

Attacking and fixing Helios: An analysis of ballot secrecy*

Véronique Cortier and Ben Smyth

Loria, CNRS & INRIA Nancy Grand Est, France

June 1, 2011

Abstract

Helios 2.0 is an open-source web-based end-to-end verifiable electronic voting system, suitable for use in low-coercion environments. In this article, we analyse ballot secrecy and discover a vulnerability which allows an adversary to compromise the privacy of voters. The vulnerability exploits the absence of ballot independence in Helios and works by replaying a voter's ballot or a variant of it, the replayed ballot influences the election outcome, introducing information that can be used to violate privacy. We demonstrated the practicality of the attack by breaking privacy in a mock election using the current Helios implementation. Moreover, the feasibility of an attack is considered in the context of French legislative elections and, based upon our findings, we believe it constitutes a real threat to ballot secrecy in such settings. We present a fix and show that our solution satisfies a formal definition of ballot secrecy using the applied pi calculus. In addition, we present attacks against other electronic voting schemes which do not assure ballot independence; in particular, we show a similar vulnerability in the protocol by Lee *et al.* and demonstrate replay attacks against the schemes by Sako & Kilian and Schoenmakers. Finally, we argue that no general relationships exist between independence and privacy properties.

Keywords. Applied Pi Calculus, Attack, Ballot Independence, Ballot Secrecy, Electronic Voting, Fiat-Shamir Heuristic, Helios, Malleability, Privacy, Replay.

1 Introduction

Paper-based elections derive security properties from physical characteristics of the real-world. For example, marking a ballot in isolation inside a polling booth

*This article is an extension of our earlier work [CS11].

and depositing the completed ballot into locked ballot box provides privacy; the polling booth also ensures that voters cannot be influenced by other votes and the locked ballot box prevents the announcement of early results, thereby ensuring fairness; and the transparency of the whole election process from ballot casting to tallying and the impossibility of altering the markings on a paper ballot sealed inside a locked ballot box gives an assurance of correctness and facilitates verifiability. Replicating these attributes in a digital setting has proven to be difficult and, hence, the provision of secure electronic voting systems is an active research topic.

Informally, *privacy* for electronic voting systems is characterised by the following requirements [KR05, DKR06, BHM08]:

- *Ballot secrecy.* A voter's vote is not revealed to anyone.
- *Receipt freeness.* A voter cannot gain information which can be used to prove, to a coercer, how she voted.
- *Coercion resistance.* A voter cannot collaborate, with a coercer, to gain information which can be used to prove how she voted.

Verifiability includes three properties [JCJ05, Dag07, KRS10]:

- *Individual verifiability.* A voter can check that her own ballot is published on the election's bulletin board.
- *Universal verifiability.* Anyone can check that all the votes in the election outcome correspond to ballots published on the election's bulletin board.
- *Eligibility verifiability.* Anyone can check that each ballot published on the bulletin board was cast by a registered voter and at most one ballot is tallied per voter.

Finally, *fairness* – summarised by the notion that *all voters are equal* – has not been thoroughly studied, but nonetheless we believe the following aspects are desirable:

- *Ballot independence.* Observing another voter's interaction with the election system does not allow a voter to cast a *related* vote.
- *No early results.* A voter cannot change her vote once partial results are available.
- *Pulling out.* Once partial results are available a voter cannot abort.

The fairness property prohibits the voting system from influencing a voter's vote; more formally, this requires that observation of the voting system (that is, observing interaction between participants) does not leak information that may affect a voter's decision. The individual, universal and eligibility verifiability properties (also called *end-to-end verifiability* [JCJ02, CRS05, Adi06, Dag07, Adi08]) allow voters and election observers to verify – independently of the

hardware and software running the election – that votes have been recorded, tallied and declared correctly. In this article, we analyse ballot secrecy in Helios 2.0 [AMPQ09].

Helios is an open-source web-based electronic voting system. The scheme is claimed to satisfy ballot secrecy [AMPQ09], but the nature of remote voting makes the possibility of satisfying stronger privacy properties difficult, and Helios does not satisfy receipt freeness nor coercion resistance. In addition to ballot secrecy, the system provides individual and universal verifiability (cf. [KRS10, SRKK10] and [Smy11, Chapter 3] for an analysis of verifiability in Helios). Helios is particularly significant due to its real-world deployment: the International Association of Cryptologic Research used Helios to elect its board members [BVQ10], following a successful trial in a non-binding poll [HBH10]; the Catholic University of Louvain adopted the system to elect the university president [AMPQ09]; and Princeton University used Helios to elect the student vice president [Pri10].

Formal definitions of ballot secrecy have been introduced in the context of the applied pi calculus by Delaune, Kremer & Ryan [KR05, DKR06, DKR09, DKR10] and Backes, Hritcu & Maffei [BHM08]. These privacy definitions consider two voters \mathcal{A} , \mathcal{B} and two candidates t , t' . Ballot secrecy is captured by the assertion that an adversary (controlling arbitrary many dishonest voters) cannot distinguish between a situation in which voter \mathcal{A} votes for candidate t and voter \mathcal{B} votes for candidate t' , from another situation in which \mathcal{A} votes t' and \mathcal{B} votes t . This can be expressed by the following equivalence.

$$\mathcal{A}(t) \mid \mathcal{B}(t') \approx_t \mathcal{A}(t') \mid \mathcal{B}(t)$$

These formal definitions of ballot secrecy have been used by their respective authors to analyse the electronic voting protocols due to: Fujioka, Okamoto & Ohta [FOO92], Okamoto [Oka98], Lee *et al.* [LBD⁺04], and Juels, Catalano & Jakobsson [JCJ02, JCJ05, JCJ10]. It therefore seems natural to check whether Helios satisfies ballot secrecy as well.

Contribution. Our analysis of Helios reveals an attack which violates ballot secrecy. The attack exploits the system’s lack of ballot independence, and works by replaying a voter’s ballot or a variant of it (without knowing the vote contained within that ballot). Replayng a voter’s ballot immediately violates ballot secrecy in an election with three voters. For example, consider the electorate Alice, Bob, and Mallory; if Mallory replays Alice’s ballot, then Mallory can reveal Alice’s vote by observing the election outcome and checking which candidate obtained at least two votes. The practicality of this attack has been demonstrated by violating privacy in a mock election using the current Helios implementation. Furthermore, the vulnerability can be exploited in more realistic settings and, as an illustrative example, we discuss the feasibility of the attack in French legislative elections. This case study suggests there is a plausible threat to ballot secrecy. We also propose a variant of the attack which abuses the malleability of ballots to ensure replayed ballots are distinct: this makes identification of replayed ballots non-trivial (that is, checking for exact duplicates is insufficient).

Nonetheless, we fix the Helios protocol by identifying and discarding replayed ballots. We believe this solution is particular well-suited because it maintains Benaloh’s principle of ballot casting assurance [Ben06, Ben07] and requires a minimal extension to the Helios code-base. The revised scheme is shown to satisfy a formal definition of ballot secrecy using the applied pi calculus.

In addition, we demonstrate that the absence of ballot independence can be exploited in other electronic voting protocols to violate privacy; in particular, a similar attack is shown against the protocol by Lee *et al.* [LBD⁺04] whereby an adversary replays a voter’s ballot or a variant of it, and verbatim replay attacks are demonstrated against two schemes presented at CRYPTO (namely, the protocols due to Sako & Kilian [SK94] and Schoenmakers [Sch99]). Finally, we present some evidence to support the hypothesis that independence does not imply privacy and vice-versa.

Related work. The concept of independence was introduced by Chor *et al.* [CGMA85] and the possibility of compromising security properties due to lack of independence has been considered, for example, by [CR87, DDN91, DDN00, Gen00]. In the context of electronic voting, Gennaro [Gen95] demonstrates that the application of the Fiat-Shamir heuristic in the Sako-Kilian electronic voting protocol [SK94] violates ballot independence, and Wikström [Wik06, Wik08] studies non-malleability for mixnets to achieve ballot independence. By comparison, we focus on the violation of ballot secrecy rather than fairness, and exploit the absence of ballot independence to compromise privacy. Similar results have been shown against mixnets [Pfi94].

Estehghari & Desmedt [ED10] claim to present an attack which undermines privacy and end-to-end verifiability in Helios. However, their attack is dependent on compromising a voter’s computer, a vulnerability which is explicitly acknowledged by the Helios specification [AMPQ09]: “*a specifically targeted virus could surreptitiously change a user’s vote and mask all of the verifications performed via the same computer to cover its tracks.*” Accordingly, [ED10] represents an exploration of known vulnerabilities rather than an attack.

Langer *et al.* [Lan10, LSBV10] and Volkamer & Grimm [VG10] also study privacy in Helios. Langer *et al.* propose a taxonomy of informal privacy requirements [Lan10, LSBV10, LSB⁺10] to facilitate a more fine-grained comparison of electronic voting systems; this framework is used to analyse Helios and the authors claim ballot secrecy is satisfied if the adversary only has access to public data [Lan10, LSBV10]. Volkamer & Grimm introduce the *k-resilience* metric [VG10, Vol09] to calculate the number of honest participants required for ballot secrecy in particular scenarios; this framework is used to analyse Helios and the authors claim ballot secrecy is satisfied if the software developers are honest and the key holders do not collude [VG10]. Contrary to these results, we show an attack against privacy. Our work highlights the necessity for rigorous mathematical analysis techniques for security protocols; in particular, we believe the erroneous results reported by Langer *et al.* were due to the use of informal methods, and the approach by Volkamer & Grimm failed because only

some particular scenarios were considered.

Delaune, Kremer & Ryan [DKR06, DKR09] have shown that the protocol by Lee *et al.* satisfies coercion resistance for two honest voters. However, for a stronger definition of coercion resistance, Küsters & Truderung [KT09] have demonstrated a forced abstention attack; in addition, Küsters & Truderung propose a variant of the scheme by Lee *et al.* which is claimed to satisfy this stronger definition. In this article, we show a new attack against the original Lee *et al.* protocol and show that the revised scheme by Küsters & Truderung might not be secure under reasonable assumptions.

Structure of this article. Section 2 presents the Helios electronic voting scheme. (We remark that this is the first cryptographic description of the Helios protocol in the literature and, hence, is an additional contribution of this article.) Section 3 describes our attack and some variants, in addition to a study of its feasibility in the context of French legislative elections. We propose several solutions for recovering privacy in Section 4 and prove that our adopted solution formally satisfies ballot secrecy in Section 5. Section 6 demonstrates that the absence of ballot independence can be similarly exploited in other electronic voting protocols to violate privacy. Finally, our conclusion appears in Section 7.

2 Background: Helios 2.0

Helios exploits the additive homomorphic [CDS94, CGS97, Sch09] and distributed decryption [Ped91, CP93] properties of ElGamal [ElG85]. We will recall these cryptographic details before presenting the Helios protocol.

2.1 Additive homomorphic ElGamal

Given cryptographic parameters (p, q, g) and a number $n \in \mathbb{N}$ of trustees, where p and q are large primes such that $q \mid p - 1$ and g is a generator of the multiplicative group \mathbb{Z}_p^* of order q , the following operations are defined by ElGamal.

Distributed key generation. Each trustee $i \in n$ selects a private key share $x_i \in_R \mathbb{Z}_q^*$ and computes a public key share $h_i = g^{x_i} \bmod p$. The public key is $h = h_1 \cdot \dots \cdot h_n \bmod p$.

Encryption. Given a message m and a public key h , select a random nonce $r \in_R \mathbb{Z}_q^*$ and derive the ciphertext $(a, b) = (g^r \bmod p, g^m \cdot h^r \bmod p)$.

Re-encryption. Given a ciphertext (a, b) and public key h , select a random nonce $r' \in_R \mathbb{Z}_q^*$ and derive the re-encrypted ciphertext $(a', b') = (a \cdot g^{r'} \bmod p, b \cdot h^{r'} \bmod p)$.

Homomorphic addition. Given two ciphertexts (a, b) and (a', b') , the homomorphic addition of plaintexts is computed by multiplication $(a \cdot a' \bmod p, b \cdot b' \bmod p)$.

Distributed decryption. Given a ciphertext (a, b) , each trustee $i \in n$ computes the partial decryption $k_i = a^{x_i}$. The plaintext $m = \log_g M$ is recovered from $M = b / (k_1 \cdot \dots \cdot k_n) \bmod p$.

