two sides of \sim are in reversed order. Moreover, the signer is also allowed to see the unblinded signatures if the user does not abort for either signature. However, in order to hide the order, the unblinded signatures are presented in the same order. `BLINDNESS` can be shown to be computationally valid if and only if blind signature scheme $(\llbracket \text{k}_b \rrbracket, \llbracket \text{r}_b \rrbracket, \llbracket \text{b} \rrbracket, \llbracket \text{ub} \rrbracket, \llbracket \text{bsign} \rrbracket, \llbracket \text{acc} \rrbracket, \llbracket \text{bver} \rrbracket)$ satisfies the blindness property.

The axiom `SHUFFLE` models the shuffling capability of the mixer. We also require further assumptions on the primitives. In particular, the commitments and pairing should be length regular, unblinded signatures on equal length messages by the same signer must have equal length, honestly generated commitment keys must have equal length, candidate names should have identical length and phase identifiers should be distinct but have identical length. For lack of space, the formalizations of these properties is given in Appendix C.

Finally, we recall the axiom representing IND-CCA2 property of encryptions from [10]. Let $\vec{t}[x]$ be a list of terms with a single variable x . For a closed term v , let $\vec{t}[v]$ denote the term that we receive from $\vec{t}[x]$ by replacing all occurrences of x by v . Let u, u', u'' be closed terms. Consider the formula

$$\vec{t}[\text{if EQL}(u, u') \text{ then } \{u\}_{\text{pk}(n_1)}^{\text{r}_e(n_2)} \text{ else } u''] \sim \vec{t}[\text{if EQL}(u, u') \text{ then } \{u'\}_{\text{pk}(n_1)}^{\text{r}_e(n_3)} \text{ else } u'']$$

in which $n_1 \in \mathcal{N}$ occurs only as $\text{k}(n_1)$, $\text{sk}(n_1)$ only occurs in decryption position (that is, as in $\text{dec}(-, \text{sk}(n_1))$), and n_2, n_3 do not occur anywhere else. We call the above formula `ENC_CCA2` if for any $t'[x]$ term with x explicitly occurring in $t'[x]$, $\text{dec}(t'[x], \text{sk}(n_1))$ occurs only as $\text{if EQ}(t'[x], x) \text{ then } t''[x] \text{ else } \text{dec}(t'[x], \text{sk}(n_1))$, where $t''[x]$ is not of the form $\text{dec}(t'''[x], \text{sk}(n_1))$. The intuition is that since in the IND-CCA2 game, after encryption, the decryption oracle decrypts only those messages that are different from the encryption, we have to make sure that the decrypted message, $t'[x]$ is different from the encryption, for which x stands. The reader is referred to [10] for additional details.

5.2 Modeling the Vote Privacy Security Game

As defined in [6], the BC technique treats protocols as abstract transition systems. We do not recall the formal definitions, we just apply them to the vote privacy games Π_b for $b = 0, 1$ between the Challenger and the Attacker as defined in Section 3.1. The transition system produces the terms of the execution, on which the indistinguishability

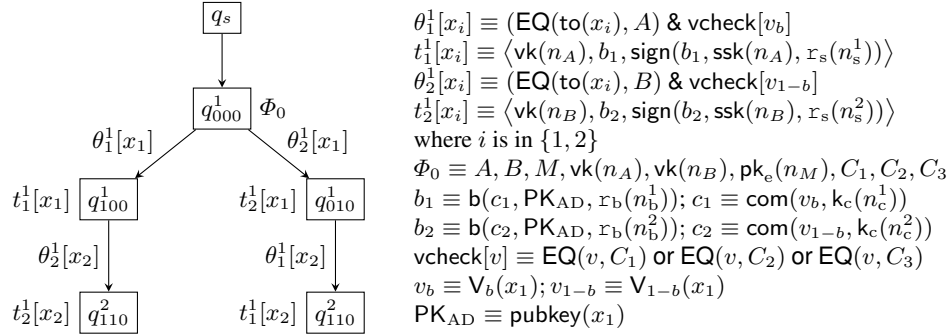


Fig. 2. Authentication phase of the security game Π_b

which increases attacker's knowledge. When the checks in a transition fail, the transition moves to state q_{exit} , which we omit in transition diagram for clarity. The transition systems for other phases can be similarly defined (see Appendix D for more details).