The computation of a discrete logarithm $\log_g M$ is hard in general. However, if M is chosen from a restricted domain, then the complexity is reduced; for example, if M is an integer such that $0 \leq M \leq n$, then the complexity is $O(n)$ by linear search or $O(\sqrt{n})$ using the baby-step giant-step algorithm [Sha71] (see also [LL90, §3.1]).

For secrecy, each trustee $i \in n$ must demonstrate knowledge of a discrete logarithm $\log_g h_i$, that is, they prove that h_i has been correctly constructed; this prevents, for example, a trustee constructing their public key share $h_i = h$. For integrity of decryption, each trustee $i \in n$ must demonstrate equality between discrete logarithms $\log_g h_i$ and $\log_a k_i$; this prevents, for example, a trustee constructing the public key share $h_i = g^{m+x_i}$ and providing the partial decryption $k_i = a^{x_i}$. In addition, the voter must demonstrate that a valid vote has been encrypted. These proofs can be achieved using signatures of knowledge (see Appendix A for details).

2.2 Protocol description

An election is created by naming an election officer, selecting a set of trustees, and generating a distributed public key pair. The election officer publishes, on the bulletin board, the public part of the trustees' key (and proof of correct construction), the candidate list $\tilde{t} = (t_1, \dots, t_l) \cup \{\epsilon\}$ (where ϵ represents a vote of abstention), and the list of eligible voters $\tilde{id} = (id_1, \dots, id_n)$; the officer also publishes the *election fingerprint*, that is, the hash of these parameters. Informally, the steps that participants take during a run of Helios are as follows.

1. The voter launches a browser script that downloads the election parameters and recomputes the election fingerprint. The voter should verify that the fingerprint corresponds to the value published on the bulletin board. (This ensures that the script is using the trustees' public key; in particular, it helps prevent encrypting a vote with an adversary's public key. Such attacks have been discussed in the context of Direct Anonymous Attestation by Rudolph [Rud07]; although, the vulnerability was discounted, in the trusted computing setting, by Leung, Chen & Mitchell [LCM08].)
2. The voter inputs her vote $v \in \tilde{t}$ to the browser script, which creates a ballot consisting of her vote encrypted by the trustees' public key, and a proof that the ballot represents a permitted vote (this is needed because the ballots are never decrypted individually, in particular, it prevents multiple votes being encoded as a single ballot). The ballot is displayed to the voter.

3. The voter can audit the ballot to check if it really represents a vote for her chosen candidate; if she decides to do this, then the script provides her with the random data used in the ballot creation. She can then independently reconstruct her ballot and verify that it is indeed well-formed. The script provides some practical resistance against vote selling by refusing to cast audited ballots. See Benaloh [Ben06, Ben07] for further details on ballot auditing.
4. When the voter has decided to cast her ballot, the script submits it to the election officer. The election officer authenticates the voter and checks that she is eligible to vote. The election officer also verifies the proof and publishes the ballot, appended with the voter's identity id , on the bulletin board. (In practice, the election officer also publishes the hash of the ballot, we omit this detail for brevity.)
5. Individual voters can check that their ballots appear on the bulletin board and, by verifying the proof, observers are assured that ballots represent permitted votes.
6. After some predefined deadline, the election officer homomorphically combines the ballots and publishes the encrypted tally on the bulletin board. Anyone can check that tallying is performed correctly.
7. Each of the trustees publishes a partial decryption of the encrypted tally, together with a signature of knowledge proving the partial decryption's correct construction. Anyone can verify these proofs.
8. The election officer decrypts the tally and publishes the result. Anyone can check this decryption.

Formally, Step 2 is defined in Figure 1. (For simplicity the ballot construction algorithm in Figure 1 considers a vote $v \in \tilde{t}$, this can be generalised [AMPQ09] to consider a vote $\tilde{v} \subseteq \tilde{t}$.) Checking voter eligibility (Step 4) is beyond the scope of Helios and Adida *et al.* [AMPQ09] propose the use of existing infrastructure. The remaining steps follow immediately from the application of cryptographic primitives (see Section 2.1 for details).

2.3 Software implementation

Helios 3.0 is an extension of Helios 2.0 which adds numerous practical features, including: integration of authentication with various web-services (for example, Facebook, GMail and Twitter), bulk voter registration using pre-existing electoral rolls, and simplification of administration with multiple trustees. Helios 3.0 has been implemented and is publicly available: <http://heliosvoting.org/>.

Figure 1 Ballot construction by the browser script

Input: Cryptographic parameters (p, q, g) , public key h , candidate list $\tilde{t} = (t_1, \dots, t_l) \cup \{\epsilon\}$ and vote v .

Output: Encrypted vote $(a_1, b_1), \dots, (a_l, b_l)$, signatures of knowledge $(\bar{a}_1, \bar{b}_1, \bar{c}_1, \bar{s}_1, \bar{a}'_1, \bar{b}'_1, \bar{c}'_1, \bar{s}'_1), \dots, (\bar{a}_l, \bar{b}_l, \bar{c}_l, \bar{s}_l, \bar{a}'_l, \bar{b}'_l, \bar{c}'_l, \bar{s}'_l)$ and signature of knowledge $(\bar{a}, \bar{b}, \bar{c}, \bar{s}, \bar{a}', \bar{b}', \bar{c}', \bar{s}')$.

1. If $v \notin \tilde{t}$ then the script terminates.
2. Encode the vote v as a bitstring. For all $1 \leq i \leq l$, let

$$m_i = \begin{cases} 1 & \text{if } v = t_i \\ 0 & \text{otherwise} \end{cases}$$

3. The bitstring representing the vote is encrypted. For all $1 \leq i \leq l$, let

$$(a_i, b_i) = (g^{r_i} \bmod p, g^{m_i} \cdot h^{r_i} \bmod p)$$

where $r_i \in_R \mathbb{Z}_q^*$.

4. For all $1 \leq i \leq l$, let $(\bar{a}_i, \bar{b}_i, \bar{c}_i, \bar{s}_i, \bar{a}'_i, \bar{b}'_i, \bar{c}'_i, \bar{s}'_i)$ be a signature of knowledge demonstrating that the ciphertext (a_i, b_i) contains either 0 or 1, that is, each candidate can receive at most one vote.
 5. Let $(\bar{a}, \bar{b}, \bar{c}, \bar{s}, \bar{a}', \bar{b}', \bar{c}', \bar{s}')$ be a signature of knowledge demonstrating that the ciphertext $(a_1 \dots a_l, b_1 \dots b_l)$ contains either 0 or 1, that is, at most one candidate receives one vote.
-

3 Attacking ballot secrecy

Ballot secrecy means a voter's vote is not revealed to anyone. We show that the Helios protocol does not satisfy this definition of ballot secrecy, by presenting an attack which allows an adversary to reveal a voter's vote. Moreover, we will show that formal definitions of ballot secrecy [KR05, DKR09, BHM08] are also violated.

Intuitively, an adversary may identify a voter's ballot on the bulletin board (using the voter's *id*) and recast this ballot by corrupting dishonest voters. The multiple occurrences of the voter's ballot will leak information in the tally and the adversary can exploit this knowledge to violate the voter's privacy. An informal description of the attack will now be presented in the case of three eligible voters and Section 3.3 considers a more realistic setting. (A formal analysis appears in Section 4.)

3.1 Attack description

Let us consider an election with candidates t_1, \dots, t_l and three eligible voters who have identities id_1, id_2 and id_3 . Suppose that voters id_1, id_2 are honest and id_3 is a dishonest voter controlled by the adversary. Further assume that the honest voters have cast their ballots. The bulletin board entries are as follows:

$$\begin{aligned} id_1, ciph_1, spk_1, spk'_1 \\ id_2, ciph_2, spk_2, spk'_2 \end{aligned}$$

where for $i \in \{1, 2\}$ we have

$$\begin{aligned} ciph_i &= (a_{i,1}, b_{i,1}), \dots, (a_{i,l}, b_{i,l}) \\ spk_i &= (\bar{a}_{i,1}, \bar{b}_{i,1}, \bar{c}_{i,1}, \bar{s}_{i,1}, \bar{a}'_{i,1}, \bar{b}'_{i,1}, \bar{c}'_{i,1}, \bar{s}'_{i,1}), \\ &\quad \dots, (\bar{a}_{i,l}, \bar{b}_{i,l}, \bar{c}_{i,l}, \bar{s}_{i,l}, \bar{a}'_{i,l}, \bar{b}'_{i,l}, \bar{c}'_{i,l}, \bar{s}'_{i,l}) \\ spk'_i &= (\bar{a}_i, \bar{b}_i, \bar{c}_i, \bar{s}_i, \bar{a}'_i, \bar{b}'_i, \bar{c}'_i, \bar{s}'_i) \end{aligned}$$

The value $ciph_i$ is the i th voter's encrypted vote, spk_i demonstrates that ciphertexts $(a_{i,1}, b_{i,1}), \dots, (a_{i,l}, b_{i,l})$ contain either 0 or 1, and spk'_i demonstrates that $(a_{i,1} \cdot \dots \cdot a_{i,l}, b_{i,1} \cdot \dots \cdot b_{i,l})$ contains either 0 or 1.

Exploiting the absence of ballot independence. The adversary observes the bulletin board and selects $ciph_k, spk_k, spk'_k$ such that id_k is the voter whose privacy will be compromised, where $k \in \{1, 2\}$. The adversary submits the ballot $ciph_k, spk_k, spk'_k$ and it immediately follows that the bulletin board is composed as follows:

$$\begin{aligned} id_1, ciph_1, spk_1, spk'_1 \\ id_2, ciph_2, spk_2, spk'_2 \\ id_3, ciph_k, spk_k, spk'_k \end{aligned}$$

It is trivial to see that each bulletin board entry represents a permitted vote; that is, $spk_1, spk'_1, spk_2, spk'_2, spk_k, spk'_k$ all contain valid signatures of knowledge.

We have informally shown that Helios does not satisfy ballot independence (observing another voter's interaction with the election system allows a voter to cast the *same* vote), and this will now be exploited to violate privacy.

Violating privacy. The homomorphic addition of ballots reveals the encrypted tally $(a_{1,1} \cdot a_{2,1} \cdot a_{k,1}, b_{1,1} \cdot b_{2,1} \cdot b_{k,1}), \dots, (a_{1,l} \cdot a_{2,l} \cdot a_{k,l}, b_{1,l} \cdot b_{2,l} \cdot b_{k,l})$ and, given the partial decryptions, these ciphertexts can be decrypted to reveal the number of votes for each candidate. Since there will be at least two votes for the candidate voter id_k voted for, the voter's vote can be revealed and hence privacy is not preserved. Moreover, the vote of the remaining honest voter will also be revealed.

A video demonstrating the attack against the Helios 3.0 implementation has been produced [SC10].

3.2 Variants exploiting malleability and key reuse

In the aforementioned attack description, the ballots cast by two voters are identical which may result in the detection of an attack. For a covert attack, the adversary may prefer to cast a distinct ballot. This can be achieved by exploiting the malleability of ballots. In particular, given a valid ballot

$$(a_1, b_1), \dots, (a_l, b_l), (\bar{a}_1, \bar{b}_1, \bar{c}_1, \bar{s}_1, \bar{a}'_1, \bar{b}'_1, \bar{c}'_1, \bar{s}'_1), \dots, \\ (\bar{a}_l, \bar{b}_l, \bar{c}_l, \bar{s}_l, \bar{a}'_l, \bar{b}'_l, \bar{c}'_l, \bar{s}'_l), (\bar{a}, \bar{b}, \bar{c}, \bar{s}, \bar{a}', \bar{b}', \bar{c}', \bar{s}') \quad (\text{B1})$$

the following ballots are also valid

$$(a_1, b_1), \dots, (a_l, b_l), (\bar{a}_1, \bar{b}_1, \bar{c}_1, \bar{s}_1 + q, \bar{a}'_1, \bar{b}'_1, \bar{c}'_1, \bar{s}'_1 + q), \dots, \\ (\bar{a}_l, \bar{b}_l, \bar{c}_l, \bar{s}_l + q, \bar{a}'_l, \bar{b}'_l, \bar{c}'_l, \bar{s}'_l + q), (\bar{a}, \bar{b}, \bar{c}, \bar{s} + q, \bar{a}', \bar{b}', \bar{c}', \bar{s}' + q) \quad (\text{B2})$$

$$(a_{\pi(1)}, b_{\pi(1)}), \dots, (a_{\pi(l)}, b_{\pi(l)}), \\ (\bar{a}_{\pi(1)}, \bar{b}_{\pi(1)}, \bar{c}_{\pi(1)}, \bar{s}_{\pi(1)}, \bar{a}'_{\pi(1)}, \bar{b}'_{\pi(1)}, \bar{c}'_{\pi(1)}, \bar{s}'_{\pi(1)}), \dots, \\ (\bar{a}_{\pi(l)}, \bar{b}_{\pi(l)}, \bar{c}_{\pi(l)}, \bar{s}_{\pi(l)}, \bar{a}'_{\pi(l)}, \bar{b}'_{\pi(l)}, \bar{c}'_{\pi(l)}, \bar{s}'_{\pi(l)}), \\ (\bar{a}, \bar{b}, \bar{c}, \bar{s}, \bar{a}', \bar{b}', \bar{c}', \bar{s}') \quad (\text{B3})$$

where π is an arbitrary permutation over $\{1, \dots, l\}$. Ballot B2 adds q to the response components of Ballot B1, this changes the ballot but not the vote. (It is also possible to modify a subset of the response components.) However, this might be considered an implementation bug in Helios 3.0, rather than a theoretical attack, because the ballots are identical if considered as group elements. Replaying Ballot B3 has the advantage of casting a theoretical distinct ballot, since Ballot B3 represents a vote for a different candidate (with the exception of an abstention vote), and it is possible to compromise ballot secrecy in elections with three voters without abstention votes; however, more than one (modified) ballot may be required in elections with abstention votes. This variant of our attack also demonstrates a further violation of ballot independence in Helios: observing another voter's interaction with the election system allows a voter to cast a *different* vote (for example, a voter can cast a distinct vote from their boss). Both variants of the attack are particularly useful when the bulletin board includes the hash of the ballot (for example, in the Helios 3.0 implementation), rather than the complete ballot, because the hashes will be distinct.