We obtain the execution terms from the transition system in the following way. For a given b bit of the Challenger, Φ_i^b lists of terms representing what the attacker has seen up to step i are created according to the rounds of the protocol. Notice that the initialization frame Φ_0 is defined in Figure 2 and it is independent of the bit b . For the first step, the attacker's computation f_1 is applied to Φ_0 , and $f_1(\Phi_0)$ is sent to the Challenger, who then carries out the checks θ_1^1 and θ_2^1 , and sends back

$$\phi_1^b := \text{if } \theta_1^1[f_1(\Phi_0)] \text{ then } t_1^1[f_1(\Phi_0)] \text{ else (if } \theta_2^1[f_1(\Phi_0)] \text{ then } t_2^1[f_1(\Phi_0)] \text{ else } \mathbf{0}),$$

and we set $\Phi_1^b \stackrel{\text{def}}{=} \Phi_0, \phi_1^b$. Similarly, in the second step, we have

$$\begin{aligned} \phi_2^b := & \text{if } \theta_1^1[f_1(\Phi_0)] \\ & \text{then (if } \theta_2^1[f_2(\Phi_1^b)] \text{ then } t_2^1[f_2(\Phi_1^b)] \text{ else } \mathbf{0}) \\ & \text{else (if } \theta_2^1[f_1(\Phi_0)] \text{ then (if } \theta_2^1[f_2(\Phi_1^b)] \text{ then } t_2^1[f_2(\Phi_1^b)] \text{ else } \mathbf{0}) \text{ else } \mathbf{0}), \end{aligned}$$

and $\Phi_2^b \stackrel{\text{def}}{=} \Phi_1^b, \phi_2^b$, which is the end of the authentication phase. The other phases are done similarly continuing from Φ_3^b .

Let Φ_m^b be the last frame. In order to establish vote-privacy for the FOO protocol we have to show that the axioms and first-order inference rules imply that $\Phi_m^0 \sim \Phi_m^1$. Then from the soundness theorem of [6], it follows that there is no successful PPT attacker break the security game. We have the following (See Section E for the proof).

Theorem 1. *The modified FOO protocol respects vote privacy for one session.*

6 Conclusions

We analyzed the FOO electronic voting protocol for vote privacy in the provable security model using the computationally complete symbolic attacker (CCSA) framework. As part of the analysis, we showed that security properties of trapdoor commitments and blind signatures can be faithfully translated into axioms in the CCSA framework. We demonstrated that the framework is effective in that it revealed new attacks on the

FOO protocol and could be used to prove the modified FOO protocol secure. As part of future work, we plan to investigate expressing and verifying stronger privacy properties of receipt-freeness and coercion-resistance for electronic voting protocols in the CCSA framework. We also plan to investigate automation of the verification tasks.

References

1. Fujioka, A., Okamoto, T., Ohta, K.: A Practical Secret Voting Scheme for Large Scale Elections. In: ASIACRYPT '92. LNCS, Springer (1993) 244–251
2. Kremer, S., Ryan, M.: Analysis of an electronic voting protocol in the applied pi calculus. In: ESOP, Springer (2005) 186–200
3. Abadi, M., Blanchet, B., Fournet, C.: The applied pi calculus: Mobile values, new names, and secure communication. *J. ACM* (2017) 1–41
4. Delaune, S., Ryan, M., Smyth, B.: Automatic verification of privacy properties in the applied pi calculus. In: IFIPTM. (2008) 263–278
5. Chadha, R., Cheval, V., Ciobăcă, c., Kremer, S.: Automated verification of equivalence properties of cryptographic protocols. *ACM Trans. Comput. Logic* (2016) 1–32
6. Bana, G., Comon-Lundh, H.: A Computationally Complete Symbolic Attacker for Equivalence Properties. In: CCS '14, ACM (2014) 609–620
7. Bana, G., Comon-Lundh, H.: Towards Unconditional Soundness: Computationally Complete Symbolic Attacker. In: POST. LNCS, Springer (2012) 189–208
8. Bana, G., Adão, P., Sakurada, H.: Computationally Complete Symbolic Attacker in Action. In: FSTTCS'12. LIPIcs, Schloss Dagstuhl (2012) 546–560
9. Bana, G., Adão, P., Sakurada, H.: Symbolic Verification of the Needham-Schroeder-Lowe Protocol (2012) <http://web.ist.utl.pt/pedro.adao/pubs/drafts/nsl-long.pdf>.
10. Bana, G., Chadha, R.: Verification methods for the computationally complete symbolic attacker based on indistinguishability. *Cryptology ePrint Archive, Report 2016/069* (2016) <http://eprint.iacr.org/2016/069>.
11. Comon, H., Koutsos, A.: Formal computational unlinkability proofs of rfid protocols. In: CSF, IEEE (2017) 100–114
12. Scerri, G., Stanley-Oakes, R.: Analysis of key wrapping apis: Generic policies, computational security. In: CSF, IEEE (2016) 281–295
13. Chaum, D.L.: Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM* (1981) 84–90
14. Benaloh, J., Tuinstra, D.: Receipt-free secret-ballot elections (extended abstract). In: ACM SoTC. STOC (1994) 544–553
15. Benaloh, J.D.C.: Verifiable Secret-ballot Elections. PhD thesis, Yale University (1987)
16. Juels, A., Luby, M., Ostrovsky, R.: Security of blind digital signatures. In: CRYPTO, Springer (1997) 150–164
17. Smyth, B., Bernhard, D.: Ballot secrecy and ballot independence coincide. In: ESORICS, Springer (2013) 463–480
18. Bernhard, D., Cortier, V., Galindo, D., Pereira, O., Warinschi, B.: Sok: A comprehensive analysis of game-based ballot privacy definitions. In: IEEE S&P. (2015) 499–516
19. Cortier, V., Dragan, C.C., Dupressoir, F., Schmidt, B., Strub, P., Warinschi, B.: Machine-checked proofs of privacy for electronic voting protocols. In: IEEE S&P. (2017) 993–1008
20. Fischlin, M., Schröder, D.: Security of blind signatures under aborts. In: Jarecki, S., Tsudik, G., eds.: *Public Key Cryptography*, Springer (2009) 297–316
21. Chaum, D., Rivest, R.L., Sherman, A.T.: Blind Signatures for Untraceable Payments. In: *Advances in Cryptology: Crypto 82*, Springer (1983) 199–203

22. Abdalla, M., Nampreppe, C., Neven, G.: On the (im)possibility of blind message authentication codes. In Pointcheval, D., ed.: Topics in Cryptology, Springer (2006) 262–279
23. Smart, N.P.: Cryptography Made Simple. Information Security and Cryptography. Springer (2016)
24. Damgård, I.B., Pedersen, T.P., Pfitzmann, B.: On the existence of statistically hiding bit commitment schemes and fail-stop signatures. Journal of Cryptology (1997) 163–194
25. Naor, M.: Bit commitment using pseudorandomness. J. Cryptol. (1991) 151–158

A Formal definitions of computationally hiding and blindness

Definition 1. Computationally Hiding [23]. A commitment scheme $\mathcal{C} = (KG, Commit, Open)$ is computationally hiding [23] if for any PPT algorithm \mathcal{A} (working in modes *find* and *guess*) the probability that the following experiment $Hiding_{\mathcal{A}}^{\mathcal{C}}$ evaluates to 1 is negligibly different from $1/2$.

Experiment $Hiding_{\mathcal{A}}^{\mathcal{C}}(\eta)$
 $(m_0, m_1, \beta_{find}) \leftarrow \mathcal{A}(find, 1^\eta)$ (m_0 and m_1 are of same length)
 $b \xleftarrow{\$} \{0, 1\}$, $k_0 \leftarrow KG(1^\eta)$, $k_1 \leftarrow KG(1^\eta)$,
 $com_b \leftarrow Commit(m_b, k_0)$, $com_{1-b} \leftarrow Commit(m_{1-b}, k_1)$
 $b^* \leftarrow \mathcal{A}(guess, \beta_{find}, com_b, com_{1-b})$
 return 1 iff $b = b^*$

Definition 2. Blindness [16]. An interactive signature scheme $\mathcal{B} = (Gen, \langle \mathcal{S}, \mathcal{U} \rangle, \text{verif})$ satisfies blindness [16] if for any efficient algorithm \mathcal{S}^* (working in modes *find*, *issue* and *guess*) the probability that the following experiment $Blind_{\mathcal{S}^*}^{\mathcal{B}}(\eta)$ evaluates to 1 is negligibly close to $1/2$, where $\mathcal{S}^*(\cdot, \mathcal{U}(pk, m_b))^1, \mathcal{U}(pk, m_{1-b})^1$ represents that the signer \mathcal{S}^* can invoke arbitrary ordered executions with $\mathcal{U}(pk, m_b)$ and $\mathcal{U}(pk, m_{1-b})$ once:

Experiment $Blind_{\mathcal{S}^*}^{\mathcal{B}}(\eta)$:
 $(pk, m_0, m_1, st_{find}) \leftarrow \mathcal{S}^*(find, 1^\eta)$
 $b \xleftarrow{\$} \{0, 1\}$, $st_{issue} \leftarrow \mathcal{S}^*(\cdot, \mathcal{U}(pk, m_b))^1, \mathcal{U}(pk, m_{1-b})^1 (issue, st_{find})$
 Let $\sigma_b \leftarrow \mathcal{U}(pk, m_b)$, $\sigma_{1-b} \leftarrow \mathcal{U}(pk, m_{1-b})$, $(\sigma_0, \sigma_1) := (\perp, \perp)$ if $\sigma_0 = \perp$ or $\sigma_1 = \perp$
 $b^* \leftarrow \mathcal{S}^*(guess, \sigma_0, \sigma_1, st_{issue})$
 return 1 iff $b = b^*$

B Core axioms

We recall the core axioms from [10] in the table 3 below. The axioms use the notion of *freshness* that has been defined in [6] and is reproduced in Section 5.

C Additional Axioms

Additional axioms used in the security game for the FOO protocol are given in Table 4. COMMEQL models the assumption that the length of commitments on equal length messages are equal. COMMKEYEQL models the assumption that the length of honestly generated commitment keys are equal. UBNOTUNDEFINED models the assumption that if

Axioms for indistinguishability.	
REFL:	$\vec{x} \sim \vec{x}$
SYM:	$\vec{x} \sim \vec{y} \longrightarrow \vec{y} \sim \vec{x}$
TRANS:	$\vec{x} \sim \vec{y} \wedge \vec{y} \sim \vec{z} \longrightarrow \vec{x} \sim \vec{z}$
RESTR:	for any projection and permutation $p, \vec{x} \sim \vec{y} \longrightarrow p(\vec{x}) \sim p(\vec{y})$
FUNCAPP:	for any $f \in \mathcal{F} \cup \mathcal{G}, \vec{x} \sim \vec{y} \longrightarrow \vec{x}, f(\vec{x}) \sim \vec{y}, f(\vec{y})$
TFDIST:	$\neg(\mathbf{true} \sim \mathbf{false})$
Axioms for equality.	
EQCONG:	$=$ is a congruence relation
EQTHEO:	$x = y$ if $x =_E y$
Axioms for if _ then _ else _ .	
IFSAME:	if b then x else $x = x$
IFEVAL:	for any t_1, t_2 terms, if b then $t_1[b]$ else $t_2[b] =$ if b then $t_1[\mathbf{true}]$ else $t_2[\mathbf{false}]$
IFTRUE:	if \mathbf{true} then x else $y = x$
IFFALSE:	if \mathbf{false} then x else $y = y$
IFBRANCH:	$\vec{z}, b, x_1, \dots, x_n \sim \vec{z}', b', x'_1, \dots, x'_n \wedge \vec{z}, b, y_1, \dots, y_n \sim \vec{z}', b', y'_1, \dots, y'_n$ \longrightarrow $\vec{z}, b, \begin{matrix} \text{if } b \text{ then } x_1 \\ \text{else } y_1 \end{matrix}, \dots, \begin{matrix} \text{if } b \text{ then } x_n \\ \text{else } y_n \end{matrix} \sim \vec{z}', b', \begin{matrix} \text{if } b' \text{ then } x'_1 \\ \text{else } y'_1 \end{matrix}, \dots, \begin{matrix} \text{if } b' \text{ then } x'_n \\ \text{else } y'_n \end{matrix}$
Axioms for names.	
FRESHIND:	Let n_1, n_2 be names and \vec{v}, \vec{w} be lists of ground terms such that $\text{fresh}(n_1; \vec{v}, \vec{w})$ and $\text{fresh}(n_2; \vec{v}, \vec{w})$ holds. We have $\vec{v} \sim \vec{w} \longrightarrow n_1, \vec{v} \sim n_2, \vec{w}$.
FRESHNEQ:	for any name n and a closed term v such that $\text{fresh}(n; v)$ holds, we have $\text{EQ}(n, v) \sim \mathbf{false}$.