An adversary may replay ballots in different elections, when the trustees' public key is reused and the candidate lists for each election are of equal length. This variant of the attack can be avoided if distinct keys are used for each election.

The variants of our attack in this section have all been successfully launched against the Helios 3.0 implementation.

3.3 Generalised attack and French election case study

Our attack demonstrates that the ballot of an arbitrary voter can be replayed by any other voter. In general, this does not reveal the voter’s vote; but, some information is leaked, and colluding voters can replay sufficiently many ballots to leak the voter’s vote. We will now discuss the feasibility of compromising ballot secrecy in a real-world election, focusing on the cost of an attack in French legislative elections, where each district elects a representative for the French National Assembly. Districts have several polling stations and each polling station individually announces its tally [Fr]; these tallies are published in local newspapers. The publication of tallies is typical of French elections at all levels; for example, from the election of mayor, to the presidential election.

In this (standard) voting configuration, an adversary can violate the ballot secrecy of a given voter by corrupting voters registered at the same polling station (for example, a coalition of neighbours or a family). The corrupted voters replay the ballot of the voter under attack, as previously explained. The motivation for restricting the selection of corrupted voters to the same polling station is twofold. Firstly, fewer corrupt voters are required to significantly influence the tally of an individual polling station (in comparison to influencing the election outcome). Secondly, it is unlikely to change the district’s elected representative, because a candidate will receive only a few additional votes in the district; it follows that coercing voters to sacrifice their vote, for the purposes of the attack, should be easier. In the remainder of this section, we discuss how many corrupt voters are required to violate ballot secrecy – by making a significant change in the tally of a polling station – in an arbitrary district of Aulnay-sous-Bois and a rural district in Toul.

3.3.1 Ballot secrecy in Aulnay-sous-Bois

Using historic data and/or polls, it is possible to construct the expected distribution of votes. For simplicity, let us assume the distribution of votes per polling station is the average of the 2010 tally (Table 1), and that if the adversary can increase the number of votes for a particular candidate by more than σ (by replaying a voter’s ballot), then this is sufficient to determine that the voter voted for that candidate. In addition, suppose that the adversary corrupts abstaining voters and therefore we do not consider the redistribution of votes. We remark that corrupting abstaining voters may be a fruitful strategy, since abstaining voters do not sacrifice their vote by participating in an attack.

Table 2 presents the expected distribution of votes, and includes the number of voters that an adversary must corrupt to determine if a voter voted for a particular candidate, for various values of σ . We shall further assume that participation in the region is consistent with 2010; that is, 291 of the 832 eligible voters are expected to participate. It follows that 50 voters corresponds to approximately 6% of the Aulnay-sous-Bois electorate, and 10 voters corresponds to approximately 1%. Our results therefore demonstrate that the privacy of a voter can be compromised by corrupting a small number of voters. In particular,

Party	Tally
PS	4120
UMP	3463
FN	1933
Europe Eco.	1921
Front de gauche	880
NPA	697
MODEM	456
Debout la République	431
Alliance école	193
LO	156
Émergence	113
Liste chrétienne	113

Table 1: 2010 legislative election results in Aulnay-sous-Bois [Fr10]

for medium-size parties (in terms of votes received) – including, for example, FN and Europe Ecologie – it is sufficient to corrupt 19 voters to see the number of votes increase by 50%. Furthermore, given the low turn-out (541 voters are expected to abstain), it seems feasible to corrupt abstaining voters, and therefore an attack can be launched without any voter sacrificing their vote.

Limitations. For such an attack based on a statistical model, we acknowledge that this model is rather naïve, but believe it is sufficiently indicative to illustrate the real threat of an attack against privacy. A definitive mathematical analysis should be considered in the future.

Cases of complete privacy breach. The probabilistic nature of these attacks may introduce sufficient uncertainty to prevent privacy violations, and we will consider voting configurations where an adversary can definitively learn a voter’s vote. Observe that if an attacker can corrupt half of the voters at a polling station, then the vote of an arbitrary voter can be revealed. Moreover, the cost of this attack can be reduced. In particular, if n dishonest voter’s replay voter \mathcal{V} ’s ballot, then it is possible to deduce that \mathcal{V} did not vote for any candidate that received strictly less than $n + 1$ votes. This leaks information about voter \mathcal{V} ’s chosen candidate and in cases where exactly one candidate received more than n votes, the voter’s vote can be deduced.

3.3.2 Ballot secrecy in small polling stations

The difficulties of large scale corruption may prohibit our attack in the majority of polling stations; however, our attack is feasible in small polling stations found in rural districts. For example, let us consider the 2007 legislative elections in the district of Toul [Est07]. This district has 75350 eligible voters registered at

Party	Expected tally	$\sigma = 200\%$	$\sigma = 150\%$	$\sigma = 50\%$	$\sigma = 20\%$
PS	81	162	122	41	17
UMP	68	136	102	34	14
FN	38	76	57	19	8
Europe Eco.	38	76	57	19	8
Front de gauche	17	34	26	9	4
NPA	14	28	21	7	3
MODEM	9	18	14	5	2
Debout la République	8	16	12	4	2
Alliance école	4	8	6	2	1
LO	3	6	5	2	1
Émergence	2	4	3	1	1
Liste chrétienne	2	4	3	1	1

Table 2: Number of duplicate ballots for a significant change in the tally

193 polling stations. Accordingly, the average polling station has 390 registered voters, but the variance is large. Indeed, 33 polling stations have between 50 and 99 voters, 9 polling stations have less than 50 voters, and the smallest two polling stations have 8, respectively 16, voters. Moreover, the attack is simplified by non-participating voters. In these small polling stations it is thus sufficient to corrupt a very small number of voters to reveal a voter’s vote while the final outcome of the election would not change as it is based on 75350 eligible voters.

4 Solution: Weeding replayed ballots

Our attack exploits the possibility of replaying a voter’s ballot without detection, and can be attributed to the lack of ballot independence in Helios. This section sketches some possible solutions to ensure ballot independence.

4.1 Weeding replayed ballots

The ballots replayed in our attacks can all be identified. First, ciphertexts and signatures of knowledge should have a unique representation as group elements, for example, by requiring that the response component of signatures of knowledge is in the interval $[0, q - 1]$. Second, a ballot should not contain a ciphertext that already exists on the bulletin board. The election officer should reject ballots that do not satisfy these conditions. This solution is simple and can easily be implemented in a future version of Helios.

4.2 Binding ballots to voters

The previous approach requires a special mechanism to handle replayed ballots. We now propose a technique that makes such actions futile. In essence, based upon inspiration from [Gen95, §4.2] and [CGS97], we ensure that proofs associated with replayed ballots are considered invalid; that is, we bind the link between a voter and her ballot. This is achieved by adding the identity of the voter in the construction of challenges used by signatures of knowledge. More precisely, for voter id , the sign algorithm (defined in Appendix A) is modified as follows: on input (a, b) , such that $a \equiv g^r \pmod{p}$ and $b \equiv h^r \cdot g^m \pmod{p}$, let challenge $c_m = \mathcal{H}(a_{\min}, b_{\min}, \dots, a_{\max}, b_{\max}, id) - \sum_{i \in \{\min, \dots, m-1, m+1, \dots, \max\}} c_i \pmod{q}$, where values $a_{\min}, b_{\min}, \dots, a_{\max}, b_{\max}$ and $c_1, \dots, c_{m-1}, c_{m+1}, \dots, c_m$ are defined as before. For correctness, the verification algorithm must also be modified. In particular, for candidate signatures constructed by voter id , the verifier should check $\mathcal{H}(a_{\min}, b_{\min}, \dots, a_{\max}, b_{\max}, id) \equiv \sum_{\min \leq i \leq \max} c_i \pmod{q}$.

In a similar direction, the electronic voting protocol proposed by Juels, Catalano & Jakobsson [JCJ05] – which has been implemented by Clarkson, Chong & Myers [CCM08, CCM07] as Civitas – requires ballots to be bound to private voter credentials. This provides *eligibility verifiability* [KRS10]: anyone can check that each ballot published on the bulletin board was cast by a registered voter and at most one ballot is tallied per voter. It is likely that eligibility

verifiability enforces ballot independence, but the provision of eligibility verifiability appears to be expensive, in particular, Juels, Catalano & Jakobsson and Clarkson, Chong & Myers assume the existence of an infrastructure for voter credentials.

4.3 Discussion

Our *weeding replayed ballots* solution is particularly attractive because it adheres to Benaloh’s notion of *ballot casting assurance* [Ben06, Ben07] which asserts that the ballot encryption device (the browser script in this instance) does not know the voter’s identity. (We remark that neither the original Helios scheme nor our proposed fix strictly satisfy Benaloh’s notion of ballot casting assurance if a voter decides to use her own computer.) The ballot casting assurance principle is important because knowledge of the voter’s identity could be used to infer the likelihood of auditing and this information can be used to influence the behaviour of the ballot encryption device; in particular, if a ballot is unlikely to be audited, then the device may act maliciously, for example, by encrypting a different vote. By comparison, the *binding ballots to voters* solution would necessarily require that the voter’s identity is revealed to the ballot encryption device. Moreover, for privacy purposes, the election officer may choose to allocate voters with pseudo-identities when casting ballots (rather than associate ballots with actual voter identities); since these pseudo-identities are unknown to voters in advance, an additional interaction with the election officer would be required. (Note that the use of pseudo-identities does not prevent the attack by breaking the link between ballots and voters, because the link is known by the election officer.) Finally, extending Helios to provide eligibility verifiability would require a considerable extension to the Helios code-base and, furthermore, finding a suitable solution is an open problem. Accordingly, we adopt the weeding replayed ballot solution and, in the next section, we show that this is sufficient to ensure ballot secrecy, in the formal setting.

5 Formal proof of ballot secrecy

We formally prove that weeding duplicate ballots ensures ballot secrecy. We make use of the applied pi calculus [AF01, RS11], due to its proven suitability for evaluating security properties of electronic voting protocols (see, for example, [DKR09, BHM08, KRS10]).

5.1 Applied pi calculus

We first recall the applied pi calculus setting [AF01]. We assume an infinite set of *names* $a, b, c, \dots, k, \dots, m, n, \dots, s, \dots$, an infinite set of *variables* x, y, z, \dots , and a *signature* Σ consisting of a finite set of *function symbols*, each with an associated arity. We use metavariables u, w to range over both names and variables. *Terms* L, M, N, T, U, V are built by applying function symbols to

names, variables, and other terms. We write $\{M/x\}$ for the *substitution* that replaces the variable x with the term M . Arbitrarily large substitutions can be written as $\{M_1/x_1, \dots, M_l/x_l\}$ and the letters σ and τ range over substitutions. We write $N\sigma$ for the result of applying σ to the free variables of term N . A term is *ground* when it does not contain variables.

The signature Σ is equipped with an *equational theory* E , that is, a set of equations of the form $M = N$, where the terms M, N are defined over the signature Σ . We define equality modulo the equational theory, written $=_E$, as the smallest equivalence relation on terms that contains E and is closed under application of function symbols, substitution of terms for variables and bijective renaming of names. We write $M =_E N$ when the equation $M = N$ is in the theory E , and keep the signature implicit. When E is clear from its usage, we may abbreviate $M =_E N$ as $M = N$. The negation of $M =_E N$ is denoted $M \neq_E N$ (and similarly abbreviated $M \neq N$).

Processes and *extended processes* are defined in the usual way (Figure 2). We write $\nu \tilde{u}$ for the (possibly empty) series of pairwise-distinct binders $\nu u_1. \dots \nu u_l$. The active substitution $\{M/x\}$ can replace the variable x for the term M in every process it comes into contact with and this behaviour can be controlled by restriction, in particular, the process $\nu x.(\{M/x\} \mid P)$ corresponds exactly to let $x = M$ in P . Arbitrarily large active substitutions can be obtained by parallel composition and we occasionally abbreviate $\{M_1/x_1\} \mid \dots \mid \{M_l/x_l\}$ as $\{M_1/x_1, \dots, M_l/x_l\}$ or $\{\tilde{M}/\tilde{x}\}$. We also use σ and τ to range over active substitutions, and write $N\sigma$ for the result of applying σ to the free variables of N . Extended processes must have at most one active substitution for each variable and there is exactly one when the variable is under restriction. The only minor change compared to [AF01] is that conditional branches now depend on formulae $\phi, \psi ::= M = N \mid M \neq N \mid \phi \wedge \psi$. If M and N are ground, we define $\llbracket M = N \rrbracket$ to be **true** if $M =_E N$ and **false** otherwise. The semantics of $\llbracket \cdot \rrbracket$ is then extended to formulae in the standard way.

The *scope* of names and variables are delimited by binders $u(x)$ and νu . The set of bound names is written $\text{bn}(A)$ and the set of bound variables is written $\text{bv}(A)$; similarly we define the set of free names $\text{fn}(A)$ and free variables $\text{fv}(A)$. Occasionally, we write $\text{fn}(M)$ (and $\text{fv}(M)$ respectively) for the set of names (and respectively variables) which appear in term M . An extended process is *closed* when every variable x is either bound or defined by an active substitution.

We define a *context* $C[\cdot]$ to be an extended process with a hole. We obtain $C[A]$ as the result of filling $C[\cdot]$'s hole with the extended process A . An *evaluation context* is a context whose hole is not in the scope of a replication, a conditional, an input, or an output. A context $C[\cdot]$ closes A when $C[A]$ is closed.

A *frame*, denoted φ or ψ , is an extended process built from the null process 0 and active substitutions $\{M/x\}$, which are composed by parallel composition and restriction. The *domain* $\text{dom}(\varphi)$ of a frame φ is the set of variables that φ exports, that is, the set of variables x for which φ contains an active substitution $\{M/x\}$ such that x is not under restriction. Every extended process A can be

Figure 2 Syntax for processes

$P, Q, R ::=$	(plain) processes
0	null process
$P \mid Q$	parallel composition
$!P$	replication
$\nu n.P$	name restriction
if ϕ then P else Q	conditional
$u(x).P$	message input
$\bar{u}\langle M \rangle.P$	message output
$A, B, C ::=$	extended processes
P	plain process
$A \mid B$	parallel composition
$\nu n.A$	name restriction
$\nu x.A$	variable restriction
$\{M/x\}$	active substitution

mapped to a frame $\varphi(A)$ by replacing every plain process in A with 0.

5.1.1 Operational semantics

The operational semantics are defined by three relations: *structural equivalence* (\equiv), *internal reduction* (\rightarrow), and *labelled reduction* ($\xrightarrow{\alpha}$). These relations satisfy the rules in Figure 3 and are defined such that: structural equivalence is the smallest equivalence relation on extended processes that is closed by α -conversion of both bound names and bound variables, and closed under application of evaluation contexts; internal reduction is the smallest relation on extended processes closed under structural equivalence and application of evaluation contexts; and for labelled reductions α is a *label* of the form $c(M)$, $\bar{c}\langle u \rangle$, or $\nu u.\bar{c}\langle u \rangle$ such that u is either a channel name or a variable of base type.

5.1.2 Equivalence

The definition of observational equivalence [AF01] quantifies over all contexts which makes proofs difficult, therefore we adopt labelled bisimilarity in this article. Labelled bisimilarity relies on an equivalence relation between frames, called static equivalence.

Definition 1 (Static equivalence). *Two closed frames φ and ψ are statically equivalent, denoted $\varphi \approx_s \psi$, if $\text{dom}(\varphi) = \text{dom}(\psi)$ and there exists a set of names \tilde{n} and substitutions σ, τ such that $\varphi \equiv \nu \tilde{n}.\sigma$ and $\psi \equiv \nu \tilde{n}.\tau$ and for all terms M, N such that $\tilde{n} \cap (\text{fn}(M) \cup \text{fn}(N)) = \emptyset$, we have $M\sigma =_E N\sigma$ holds if and only if $M\tau =_E N\tau$ holds. Two closed extended processes A, B are statically equivalent, written $A \approx_s B$, if their frames are statically equivalent; that is, $\varphi(A) \approx_s \varphi(B)$.*

Figure 3 Semantics for processes

PAR-0	$A \equiv A \mid 0$
PAR-A	$A \mid (B \mid C) \equiv (A \mid B) \mid C$
PAR-C	$A \mid B \equiv B \mid A$
REPL	$!P \equiv P \mid !P$
NEW-0	$\nu n.0 \equiv 0$
NEW-C	$\nu u.\nu w.A \equiv \nu w.\nu u.A$
NEW-PAR	$A \mid \nu u.B \equiv \nu u.(A \mid B)$ where $u \notin \text{fv}(A) \cup \text{fn}(A)$
ALIAS	$\nu x.\{M/x\} \equiv 0$
SUBST	$\{M/x\} \mid A \equiv \{M/x\} \mid A\{M/x\}$
REWRITE	$\{M/x\} \equiv \{N/x\}$ where $M =_E N$
COMM	$\bar{c}\langle x \rangle.P \mid c(x).Q \rightarrow P \mid Q$
THEN	if ϕ then P else $Q \rightarrow P$ if $\llbracket \phi \rrbracket = \text{true}$
ELSE	if ϕ then P else $Q \rightarrow Q$ otherwise
IN	$c(x).P \xrightarrow{c(M)} P\{M/x\}$
OUT-ATOM	$\bar{c}\langle u \rangle.P \xrightarrow{\bar{c}\langle u \rangle} P$
OPEN-ATOM	$\frac{A \xrightarrow{\bar{c}\langle u \rangle} A' \quad u \neq c}{\nu u.A \xrightarrow{\nu u.\bar{c}\langle u \rangle} A'}$
SCOPE	$\frac{A \xrightarrow{\alpha} A' \quad u \text{ does not occur in } \alpha}{\nu u.A \xrightarrow{\alpha} \nu u.A'}$
PAR	$\frac{A \xrightarrow{\alpha} A' \quad \text{bv}(\alpha) \cap \text{fv}(B) = \text{bn}(\alpha) \cap \text{fn}(B) = \emptyset}{A \mid B \xrightarrow{\alpha} A' \mid B}$
STRUCT	$\frac{A \equiv B \quad B \xrightarrow{\alpha} B' \quad B' \equiv A'}{A \xrightarrow{\alpha} A'}$

The relation \approx_s is called *static* equivalence because it only examines the current state of the processes, and not the processes' dynamic behaviour. The following definition of labelled bisimilarity captures the dynamic part.

Definition 2 (Labelled bisimilarity). *Labelled bisimilarity (\approx_l) is the largest symmetric relation \mathcal{R} on closed extended processes such that $A \mathcal{R} B$ implies:*

1. $A \approx_s B$;
2. if $A \rightarrow A'$, then $B \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' ;
3. if $A \xrightarrow{\alpha} A'$ such that $\text{fv}(\alpha) \subseteq \text{dom}(A)$ and $\text{bn}(\alpha) \cap \text{fn}(B) = \emptyset$, then $B \rightarrow^* \xrightarrow{\alpha} \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' .

5.2 Modelling Helios in applied pi

We start by constructing a suitable signature Σ to capture the cryptographic primitives used by Helios and define an equational theory E to capture the relationship between these primitives.

5.2.1 Signature

We adopt the following signature.

$$\Sigma = \{\text{ok}, \text{zero}, \text{one}, \perp, \text{fst}, \text{snd}, \text{pair}, *, +, \circ, \text{partial}, \text{checkspk}, \text{penc}, \text{spk}, \}$$

Functions ok , zero , one , \perp are constants; fst , snd are unary functions; dec , pair , partial , $*$, $+$, \circ are binary functions; checkspk , penc are ternary functions; and spk is a function of arity four. We adopt infix notation for $*$, $+$, and \circ .

The term $\text{penc}(T, N, M)$ denotes the encryption of plaintext M , using random nonce N and key T . The term $U * U'$ denotes the homomorphic combination of ciphertexts U and U' , the corresponding operation on plaintexts is written $M + M'$ and $N \circ N'$ on nonces. The partial decryption of ciphertext U using key L is denoted $\text{partial}(L, U)$. The term $\text{spk}(T, N, M, U)$ represents a signature of knowledge that proves U is a ciphertext under the public key T on the plaintext M using nonce N and such that M is either the constant zero or one . We introduce tuples using pairings and, for convenience, $\text{pair}(M_1, \text{pair}(\dots, \text{pair}(M_n, \perp)))$ is occasionally abbreviated as (M_1, \dots, M_n) , and $\text{fst}(\text{snd}^{i-1}(M))$ is denoted $\pi_i(M)$, where $i \in \mathbb{N}$. We use the equational theory E that asserts functions $+$, $*$, \circ are commutative and associative, and includes the equations:

$$\text{fst}(\text{pair}(x, y)) = x \tag{E1}$$

$$\text{snd}(\text{pair}(x, y)) = y \tag{E2}$$

$$\text{zero} + \text{one} = \text{one} \tag{E3}$$

$$\text{dec}(x_{\text{sk}}, \text{penc}(\text{pk}(x_{\text{sk}}), x_{\text{rand}}, x_{\text{plain}})) = x_{\text{plain}} \tag{E4}$$

$$\text{dec}(\text{partial}(x_{\text{sk}}, \text{ciph}), \text{ciph}) = x_{\text{plain}} \tag{E5}$$

$$\text{where } \text{ciph} = \text{penc}(\text{pk}(x_{\text{sk}}), x_{\text{rand}}, x_{\text{plain}})$$

$$\begin{aligned} \text{penc}(x_{\text{pk}}, y_{\text{rand}}, y_{\text{plain}}) * \text{penc}(x_{\text{pk}}, z_{\text{rand}}, z_{\text{plain}}) \\ = \text{penc}(x_{\text{pk}}, y_{\text{rand}} \circ z_{\text{rand}}, y_{\text{plain}} + z_{\text{plain}}) \end{aligned} \tag{E6}$$

$$\text{checkspk}(x_{\text{pk}}, \text{ball}, \text{spk}(x_{\text{pk}}, x_{\text{rand}}, \text{zero}, \text{ball})) = \text{ok} \tag{E7}$$

$$\text{where } \text{ball} = \text{penc}(x_{\text{pk}}, x_{\text{rand}}, \text{zero})$$

$$\begin{aligned} \text{checkspk}(x_{\text{pk}}, \text{ball}, \text{spk}(x_{\text{pk}}, x_{\text{rand}}, \text{one}, \text{ball})) &= \text{ok} & (\text{E8}) \\ \text{where } \text{ball} &= \text{penc}(x_{\text{pk}}, x_{\text{rand}}, \text{one}) \end{aligned}$$

Equation E5 allows plaintext M to be recovered from ciphertext $\text{penc}(\text{pk}(L), N, M)$ given partial decryption $\text{partial}(L, \text{penc}(\text{pk}(L), N, M))$, when the partial decryption is constructed using the private key L . Equation E6 represents the homomorphic combination of ciphertexts. The Equations E7 and E8 allow the verification of signatures of knowledge $\text{spk}(T, N, M, \text{penc}(T, N, M))$, when $M \in \{\text{zero}, \text{one}\}$. The remaining equations are standard.