Table 3: Core Axioms

a user successfully completes the blind signature protocol then the output is different from \perp . UBEQL models the assumption that blind signatures on equal length message are of equal length. PAIREQL models the length regularity of pairing. CANDEQL (resp. PHASEQL) models the assumptions that candidate identities (resp. phase identities) are of equal length. AGENTDIST (resp. PHASEDIST) models the assumption that identities of honest voters (resp. phase identifiers) are distinct.

D Transition Systems for Opening and Voting phases

The Voting Phase is in Figure 2 where we use similar conventions. It starts from the state q_{110}^2 which is final state of Authentication Phase. Again, the left branch simulates the situation where the voter A receives message x_3 from an attacker and it is accepted as a blinded commitment which is signed by administrator. Then A unblinds, encrypts the unblinded message with mixer's public key $k_e(n_M)$, and sends out the

Commitments

Let t_1 and t_2 be two ground terms. Let n_1 and n_2 be names such that $\text{fresh}(n_1, n_2; t_1, t_2)$ holds.

COMMEQL : if EQL(t_1, t_2) then EQL($\text{com}(t_1, k_c(n_1)), \text{com}(t_2, k_c(n_2))$) else **true** = **true**

COMMKEYEQL : EQL($k_c(n_1), k_c(n_2)$) = **true**

Blind Digital Signatures

For $i = 0, 1$, let t, t_1^i , and t_2^i be ground terms. Let n^i be a name with

$\text{fresh}(n^i; t, t_1^i, t_2^i, t_1^{1-i})$, such that it occurs in t_2^i , only as $b(t_1^i, t, r_b(n^i))$:

UBNOTUNDEFINED :

if $\text{acc}(t_1^0, t, r_b(n), t_2^0)$ then EQ($\text{ub}(t_1^0, t, r_b(n^0), t_2^0), \perp$) else **false** = **false**

UBEQL :

if EQL(t_1^0, t_1^1) & EQL(t_2^0, t_2^1) & $\text{acc}(t_1^0, t, r_b(n^0), t_2^0)$ & $\text{acc}(t_1^1, t, r_b(n^0), t_2^1)$
then EQL($\text{ub}(t_1^0, t, r_b(n^0), t_2^0), \text{ub}(t_1^1, t, r_b(n^1), t_2^1)$) = **true**
else **true**

Further length regularity

PAIREQL : if EQL(x_1, y_1) & EQL(x_2, y_2) then EQL($\langle x_1, x_2 \rangle, \langle y_1, y_2 \rangle$) else **true** = **true**

CANDEQL : EQL(C_1, C_2) = **true** for $i, j \in \{1, 2, 3\}$

PHASEEQL : EQL(ph_2, ph_3) = **true**

Distinctness of constants

AGENTDIST : EQ(A, B) = EQ(A, M) = EQ(B, M) = **false**

PHASEDIST : EQ(ph_2, ph_3) = **false**

Table 4: Additional Axioms for the FOO protocol

encrypted message $t_1^2[x_3]$. And so forth, following the security game. The abbreviation $\langle x_1, x_2, x_3 \rangle \stackrel{\text{def}}{=} \langle x_1, \langle x_2, x_3 \rangle \rangle$ encodes triples and $\tau_1(x) \stackrel{\text{def}}{=} \pi_1(x), \tau_2(x) \stackrel{\text{def}}{=} \pi_1(\pi_2(x))$ and $\tau_3(x) \stackrel{\text{def}}{=} \pi_2(\pi_2(x))$ encode projections from triples.

In order to formalize the Opening Phase, we introduce the following abbreviations:

- $\text{dist}(x, y, z) \stackrel{\text{def}}{=} \text{not}(\text{EQ}(x, y)) \ \& \ \text{not}(\text{EQ}(x, z)) \ \& \ \text{not}(\text{EQ}(y, z)).$
- $\text{isin}(x, y) \stackrel{\text{def}}{=} \text{EQ}(x, \tau_1(y)) \ \text{or} \ \text{EQ}(x, \tau_2(y)) \ \text{or} \ \text{EQ}(x, \tau_3(y))$
- $\text{bcheck}(x, y) \stackrel{\text{def}}{=} \text{isin}(x, \langle \pi_2(\tau_1(y)), \pi_2(\tau_2(y)), \pi_2(\tau_3(y)) \rangle)$
-

$\text{label}(x, y) \stackrel{\text{def}}{=} \text{if} \ \text{EQ}(x, \pi_2(\tau_1(y))) \ \text{then} \ \pi_1(\tau_1(y))$
 $\text{else if} \ \text{EQ}(x, \pi_2(\tau_2(y))) \ \text{then} \ \pi_1(\tau_2(y))$
 $\text{else if} \ \text{EQ}(x, \pi_3(\tau_2(y))) \ \text{then} \ \pi_1(\tau_3(y))$
 $\text{else } 0$

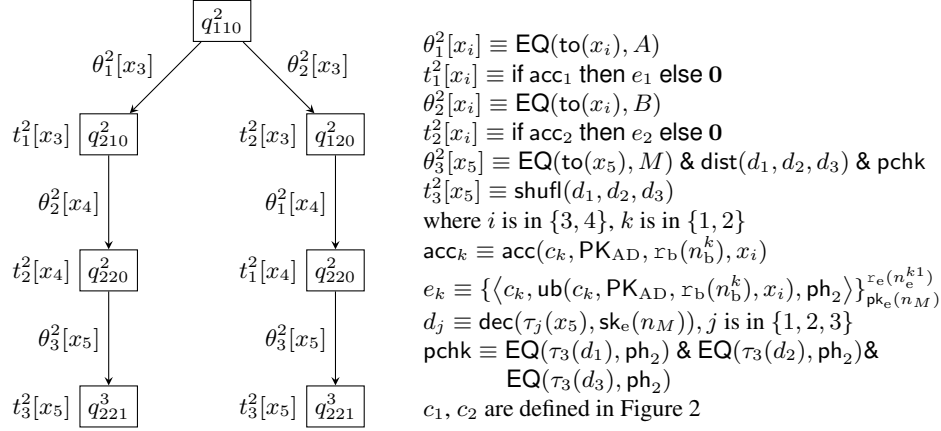


Fig. 3. Voting Phase of the FOO protocol

Here $\text{isin}(x, y)$ checks that if x occur in message y which itself is a triple. $\text{bcheck}(-, -)$ models the voter's check that if his (her) ballot presents in the list of label-ballot pairs. $\text{label}(x, y)$ is the finding the label of x in the message y which itself is a tuple of label-ballot pairs.

Finally, the Opening Phase is in Figure 4.

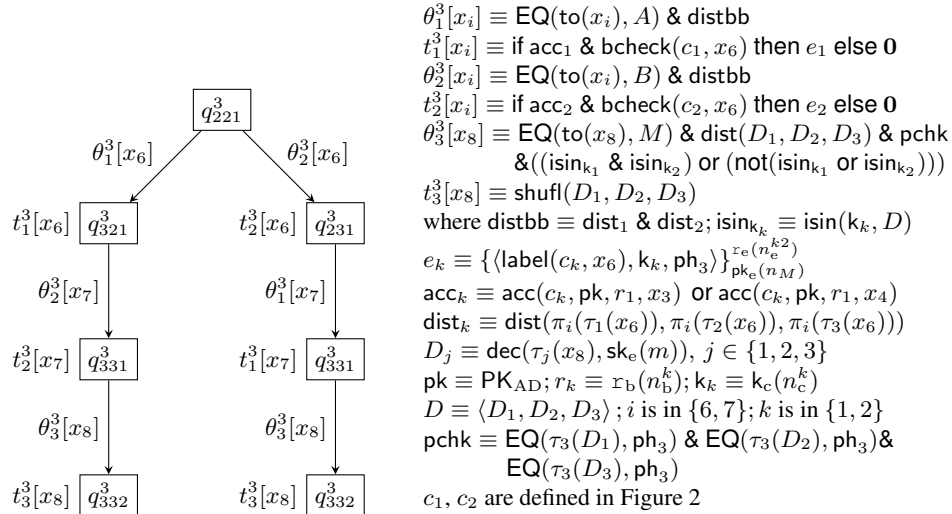


Fig. 4. Opening Phase of the FOO protocol

E Sketch of the Proof of Vote Privacy

Instead of showing the full security proof, we show how the combination of `COMP hid`, `COMMEQL`, `UBNOTUNDEFINED` and `BLINDNESS` allows us to fix the inadequacy of the blindness property. The key idea is that as the length of candidate identities are equal, the commitments hide the underlying vote. Thus, the probability that the attacker can cause the blind signature for candidate v_b is accepted but rejected for v_{1-b} is negligibly small. This is formalized in the following Proposition which combines `COMP hid` and `BLINDNESS`.