Example 1. *Given randomness N, N' , plaintexts $(M, M') \in \{(\text{zero}, \text{zero}), (\text{zero}, \text{one}), (\text{one}, \text{zero})\}$, and public key T , one can construct a signature of knowledge $L = \text{spk}(T, N \circ N', M + M', \text{penc}(T, N, M) * \text{penc}(T, N', M'))$. Then checkspk applied to the public key T , the homomorphically combined ciphertexts $\text{penc}(T, N, M) * \text{penc}(T, N', M')$, and the signature L is equal to ok using Equations E3, E6, E7, and E8*

5.2.2 Helios process specification

In the applied pi calculus, it is sufficient to model the parts of the voting system which need to be trusted for ballot secrecy; all the remaining parts of the system are controlled by the adversarial environment. Accordingly, we assume the existence of at least two honest voters \mathcal{A}, \mathcal{B} ; since this avoids the scenario where ballot secrecy of an individual voter is compromised by collusion amongst all the remaining voters. In addition, the following trust assumptions are required.

- At least one trustee is honest
- The election officer runs the bulletin board honestly:
 - Voters \mathcal{A}, \mathcal{B} have authentic channels with the bulletin board
 - Signatures of knowledge are checked for dishonest voters*
 - Replays of honest ballots (that is, those cast by \mathcal{A} or \mathcal{B}) are rejected*
 - The tally is correctly computed*
 - The trustees have an authentic channel with the bulletin board
- The browser script is trusted and has the correct public key of the election

(Assumptions marked with * could be performed by an honest trustee, rather than the bulletin board.) Although neither voters nor observers can verify that there exists an honest trustee, an assurance of trust is provided by distribution. The necessity to trust the election officer to run the bulletin board is less desirable and work-in-progress [PAM10] aims to weaken this assumption; moreover, to further distribute trust assumptions, the trustees could also check signatures and tallying. Finally, trust in the browser script can be obtained by using software written by a reputable source or writing your own code.

Figure 4 Helios process specification

Given the number of voters $n \geq 2$ and Helios process specification ϕ , the administration process A_n^ϕ and voting process V are defined below.

$$\begin{aligned}
V &= \nu r . \\
&\quad \text{let } ciph = \text{penc}(z_{\text{pk}}, r, x_{\text{vote}}) \text{ in} \\
&\quad \text{let } spk = \text{spk}(z_{\text{pk}}, r, x_{\text{vote}}, ciph) \text{ in} \\
&\quad \nu r' . \\
&\quad \text{let } ciph' = \text{penc}(z_{\text{pk}}, r', x'_{\text{vote}}) \text{ in} \\
&\quad \text{let } spk' = \text{spk}(z_{\text{pk}}, r', x'_{\text{vote}}, ciph') \text{ in} \\
&\quad \text{let } \widehat{spk} = \text{spk}(z_{\text{pk}}, r \circ r', x_{\text{vote}} + x'_{\text{vote}}, ciph * ciph') \text{ in} \\
&\quad \overline{x_{\text{auth}}} \langle (ciph, ciph', spk, spk', \widehat{spk}) \rangle \\
\\
A_n^\phi &= \nu sk_T, a_1, a_2, d . (- \mid BB_n^\phi \mid T \mid \{\text{pk}(sk_T)/z_{\text{pk}}\}) \\
\\
BB_n^\phi &= a_1(y_1) . \bar{c}\langle y_1 \rangle . a_2(y_2) . \bar{c}\langle y_2 \rangle . \\
&\quad a_3(y_3) . \text{if } \phi\{y_3/y_{\text{ballot}}\} \text{ then} \\
&\quad \dots a_n(y_n) . \text{if } \phi\{y_n/y_{\text{ballot}}\} \text{ then} \\
&\quad \text{let } tally = \pi_1(y_1) * \dots * \pi_1(y_n) \text{ in} \\
&\quad \text{let } tally' = \pi_2(y_1) * \dots * \pi_2(y_n) \text{ in} \\
&\quad \bar{d}\langle (tally, tally') \rangle . \\
&\quad d(y_{\text{partial}}) . \\
&\quad \bar{c}\langle y_{\text{partial}} \rangle . \\
&\quad \bar{c}\langle (\text{dec}(\pi_1(y_{\text{partial}}), tally), \text{dec}(\pi_2(y_{\text{partial}}), tally')) \rangle \\
\\
T &= d(y_{\text{tally}}) . \\
&\quad \bar{d}\langle (\text{partial}(sk_T, \pi_1(y_{\text{tally}})), \text{partial}(sk_T, \pi_2(y_{\text{tally}}))) \rangle
\end{aligned}$$

In an election with two candidates, the trusted components are modelled by the administration process A_n^ϕ and voting process V defined in Figure 4. For generality, the administration process A_n^ϕ is parametrised by the number of voters n and a formula ϕ ; the latter corresponds to the checks performed by the bulletin board before accepting a ballot. We will consider several variants of Helios (including the original Helios 2.0 protocol and our fixed scheme) by considering suitable formula that we call *Helios process specifications*.

Definition 3 (Helios process specification). *A formula ϕ is a Helios process specification, if $\text{fv}(\phi) \subseteq \{y_1, y_2, y_{\text{ballot}}, z_{\text{pk}}\}$.*

The voting process V contains free variables $x_{\text{vote}}, x'_{\text{vote}}$ to represent the voter's vote (which is expected to be encoded as constants **zero** and **one**), and the free variable x_{auth} represents the channel shared by the voter and the bulletin board. The definition of the process V corresponds to the description of the browser script (Figure 1). The administration process A_n^ϕ is parametrised

by the number of voters n and Helios process specification ϕ . The restricted name sk_T models the tallier's secret key and the public part $\text{pk}(sk_T)$ is included in the process's frame. The restricted names a_1, a_2 model authentic channels between the two honest voters and the bulletin board, and the channel name d captures the authentic channel with the honest trustee. To ensure the adversary has access to messages sent on private channels, communication is relayed on the public channel c . The sub-process BB_n^ϕ represents the bulletin board and T represents the tallier. The bulletin board accepts ballots from each voter and checks they are valid using the Helios process specification ϕ (this predicate will be discussed in more detail below). Once all ballots have been submitted, the bulletin board homomorphically combines the ciphertexts and sends the encrypted tallies to the tallier for decryption. (The necessity for all voters to participate is included for simplicity, in particular, our bulletin board does not weed ballots containing invalid proofs.) The tallier receives the homomorphic combinations of ballots y_{tally} and derives a partial decryption for each candidate; these partial decryptions are sent to the bulletin board and the election result is published.

Example 2. *Given a Helios process specification ϕ , an election with voters \mathcal{A} and \mathcal{B} who select votes $(m_1, m'_1), (m_2, m'_2) \in \{(\text{zero}, \text{zero}), (\text{zero}, \text{one}), (\text{one}, \text{zero})\}$ and such that the other $n - 2$ voters are controlled by the adversary, can be modelled by the process $A_n^\phi[V\{a_1/x_{\text{auth}}\}\sigma \mid V\{a_2/x_{\text{auth}}\}\tau]$, where $\sigma = \{m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\}$ and $\tau = \{m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\}$.*

Ballot validity. In Helios 2.0, the election officer considers a ballot to be valid if the signature proofs of knowledge hold. Accordingly, we can model the Helios administration by the process $A_n^{\phi_{\text{orig}}}$ where the Helios process specification ϕ_{orig} is defined as follows.

$$\begin{aligned} \phi_{\text{orig}} \triangleq & \text{checkspk}(z_{\text{pk}}, \pi_1(y_{\text{ballot}}), \pi_3(y_{\text{ballot}})) = \text{ok} \\ & \wedge \text{checkspk}(z_{\text{pk}}, \pi_2(y_{\text{ballot}}), \pi_4(y_{\text{ballot}})) = \text{ok} \\ & \wedge \text{checkspk}(z_{\text{pk}}, \pi_1(y_{\text{ballot}}) * \pi_2(y_{\text{ballot}}), \pi_5(y_{\text{ballot}})) = \text{ok} \end{aligned}$$

We have shown that these checks are insufficient to ensure ballot secrecy. Our weeding replayed ballots solution proposed in Section 4.1 additionally requires that the ciphertexts inside the ballot do not appear on the bulletin board. This revised scheme can be modelled using the Helios process specification ϕ_{sol} is defined as follows.

$$\phi_{\text{sol}} \triangleq \phi_{\text{orig}} \wedge \pi_6(y_{\text{ballot}}) = \perp \wedge \bigwedge_{i,j \in \{1,2\}} \pi_i(y_j) \neq \pi_1(y_{\text{ballot}}) \wedge \pi_i(y_j) \neq \pi_2(y_{\text{ballot}})$$

We can also model a naïve solution that would consist in weeding only identical ballots by considering the Helios process specification ϕ_{ident} defined below.

$$\phi_{\text{ident}} \triangleq \phi_{\text{orig}} \wedge \pi_6(y_{\text{ballot}}) = \perp \wedge y_{\text{ballot}} \neq y_1 \wedge y_{\text{ballot}} \neq y_2$$

We have already shown that removing exact duplicates is insufficient because it would fail to detect variants of our attack whereby the contents of a ballot are permuted. In the next section, we formally show that Helios 2.0 (modelled using ϕ_{orig}) and the naïve solution (modelled using ϕ_{ident}) do not satisfy ballot secrecy, and that our proposed solution (modelled using ϕ_{sol}) does satisfy ballot secrecy.

5.3 Formal analysis: Ballot secrecy

Based upon [KR05, DKR06, DKR09], and as previously discussed (see *related work* in Section 1), we formalise ballot secrecy for two voters \mathcal{A} , \mathcal{B} and two candidates t , t' with the assertion that an adversary cannot distinguish between a situation in which voter \mathcal{A} votes for candidate t and voter \mathcal{B} votes for candidate t' , from another situation in which \mathcal{A} votes t' and \mathcal{B} votes t . Formally, this is captured by Definition 4.

Definition 4 (Ballot secrecy). *Given a Helios process specification ϕ , we say ballot secrecy is satisfied if for all $(m_1, m'_1), (m_2, m'_2) \in \{(\text{zero}, \text{zero}), (\text{zero}, \text{one}), (\text{one}, \text{zero})\}$ and integers $n \geq 2$, we have*

$$A_n^\phi[V\{a_1/x_{\text{auth}}\}\sigma \mid V\{a_2/x_{\text{auth}}\}\tau] \approx_l A_n^\phi[V\{a_1/x_{\text{auth}}\}\tau \mid V\{a_2/x_{\text{auth}}\}\sigma]$$

where $\sigma = \{m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\}$ and $\tau = \{m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\}$.

The ballot secrecy definition proposed by Delaune, Kremer & Ryan considered a vote to be an arbitrary name, whereas a vote in our setting must be a pair $(m, n) \in \{(\text{zero}, \text{zero}), (\text{zero}, \text{one}), (\text{one}, \text{zero})\}$; it follows that Definition 4 is a straightforward variant of the original.

The Helios 2.0 protocol does not satisfy our privacy definition (Lemma 1) and naïve ballot weeding solutions are also insufficient (Lemma 2).

Lemma 1. *The Helios process specification ϕ_{orig} does not satisfy ballot secrecy.*

Intuitively, the proof of Lemma 1 is due to the environment's ability to replay \mathcal{A} 's ballot, therefore introducing an observable difference: the result will include two instances of \mathcal{A} 's vote. Formally, this follows immediately from the proof Lemma 2.