Proposition 1. Let v_0, v_1, t be ground terms, let \vec{z} be a list of ground terms. Assume $\text{EQL}(v_0, v_1) = \text{true}$. Let t_0 and t_1 be terms containing two variables x and y . Let $n_{b0}, n_{b1}, n_{c0}, n_{c1}$ be names such that $\text{fresh}(n_{b0}, n_{b1}, n_{c0}, n_{c1}; \vec{z}, t, v_0, v_1, t_0, t_1)$ holds. Suppose that the blind signatures satisfy `BLINDNESS` and `UBNOTUNDEFINED`, and commitments satisfy `COMP hid` and `COMMEQL` properties. Then

$$\vec{z}, \mathbf{b}(c_0^0[v_0], t, r_b(n_{b0})), \mathbf{b}(c_1^0[v_1], t, r_b(n_{b1})), t^0 \sim \vec{z}, \mathbf{b}(c_0^1[v_1], t, r_b(n_{b0})), \mathbf{b}(c_1^1[v_0], t, r_b(n_{b1})), t^1 \quad (1)$$

$$\begin{aligned} \text{where } t^0 &\equiv \text{if } \text{acc}_0^0 \ \& \ \text{acc}_1^0 \ \text{then } \langle \langle \text{ub}_0^0[v_0], c_0^0[v_0], k_c(n_{c0}) \rangle, \langle \text{ub}_1^0[v_1], c_1^0[v_1], k_c(n_{c1}) \rangle \rangle \\ &\quad \text{else if } \text{acc}_0^0 \ \text{then } \langle \langle \text{ub}_0^0[v_0], c_0^0[v_0], k_c(n_{c0}) \rangle, \perp \rangle \\ &\quad \quad \text{else if } \text{acc}_1^0 \ \text{then } \langle \perp, \langle \text{ub}_1^0[v_1], c_1^0[v_1], k_c(n_{c1}) \rangle \rangle \\ &\quad \quad \quad \text{else } \langle \perp, \perp \rangle \\ t^1 &\equiv \text{if } \text{acc}_0^1 \ \& \ \text{acc}_1^1 \ \text{then } \langle \langle \text{ub}_1^1[v_0], c_1^1[v_0], k_c(n_{c1}) \rangle, \langle \text{ub}_0^1[v_1], c_0^1[v_1], k_c(n_{c0}) \rangle \rangle \\ &\quad \text{else if } \text{acc}_1^1 \ \text{then } \langle \langle \text{ub}_1^1[v_0], c_1^1[v_0], k_c(n_{c1}) \rangle, \perp \rangle \\ &\quad \quad \text{else if } \text{acc}_0^1 \ \text{then } \langle \perp, \langle \text{ub}_0^1[v_1], c_0^1[v_1], k_c(n_{c0}) \rangle \rangle \\ &\quad \quad \quad \text{else } \langle \perp, \perp \rangle \end{aligned}$$

Where $t_i^j = t_i[x \leftarrow \mathbf{b}(c_0^j[v_j], t, r_b(n_{b0})), y \leftarrow \mathbf{b}(c_1^j[v_{1-j}], t, r_b(n_{b1}))]$, $i, j \in \{0, 1\}$,

$$\begin{aligned} c_i^0[v_i] &\equiv \text{com}(v_i, k_c(n_{ci})) & c_i^1[v_{1-i}] &\equiv \text{com}(v_{1-i}, k_c(n_{ci})) \\ \text{acc}_i^0 &\equiv \text{acc}(c_i^0[v_i], t, r_b(n_{bi}), t_i^0) & \text{acc}_i^1 &\equiv \text{acc}(c_i^1[v_{1-i}], t, r_b(n_{bi}), t_i^1) \\ \text{ub}_i^0[v_i] &\equiv \text{ub}(c_i^0[v_i], t, r_b(n_{bi}), t_i^0) & \text{ub}_i^1[v_{1-i}] &\equiv \text{ub}(c_i^1[v_{1-i}], t, r_b(n_{bi}), t_i^1) \end{aligned}$$