Lemma 2. *The Helios process specification ϕ_{ident} does not satisfy ballot secrecy.*

Proof. Consider $n = 3$, $(m_1, m'_1) = (\text{zero}, \text{one})$ and $(m_2, m'_2) = (\text{one}, \text{zero})$. Let $\sigma = \{m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\}$ and $\tau = \{m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\}$. We consider a sequence of transitions where the two voters output their ballots and then the adversary chooses its ballots to be a permutation of the first voter's ballot. Namely, if the first voter's ballot is $(\text{ciph}, \text{ciph}', \text{spk}, \text{spk}', \widehat{\text{spk}})$ then the adversary outputs $(\text{ciph}', \text{ciph}, \text{spk}', \text{spk}, \widehat{\text{spk}})$. Formally, this corresponds to the transitions.

$$\begin{aligned} & A_n^\phi[V\{a_1/x_{\text{auth}}\}\sigma \mid V\{a_2/x_{\text{auth}}\}\tau] \\ & \xrightarrow{\nu x.\bar{c}\langle x \rangle} \xrightarrow{\nu y.\bar{c}\langle y \rangle} \xrightarrow{c(\langle \pi_2(x), \pi_1(x), \pi_4(x), \pi_3(x), \pi_5(x) \rangle)} \xrightarrow{*} \xrightarrow{\nu z.\bar{c}\langle z \rangle} \nu \tilde{n}.\tau_1 \end{aligned}$$

for some names \tilde{n} and substitution τ_1 such that:

$$\text{dec}(\pi_1(z), \pi_1(x) * \pi_1(y) * \pi_2(x))\tau_1 =_E \text{one} + \text{one}$$

Then this labeled transition has to be matched:

$$A_n^\phi[V\{a_1/x_{\text{auth}}\}\tau \mid V\{a_2/x_{\text{auth}}\}\sigma] \xrightarrow{\nu x.\bar{c}\langle x \rangle} \xrightarrow{\nu y.\bar{c}\langle y \rangle} \xrightarrow{c(\langle \pi_2(x), \pi_1(x), \pi_4(x), \pi_3(x), \pi_5(x) \rangle)} \xrightarrow{*} \xrightarrow{\nu z.\bar{c}\langle z \rangle} \nu \tilde{n}.\tau_2$$

for some names \tilde{n} and substitution τ_2 , such that:

$$\text{dec}(\pi_1(z), \pi_1(x) * \pi_1(y) * \pi_2(x))\tau_2 =_E \text{one}$$

It follows immediately that $\nu \tilde{n}.\tau_1 \not\approx_s \nu \tilde{n}.\tau_2$ and, hence, ϕ_{ident} does not satisfy ballot secrecy. \square

In contrast, removing duplicates up to permutation ensures ballot secrecy.

Theorem 1. *The Helios process specification ϕ_{sol} satisfies ballot secrecy.*

ProVerif is an automatic tool that can check equivalence in the applied pi calculus [BAF08]. Although ProVerif has been successfully used to prove ballot secrecy (for example, in the Fujioka, Okamoto & Ohta protocol [DRS08]), it cannot prove Theorem 1 for two main reasons. First, ProVerif cannot prove equivalences under the homomorphic equation (Equation E6). Second, our theorem states ballot secrecy for any number n of participants and ProVerif cannot handle parametrised processes. We proceed by constructing a relation that relates $A_n^\phi[V\{a_1/x_{\text{auth}}\}\sigma \mid V\{a_2/x_{\text{auth}}\}\tau]$ and $A_n^\phi[V\{a_1/x_{\text{auth}}\}\tau \mid V\{a_2/x_{\text{auth}}\}\sigma]$, and all their successors, such that it satisfies the three properties of Definition 2. In particular, the two final frames (containing the result of the election) should be statically equivalent. A key step is to show that ballots accepted by the bulletin board must have a particular form due to the checks performed by ϕ_{sol} .

Definition 5 (Valid ballot). *A term N is said to be a valid ballot if $N = \langle N_1, N_2, N_3, N_4, N_5 \rangle$ for some N_i such that $N_1 = \text{penc}(z_{\text{pk}}, N_1^1, N_1^2)$ and $N_1 = \text{penc}(z_{\text{pk}}, N_2^1, N_2^2)$ with $N_1^1, N_2^2 \in \{\text{zero}, \text{one}\}$.*

Having shown that the bulletin board accepts only valid ballots, we can deduce that the outcome of the election at the end of the execution of $A_n^\phi[V\{a_1/x_{\text{auth}}\}\sigma \mid V\{a_2/x_{\text{auth}}\}\tau]$ is exactly the same as in $A_n^\phi[V\{a_1/x_{\text{auth}}\}\tau \mid V\{a_2/x_{\text{auth}}\}\sigma]$. We can then conclude the proof of Theorem 1 by showing that the partial decryptions and the encrypted ballots of honest voters do not leak any extra information to the adversary. The full proof appears in Appendix B.

5.4 Limitations

The limitations of our model, which we introduced to simplify the presentation and proof, are detailed below; we believe a full security proof should follow using

similar reasoning. We consider a model with only two candidates and, moreover, we make use of the (standard) definition of ballot secrecy which is limited to elections with two honest voters [KR05, DKR06, DKR09]. In addition, the definition of ballot secrecy does not consider parallel composition of protocol executions and we therefore recommend using distinct keys for each election (although we believe it should be sufficient to include an election identifier – for example, the election fingerprint – within the challenge hashes included within signatures of knowledge, similar to the methodology in Section 4.2). The administrative process A_n^ϕ enforces an ordering on voters (namely, the voter using private channel a_1 must vote first, followed by the voter using private channel a_2 , and then any remaining voters – controlled by the adversarial environment – can vote) and A_n^ϕ does not permit revoting. The signature and equational theory do not capture low-level technical details surrounding the correct construction of public keys; in particular, we do not use signatures of knowledge to verify correct key construction. We also omit signatures of knowledge that demonstrate correct construction of partial decryptions. Finally, we offer the usual caveat to formal analysis and acknowledge that our result does not imply the absence of real-world attacks (see, for example, [RS98, AR00, AR02, War03, War05]). It may, therefore, be possible to modify the ballot in a way that would not be captured by our analysis. (In particular, it is important to notice that the scheme used for signatures of knowledge is not provably non-malleable.) We partly overcome these limitations in our complimentary work [BCP⁺11] by presenting a variant of Helios that is provably secure in a cryptographic setting.

6 Attacks against other schemes

This section demonstrates that the absence of ballot independence can be exploited in other electronic voting protocols to violate privacy. In particular, we demonstrate replay attacks against schemes by Sako & Kilian [SK94] and Schoenmakers [Sch99], both of which were presented at CRYPTO, and we show that the malleable cryptographic scheme adopted by Lee *et al.* [LBD⁺04] can be exploited to launch attacks. In addition, perhaps contrary to intuition, we will argue that no general relationships exist between independence and privacy properties.

6.1 Exploiting replays in the protocol by Sako & Kilian

The Sako & Kilian [SK94] electronic voting scheme capitalises upon advances in cryptography to improve the Banaloh & Yung protocol [BY86]. The scheme is interesting because it was one of the first electronic voting protocols to adopt the Fiat-Shamir heuristic to derive non-interactive proofs (this evolution was key for the development of end-to-end verifiable electronic voting systems). However, we will show that the application of the Fiat-Shamir heuristic compromises ballot secrecy. In particular, the interactive nature of zero-knowledge proofs guarantees ballot independence; whereas, non-interactive proofs, derived using

the Fiat-Shamir heuristic, do not assure independence. This can be exploited by a replay attack to violate ballot secrecy.

6.1.1 Protocol description

The scheme is based upon a pair of *partially compatible homomorphic encryption* functions, that is, a pair of functions f_1, f_2 over \mathbb{Z}_q , where q is prime, such that for all $i, j \in \{1, 2\}$ the following properties are satisfied:

- $f_i(x + y) = f_i(x) \cdot f_i(y)$, where $x, y \in \mathbb{Z}_q$
- Distributions $(f_i(x), f_j(y))$ and $(f_i(x), f_j(x))$ are computationally indistinguishable, where x and y are chosen uniformly in \mathbb{Z}_q .

The voting protocol is defined for $m \in \mathbb{N}$ voters as follows.

Setup. Talliers \mathcal{T} and \mathcal{T}' publish public keys k and k' for a public key encryption scheme E (which need not be homomorphic).

Voting. Given vote $v_i \in \{-1, 1\}$, the voter generates nonces $x_i, x'_i \in \mathbb{Z}_q$ such that $v_i = x_i + x'_i$ and constructs her ballot as follows:

$$\begin{aligned} Y_i &= f_1(x_i) \\ Y'_i &= f_2(x'_i) \\ Z_i &= E(k, x_i) \\ Z'_i &= E(k', x'_i) \end{aligned}$$

In addition, the voter is required to prove $x_i + x'_i \in \{1, -1\}$ in zero-knowledge. However, to avoid an interactive proof, the Fiat-Shamir heuristic is applied to derive a signature of knowledge σ_i . (For brevity we omit the construction of σ_i , see [SK94, Figure 1] for details.)

Tallying. Given ballots $Y_1, Y'_1, Z_1, Z'_1, \sigma_1, \dots, Y_n, Y'_n, Z_n, Z'_n, \sigma_n$, tallier \mathcal{T} decrypts each Z_i to recover \hat{x}_i and checks $Y_i = f_1(\hat{x}_i)$, similarly, tallier \mathcal{T}' decrypts Z'_i to recover \hat{x}'_i and checks $Y'_i = f_2(\hat{x}'_i)$; the talliers also check the signature of knowledge σ_i . The talliers publish $V = \sum_{i=1}^m \hat{x}_i$ and $V' = \sum_{i=1}^m \hat{x}'_i$, and the result is $T = V + V'$, which can be verified by checking $f_1(V) = \prod_{i=1}^m Y_i$ and $f_2(V') = \prod_{i=1}^m Y'_i$.

6.1.2 Attacking ballot secrecy

We show that the voting protocol by Sako & Kilian does not satisfy ballot secrecy, by presenting a replay attack which allows an adversary to reveal a voter's vote. Intuitively, an adversary may observe the ballot posted by a particular voter and recast this ballot by corrupting dishonest voters. The multiple occurrences of the voter's ballot will leak information in the tally and the adversary can exploit this knowledge to violate the voter's privacy. An informal description of the attack will now be presented in the case of three eligible voters.

Let us consider an election with three eligible voters who have identities id_1 , id_2 and id_3 . Suppose that voters id_1 , id_2 are honest and id_3 is a dishonest voter controlled by the adversary. Further assume that the adversary has observed the ballot

$$Y_k, Y'_k, Z_k, Z'_k, \sigma_k$$

being cast by the voter whose privacy will be compromised.

Exploiting the absence of ballot independence. As shown by Gennaro [Gen95], an adversary can replay the ballot $Y_k, Y'_k, Z_k, Z'_k, \sigma_k$, thereby violating ballot independence. (The violation of ballot independence is due to the adversary's ability to cast the *same* vote as the honest voter.) Since the ballot was constructed by an honest voter, it is trivial to see that it will be considered valid by the talliers. We will now show how the lack of ballot independence can be exploited to violate privacy.

Violating privacy. The bulletin board will be constructed as follows

$$Y_1, Y'_1, Z_1, Z'_1, \sigma_1, Y_2, Y'_2, Z_2, Z'_2, \sigma_2, Y_k, Y'_k, Z_k, Z'_k, \sigma_k, V, V'$$

where $k \in \{1, 2\}$, $V = x_1 + x_2 + x_k$ and $V' = x'_1 + x'_2 + x'_k$. It follows from the protocol description that $v_i = x_i + x'_i$, where $i \in \{1, 2, k\}$, and the result $T = V + V' = v_1 + v_2 + v_k$. Since there will be at least two votes for the candidate voter id_k voted for, the voter's vote can be revealed: if $T \geq 2$, then $v_k = 1$; otherwise $v_k = -1$. It follows that the voter's privacy has been compromised; moreover, the vote of the remaining honest voter is $T - 2 \cdot v_k$.

6.1.3 Independence and the Fiat-Shamir heuristic

The interactive nature of zero-knowledge proofs guarantee independence; by comparison, non-interactive proofs, derived using the Fiat-Shamir heuristic, do not assure independence. As a consequence, application of the Fiat-Shamir heuristic may compromise the security of cryptographic protocols and we have shown how application of the heuristic erodes privacy in the electronic voting scheme by Sako & Kilian. This demonstrates that the use of the Fiat-Shamir heuristic requires some care and highlights the necessity for thorough security analysis.

6.1.4 Generalising replay attacks

The replay attack against the voting protocol by Sako & Kilian can be generalised to other schemes where an adversary can observe a ballot cast by a particular voter and replay this ballot verbatim. In particular, the voting protocol by Schoenmakers [Sch99] fits this description.

Exploiting replays in the protocol by Schoenmakers. The electronic voting protocol by Schoenmakers [Sch99] is based upon [CFSY96, CGS97]. The scheme explicitly aims to provide efficient small-scale elections (for example, boardroom elections) and, given that our attack is particularly well suited to small-scale elections, we find it interesting to study the possibility of violating ballot secrecy in this setting. Ballot independence is not provided [Sch99, §5] and we exploit privacy using a replay attack. The attack description is straightforward and follows immediately from our discussion; accordingly, we omit the details and refer the interested reader to our technical report [SC11, §3].

6.1.5 Possible solutions: Weeding duplicate ballots

Our verbatim replay attacks against the voting protocols by Sako & Kilian and Schoenmakers exploit the possibility of replaying a voter’s ballot without detection. We believe it should be sufficient for the election officer to reject any duplicate ballots to ensure ballot secrecy, alternatively, the *binding ballots to voters* solution (Section 4.2) may also be suitable, although proving the security of these solutions remains an open problem.