The proof idea is the following. Let us further introduce the notation

$u_i^j[v_{j \oplus i}] \equiv \langle \text{ub}_i^j[v_{j \oplus i}], c_i^j[v_{j \oplus i}], k_c(n_{ci}) \rangle$ and $b_i^j[v_{j \oplus i}] \equiv \mathbf{b}(c_i^j[v_{j \oplus i}], t, r_b(n_{bi}))$ where \oplus is XOR. Then using the axiom `IFBRANCH`, it is sufficient to show the following equivalences:

$$\begin{aligned} &\vec{z}, b_0^0[v_0], b_1^0[v_1], \text{acc}_0^0 \ \& \ \text{acc}_1^0, \ \text{if } \text{acc}_0^0 \ \& \ \text{acc}_1^0 \ \text{then } \langle u_0^0[v_0], u_1^0[v_1] \rangle \ \text{else } \langle \perp, \perp \rangle \\ &\sim \vec{z}, b_0^1[v_1], b_1^1[v_0], \text{acc}_0^1 \ \& \ \text{acc}_1^1, \ \text{if } \text{acc}_0^1 \ \& \ \text{acc}_1^1 \ \text{then } \langle u_1^1[v_0], u_0^1[v_1] \rangle \ \text{else } \langle \perp, \perp \rangle \end{aligned}$$

which follows from the `BLINDNESS`, `UBNOTUNDEFINED` and `FUNCAPP`; and the equivalences

$$\begin{aligned} &\vec{z}, b_0^0[v_0], b_1^0[v_1], \text{acc}_0^0 \ \& \ \text{acc}_1^0, \ \text{if } \text{acc}_i^0 \ \text{then } u_i^0[v_i] \ \text{else } \perp \\ &\sim \vec{z}, b_0^1[v_1], b_1^1[v_0], \text{acc}_0^1 \ \& \ \text{acc}_1^1, \ \text{if } \text{acc}_{1-i}^1 \ \text{then } u_{1-i}^1[v_i] \ \text{else } \perp \end{aligned}$$

which follow from `COMP hid` and `FUNCAPP`. However, the axiom itself is not sufficient, we need the condition that the votes have equal length, otherwise these cannot be proven and an attack is constructed, which is the first attack described in Section 3.2.

Of course, in the actual security game, encryptions are sent too where the votes are also in reversed order, and decryptions are done by the mixer. This is dealt with the `ENCCCA2` axiom. For the applicability of the axiom, the terms sent have to be rewritten using the equality to introduce cases when the adversarial function symbols on

which decryptions are applied are either of the previous encryptions (namely the ones sent by voter A and B), or when they are neither. More precisely, a term of the form $\text{dec}(f(\Phi), \text{sk}(n_1))$, assuming that Φ has say two honest encryptions in them, $\{t\}_{\text{pk}(n_1)}^{r(n)}$ and $\{t'\}_{\text{pk}(n_1)}^{r(n')}$, has to be rewritten as

$$\text{dec}(f(\Phi), \text{sk}(n_1)) = \begin{array}{l} \text{if EQ}(f(\Phi), \{t\}_{\text{pk}(n_1)}^{r(n)}) \\ \text{then } t \\ \text{else if EQ}(f(\Phi), \{t'\}_{\text{pk}(n_1)}^{r(n')}) \text{ then } t' \text{ else } \text{dec}(f(\Phi), \text{sk}(n_1)) \end{array} .$$

Thus, when the attacker messages are equal to some encrypted messages sent, the decryptions can be replaced by the plaintext. When they are not equal, then the decryption on the function symbol is kept. Once this is done, the ENC_{CCA2} axiom can be applied, and the plaintexts inside the encryptions can now be switched as long as they are of equal length. This is why we needed to assume that the unblinding results equal length. Without that assumption we obtain the second attack in Section 3.2.

Finally, when the decryptions are applied on the correct previous encryptions, there are still several possibilities for each decryption because a priori the adversary could redirect the same encryption twice. A priori, each decryption will succeed on any of the previous (four) encryptions sent by A and B . To ensure that there is no clash, we made sure that the mixer checks both the phase and that all the decryptions are different. Otherwise we obtain the third attack of Section 3.2.