6.2 Exploiting malleability in the protocol by Lee *et al.*

The Lee *et al.* [LBD⁺04] electronic voting scheme adopts an offline tamper-resistant hardware device to ensure receipt freeness; more precisely, the hardware device takes an ElGamal encrypted vote as input and outputs a re-encrypted ciphertext, this prevents a voter proving how she voted by reconstruction as she does not know the nonce introduced for re-encryption. In addition, the hardware device provides a Designated Verifier Proof of re-encryption, thereby allowing the voter to verify that the device behaved correctly. The device is assumed to be offline and, hence, communication between the voter and the device is assumed to be untappable.

6.2.1 Background: Multiplicative homomorphic ElGamal

The scheme uses multiplicative homomorphic ElGamal, rather than the additive variant presented in Section 2.1. The operations for key generation, homomorphic combination and re-encryption are standard; albeit, the result of homomorphic combination is the multiplication of plaintexts, rather than the addition of plaintexts. We recall the operations for encryption and decryption below.

Encryption. Given a message m and a public key h , select a random nonce $r \in_R \mathbb{Z}_q^*$ and derive the ciphertext $(a, b) = (g^r \bmod p, m \cdot h^r \bmod p)$.

Distributed decryption. Given a ciphertext (a, b) , each trustee $i \in n$ computes the partial decryption $k_i = a^{x_i}$. The plaintext $m = b / (k_1 \cdot \dots \cdot k_n) \bmod p$.

The application of these primitives to derive the scheme by Lee *et al.* will be discussed in the next section.

6.2.2 Protocol description

An election is created by naming an election officer, selecting a set of mixers, and choosing a set of trustees. The trustees generate a distributed public key pair and the election officer publishes the public key on the bulletin board. (For robustness, threshold ElGamal may be used; we omit these details for brevity.) The election officer also publishes the candidate list, the public keys of eligible voters, and the public keys of the tamper-resistant hardware devices. Informally, the steps that the participants take during an election are as follows.

1. The voter constructs an ElGamal ciphertext (a, b) containing her vote v and sends the ciphertext to her tamper-resistant hardware device.
2. The hardware device re-encrypts the voter's ciphertext to produce (a', b') and computes a Designated Verifier Proof of re-encryption τ . The device also derives a signature σ on the re-encryption. The hardware device returns $(a', b'), \sigma, \tau$ to the voter.
3. If the signature and proof are valid, then the voter generates a signature σ' on the message σ using her private key. The voter submits her ballot $(a', b'), \sigma, \sigma'$ to the bulletin board.
4. Individual voters can check that their ballots appear on the bulletin board and can be assured that the ciphertext (a', b') contains their vote v by verifying the Designated Verifier Proof τ .
5. Voters and observers can check that ballots were cast by registered voters by verifying signatures σ' , and are assured that each voter cast at most one ballot by checking that no voter signed two values. In addition, voters and observers should verify signatures σ for receipt freeness.
6. After some predefined deadline, valid ballots (that is, ballots associated with valid signatures σ and σ') are submitted to the mixers. Anyone can check that mixing is performed correctly.
7. Each of the trustees publishes a partial decryption for every ciphertext output by the mix. Anyone can verify these proofs.
8. The election officer decrypts each ciphertext and publishes the election result. Anyone can check these decryptions.

See Lee *et al.* [LBD⁺04] for further details.

6.2.3 Attacking ballot secrecy

We show that the voting protocol by Lee *et al.* [LBD⁺04] does not satisfy ballot secrecy, by presenting a replay attack that exploits malleability to reveal a voter’s vote. Intuitively, an adversary may identify a voter’s encrypted vote on the bulletin board, since it is signed by the voter. This ciphertext can be submitted to a tamper-resistant hardware device (possibly after re-encryption) and the device will return $(\hat{a}, \hat{b}), \hat{\sigma}, \hat{\tau}$; the ballot $(\hat{a}, \hat{b}), \hat{\sigma}, \hat{\sigma}'$ can then be submitted by the adversary to the bulletin board, where $\hat{\sigma}'$ is a signature on $\hat{\sigma}$ constructed by a registered voter under the adversary’s control. As explain in Section 6.1.2, the multiple occurrences of the voter’s ballot will leak information in the tally and the adversary can exploit this knowledge to violate the voter’s privacy.

Variant exploiting homomorphic encryption. The adversary can exploit the homomorphic properties of ElGamal to avoid casting the *same* vote as an honest voter. In this variant, suppose the adversary wants to recover the vote from ballot $(a'_k, b'_k), \sigma_k, \sigma'_k$, the adversary derives the ciphertext $(c, d) = (a'_k, b'_k) \cdot (c', d')$, where (c', d') is an ElGamal ciphertext containing some message m . The adversary submits the ciphertext (c, d) to a tamper-resistant hardware device and the device will return $(\hat{c}, \hat{d}), \hat{\sigma}, \hat{\tau}$; the ballot $(\hat{c}, \hat{d}), \hat{\sigma}, \hat{\sigma}'$ can be submitted by the adversary to the bulletin board, where $\hat{\sigma}'$ is a signature on $\hat{\sigma}$ constructed by a registered voter. The output of the mix will include the adversaries re-encrypted ciphertext and the election officer will publish $m \cdot v$ on the bulletin board, where ciphertext (a'_k, b'_k) includes the vote v . This variant of the attack is particularly interesting because the replayed ballots are undetectable; in particular, weeding duplicate ballots would clearly not be sufficient to ensure privacy.

6.2.4 Privacy in the variant by Küsters & Truderung

Independently, the scheme by Lee *et al.* is vulnerable to a forced abstention attack [KT09, §6.2], whereby the voter votes for a candidate not included on the candidate list. This can be trivially witnessed since casting a vote for some random nonce will result in a unique occurrence of that nonce being published on the bulletin board by the election officer. Accordingly, Küsters & Truderung [KT09, §6.3] propose a variant of the protocol by Lee *et al.* which they claim satisfies coercion resistance. In the revised scheme the voter must prove to her tamper-resistant hardware device that her ciphertext contains a valid vote. However, our attack is still valid under one of the following assumptions: 1) there exists a tamper-resistant hardware device which does not check the proof that the voter’s ciphertext is correctly formed; 2) there exists a signing key which has been extracted from a hardware device; or 3) the election officer publishes the public part of a signing key where the private part is known to the adversary. The attack follows immediately from our original description under assumption 1 and, under assumptions 2 or 3, the adversary signs the encrypted ciphertext without using the tamper-resistant hardware device.

Figure 5 Helios administrator that preserves independence but not privacy

Given the number of voters $n \geq 2$ the administration process $\bar{A}_n^{\phi_{\text{sol}}}$ is defined below, where process T is presented in Figure 4.

$$\begin{aligned} \bar{A}_n^{\phi_{\text{sol}}} &= \nu \text{sk}_T, a_1, a_2, d. (- \mid \bar{B}\bar{B}_n^{\phi_{\text{sol}}} \mid !T \mid \{\text{pk}(\text{sk}_T)/z_{\text{pk}}\}) \\ \bar{B}\bar{B}_n^{\phi_{\text{sol}}} &= a_1(y_1) . \bar{c}\langle y_1 \rangle . a_2(y_2) . \bar{c}\langle y_2 \rangle . \\ &\quad a_3(y_3) . \text{if } \phi_{\text{sol}}\{y_3/y_{\text{ballot}}\} \text{ then} \\ &\quad \dots a_n(y_n) . \text{if } \phi_{\text{sol}}\{y_n/y_{\text{ballot}}\} \text{ then} \\ &\quad \bar{d}\langle (\pi_1(y_1), \pi_2(y_1)) \rangle . d\langle z_1 \rangle . \bar{c}\langle z_1 \rangle . \\ &\quad \dots . \bar{d}\langle (\pi_1(y_n), \pi_2(y_n)) \rangle . d\langle z_n \rangle . \bar{c}\langle z_n \rangle \end{aligned}$$

6.3 Independence and privacy are unrelated properties

Intuitively, these attacks may be suggestive of a general relationship between independence and privacy properties, however, we shall now present examples that suggest independence does not imply privacy and vice-versa.

A protocol with independence but no privacy. Consider a variant of the fixed Helios voting scheme in which each of the trustees publish a partial decryption of individual ciphertexts (rather than a partial decryption of the homomorphically combined ciphertexts, that is, the encrypted tally). Intuitively, this variant preserves ballot independence but does not satisfy ballot secrecy, since the partial decryptions allow votes to be recovered from ballots and the link between a voter and her ballot is known. Formally, this variant is captured by modelling the Helios administrator process as $\bar{A}_n^{\phi_{\text{sol}}}$, defined in Figure 5. The violation of ballot secrecy can be witnessed since

$$\bar{A}_2^{\phi_{\text{sol}}}[V\{a_1/x_{\text{auth}}\}\sigma \mid V\{a_2/x_{\text{auth}}\}\tau] \not\approx_l \bar{A}_2^{\phi_{\text{sol}}}[V\{a_1/x_{\text{auth}}\}\tau \mid V\{a_2/x_{\text{auth}}\}\sigma]$$

where $\sigma = \{\text{zero}/x_{\text{vote}}, \text{one}/x'_{\text{vote}}\}$ and $\tau = \{\text{one}/x_{\text{vote}}, \text{zero}/x'_{\text{vote}}\}$. Similarly, a further variant of the fixed Helios scheme in which each of the trustees publishes their private key at the end of the voting phase, rather than a partial decryption of the encrypted tally, also satisfies independence but not ballot secrecy.

A protocol with privacy but no independence. Consider a voting scheme in which each voter broadcasts their vote on an anonymous communication channel. Formally, the voter is modelled by the process $P = \bar{c}\langle x_{\text{vote}} \rangle$, where variable x_{vote} is parametrised by the voter's vote. For ballot secrecy it is sufficient to show $P\{M/x_{\text{vote}}\} \mid P\{N/x_{\text{vote}}\} \approx_l P\{N/x_{\text{vote}}\} \mid P\{M/x_{\text{vote}}\}$ for all ground terms M and N ; this result trivially holds by structural equivalence and hence the scheme satisfies ballot secrecy. However, independence is intuitively violated in this setting because an adversary may observe the voting system and replay a previously cast vote, that is, an adversary can cast the same vote as another

voter (without knowing which voter). In addition, it follows that early results are available in this scheme.

We also expect some published electronic voting schemes based upon blind signatures to satisfy ballot secrecy but not independence; in particular, a more realistic example of a protocol that satisfies this property is the protocol by Fujioka, Okamoto & Ohta [FOO92] under the assumption that duplicates are not rejected. Indeed, independence can be violated by a verbatim replay of the signed committed vote.

In this article, we cannot make any definitive mathematical statements about the relationship between independence and privacy properties, because independence has not been formally defined; however, we hope these examples provide some insight into the relationships we expect.

7 Conclusion and further discussion

This article identifies a vulnerability in the Helios 2.0 electronic voting protocol which can be used to violate ballot secrecy. Critics may argue that an attack is unrealistic due its high cost; indeed, in some cases, the attack may change the outcome of an election (that is, the votes introduced for the purposes of violating privacy may swing the result), and large scale privacy invasions would be expensive in terms of the required number of dishonest voters. However, if the views of these critics are to be entertained, then we must revise the standard definitions of ballot secrecy in the literature (for example, [KR05, DKR06, BHM08]) because Helios cannot satisfy them. Furthermore, we believe all voters should be considered equally and, hence, the preservation of ballot secrecy should be universal. But, for elections using Helios, our case study demonstrates the contrary: in French legislative elections a coalition of voters can gain some information about a voter’s vote in an arbitrary polling station and, moreover, if the number of voters registered at a particular polling station is small (for example, in a rural setting), then a voter’s privacy can be violated by a few dishonest voters. It follows that privacy of individual voters can be compromised by a few dishonest voters and, accordingly, we believe our attack is significant. To address the problem, we have introduced a variant of the Helios protocol which has been shown to satisfy definitions of ballot secrecy in the applied pi calculus and in our complimentary work [BCP⁺11] we present a security proof in the cryptographic setting. The vulnerability in Helios has been acknowledged by Adida & Pereira [Adi10, AP10] and they have scheduled a fix for future Helios releases. We have also shown that the absence of ballot independence can be similarly exploited in other electronic voting protocols to violate privacy; in particular, we demonstrate verbatim replay attacks against the schemes by Sako & Kilian [SK94] and Schoenmakers [Sch99], and we show that the malleable cryptographic scheme adopted by Lee *et al.* [LBD⁺04] can be exploited to replay a voter’s ballot or a variant of it, thereby violating ballot secrecy. In addition, we argue that independence does not imply privacy and

vice-versa.

Acknowledgements

We are grateful to Ben Adida and Olivier Pereira for their constructive comments, and hope this research will enhance future Helios releases. Discussion with Mark D. Ryan helped clarify the presentation of this article, and Ben Adida informed us that Douglas Wikström is the contemporaneous discoverer of the Helios attack. David Bernhard gave useful feedback on the variants of our attack against Helios. The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement n° 258865, project ProSecure, and the ANR-07-SeSur-002 AVOTÉ project.

A Signatures of knowledge

Helios is reliant on signatures of knowledge to ensure secrecy and integrity of the ElGamal scheme, and to ensure voters encrypt valid votes. This appendix presents suitable cryptographic primitives. Let \mathcal{H} denote a hash function. In Helios, \mathcal{H} is defined to be SHA-256.

A.1 Knowledge of discrete logs

Given the aforementioned cryptographic parameters (p, q, g) , a signature of knowledge demonstrating knowledge of a discrete logarithm $h = \log_g g^x$ can be derived, and verified, as defined by [CEGP87, CEG88, Sch90].

Sign. Given x , select a random nonce $w \in_R \mathbb{Z}_q^*$. Compute witness $g' = g^w \bmod p$, challenge $c = \mathcal{H}(g') \bmod q$ and response $s = w + c \cdot x \bmod q$.

Verify. Given h and signature g', s , check $g^s \equiv g' \cdot h^c \pmod{p}$, where $c = \mathcal{H}(g') \bmod q$.

A valid proof asserts knowledge of x such that $x = \log_g h$; that is, $h \equiv g^x \bmod p$.

A.2 Equality between discrete logs

Given the aforementioned cryptographic parameters (p, q, g) , a signature of knowledge demonstrating equality between discrete logarithms $\log_f f^x$ and $\log_g g^x$ can be derived, and verified, as defined by [Ped91, CP93].

Sign. Given f, g, x , select a random nonce $w \in_R \mathbb{Z}_q^*$. Compute witnesses $f' = f^w \bmod p$ and $g' = g^w \bmod p$, challenge $c = \mathcal{H}(f', g') \bmod q$ and response $s = w + c \cdot x \bmod q$.

Verify. Given f, g, h, k and signature f', g', s , check $f^s \equiv f' \cdot h^c \pmod{p}$ and $g^s \equiv g' \cdot k^c \pmod{p}$, where $c = \mathcal{H}(f', g') \pmod{q}$.

A valid proof asserts $\log_f h = \log_g k$; that is, there exists x , such that $h \equiv f^x \pmod{p}$ and $k \equiv g^x \pmod{p}$. This signature of knowledge scheme can be extended to a disjunctive proof of equality between discrete logs (see below).

For our purposes, given a ciphertext (a, b) , each trustee would derive a signature on g, a, x_i , where x_i is the trustee's private key share. The i th trustee's signature g'_i, a'_i, c_i, s_i would be verified with respect to g, a, h_i, k_i , where h_i is the trustee's share of the public key and k_i is the trustee's partial decryption; that is, the proof asserts $\log_g h_i = \log_a k_i$, as required for integrity of decryption.

A.3 Disjunctive proof of equality between discrete logs

Given the aforementioned cryptographic parameters (p, q, g) , a signature of knowledge demonstrating that a ciphertext (a, b) contains either 0 or 1 (without revealing which), can be constructed by proving that either $\log_g a = \log_h b$ or $\log_g a = \log_h b/g^m$; that is, a signature of knowledge demonstrating a disjunct proof of equality between discrete logarithms [CDS94, Sch09]. Observe for a valid ciphertext (a, b) that $a \equiv g^r \pmod{p}$ and $b \equiv h^r \cdot g^m \pmod{p}$ for some nonce $r \in \mathbb{Z}_q^*$; hence the former disjunct $\log_g g^r = \log_h h^r \cdot g^m$ is satisfied when $m = 0$, and the latter $\log_g g^r = \log_h (h^r \cdot g^m)/g^m$ when $m = 1$.

This technique is generalised by [AMPQ09] to allow a signature of knowledge demonstrating that a ciphertext (a, b) contains message m , where $m \in \{\min, \dots, \max\}$ for some system parameters $\min, \max \in \mathbb{N}$. Formally, a signature of knowledge demonstrating a disjunct proof of equality between discrete logarithms can be derived, and verified, as follows [AMPQ09, CDS94, Sch09].

Sign. Given ciphertext (a, b) such that $a \equiv g^r \pmod{p}$ and $b \equiv h^r \cdot g^m \pmod{p}$ for some nonce $r \in \mathbb{Z}_q^*$, where plaintext $m \in \{\min, \dots, \max\}$. For all $i \in \{\min, \dots, m-1, m+1, \dots, \max\}$, compute challenge $c_i \in_R \mathbb{Z}_q^*$, response $s_i \in_R \mathbb{Z}_q^*$ and witnesses $a_i = g^{s_i}/a^{c_i} \pmod{p}$ and $b_i = h^{s_i}/(b/g^i)^{c_i} \pmod{p}$. Select a random nonce $w \in_R \mathbb{Z}_q^*$. Compute witnesses $a_m = g^w \pmod{p}$ and $b_m = h^w \pmod{p}$, challenge $c_m = \mathcal{H}(a_{\min}, b_{\min}, \dots, a_{\max}, b_{\max}) - \sum_{i \in \{\min, \dots, m-1, m+1, \dots, \max\}} c_i \pmod{q}$ and response $s_m = w + r \cdot c_m \pmod{q}$.

Verify. Given (a, b) and $(a_{\min}, b_{\min}, c_{\min}, s_{\min}, \dots, a_{\max}, b_{\max}, c_{\max}, s_{\max})$, for each $\min \leq i \leq \max$ check $g^{s_i} \equiv a_i \cdot a^{c_i} \pmod{p}$ and $h^{s_i} \equiv b_i \cdot (b/g^i)^{c_i} \pmod{p}$. Finally, check $\mathcal{H}(a_{\min}, b_{\min}, \dots, a_{\max}, b_{\max}) \equiv \sum_{\min \leq i \leq \max} c_i \pmod{q}$.

A valid proof asserts that (a, b) is a ciphertext containing the message m such that $m \in \{\min, \dots, \max\}$.

B Proof of Theorem 1

B.1 Preliminaries

Before commencing our proof, let us first introduce some useful lemmas for the applied pi calculus.

Lemma 3. *Given frames φ, ψ , ground term M and variable $x \notin \text{dom}(\varphi) \cup \text{dom}(\psi)$, we have $\varphi \approx_s \psi$ iff $\varphi \mid \{M/x\} \approx_s \psi \mid \{M/x\}$.*

Lemma 4. *Given frames φ, ψ , terms M, N , and a variable $x \notin \text{dom}(\varphi) \cup \text{dom}(\psi)$, such that $\varphi = \nu \tilde{m}.\sigma$ and $\psi = \nu \tilde{n}.\tau$ for some names \tilde{m}, \tilde{n} and substitutions σ, τ , we have $\nu \tilde{m}.\sigma \mid \{M/x\} \approx_s \nu \tilde{n}.\tau \mid \{N/x\}$ implies $\varphi \approx_s \psi$.*

The proofs of these lemmas are straightforward.

The following lemma shows when static equivalence implies the same branching behaviour for conditionals.

Lemma 5. *Given extended processes $A \equiv C[\text{if } M = N \text{ then } P \text{ else } Q]$ and $B \equiv C'[\text{if } M = N \text{ then } P' \text{ else } Q']$ such that $A \approx_s B$, $(\text{bn}(C) \cup \text{bn}(C')) \cap (\text{fn}(M) \cup \text{fn}(N)) = \emptyset$, $\text{fv}(M) \cup \text{fv}(N) \subseteq \text{dom}(C)$ and $\text{fv}(M) \cup \text{fv}(N) \subseteq \text{dom}(C')$, for some closing evaluation context C, C' , terms M, N and processes P, P', Q, Q' , then $A \rightarrow C[P]$ iff $B \rightarrow C'[P']$ and $A \rightarrow C[Q]$ iff $B \rightarrow C'[Q']$.*

Proof. Suppose $A \equiv C[\text{if } M = N \text{ then } P \text{ else } Q]$ and $B \equiv C'[\text{if } M = N \text{ then } P' \text{ else } Q']$ such that $A \approx_s B$, $(\text{bn}(C) \cup \text{bn}(C')) \cap (\text{fn}(M) \cup \text{fn}(N)) = \emptyset$, $\text{fv}(M) \cup \text{fv}(N) \subseteq \text{dom}(C)$ and $\text{fv}(M) \cup \text{fv}(N) \subseteq \text{dom}(C')$, for some closing evaluation context C, C' , terms M, N and processes P, P', Q, Q' . Further suppose $\varphi(C[\text{if } M = N \text{ then } P \text{ else } Q]) = \nu \tilde{m}.\sigma$ and $\varphi(C'[\text{if } M = N \text{ then } P' \text{ else } Q']) = \nu \tilde{n}.\tau$, for some names \tilde{m} and \tilde{n} . By Lemma 6 we have $\nu \tilde{m}.\sigma \approx_s \nu \tilde{n}.\tau$, because static equivalence is closed under structural equivalence. Moreover, by the definition of static equivalence, for all terms U, V such that $(\tilde{m} \cup \tilde{n}) \cap (\text{fn}(U) \cup \text{fn}(V)) = \emptyset$, we have $U\sigma =_E V\sigma$ iff $U\tau =_E V\tau$.

Let us first show $A \rightarrow C[P]$ iff $B \rightarrow C'[P']$. For the \Rightarrow implication, suppose $A \rightarrow C[P]$. Since $\text{fv}(M) \cup \text{fv}(N) \subseteq \text{dom}(C)$, it must be the case that $M\sigma =_E N\sigma$. We have $\tilde{m} \cup \tilde{n} \subseteq \text{bn}(C) \cup \text{bn}(C')$ by definition of the function φ , and we derive $(\tilde{m} \cup \tilde{n}) \cap (\text{fn}(M) \cup \text{fn}(N)) = \emptyset$ because $(\text{bn}(C) \cup \text{bn}(C')) \cap (\text{fn}(M) \cup \text{fn}(N)) = \emptyset$; it follows that $M\sigma =_E N\sigma$ is a special case of $U\sigma =_E V\sigma$. We derive $M\tau =_E N\tau$ from the implication $(U\sigma =_E V\sigma) \Rightarrow (U\tau =_E V\tau)$. It trivially follows that $B \equiv C'[\text{if } M\tau = N\tau \text{ then } P' \text{ else } Q']$, and by closure of internal reduction under structural equivalence we derive $B \rightarrow C'[P']$. The \Leftarrow implications follows by symmetry.

We will now show $A \rightarrow C[Q]$ iff $B \rightarrow C'[Q']$. For the \Rightarrow implication, suppose $A \rightarrow C[Q]$. It must be the case that $M\sigma \neq_E N\sigma$ and, as before, we derive $M\tau \neq_E N\tau$. It trivially follows that $B \equiv C'[\text{if } M\tau = N\tau \text{ then } P' \text{ else } Q']$, and since $\text{fv}(M) \cup \text{fv}(N) \subseteq \text{dom}(C')$ we are assured that terms $M\tau, N\tau$ are ground; by closure of internal reduction under structural equivalence we derive $B \rightarrow C'[Q']$. The \Leftarrow implications follows by symmetry. \square

This result can naturally be extended to formula. Given ϕ , let us denote the set of free names, respectively variables, in ϕ as $\text{fn}(\phi)$, respectively $\text{fv}(\phi)$.

Corollary 1. *Given extended processes $A \equiv C[\text{if } \phi \text{ then } P \text{ else } Q]$ and $B \equiv C'[\text{if } \phi \text{ then } P' \text{ else } Q']$ such that $A \approx_s B$, $(\text{bn}(C) \cup \text{bn}(C')) \cap \text{fn}(\phi) = \emptyset$, $\text{fv}(\phi) \subseteq \text{dom}(C)$ and $\text{fv}(\phi) \subseteq \text{dom}(C')$, for some closing evaluation context C, C' , formulae ϕ and processes P, P', Q, Q' , then $A \rightarrow C[P]$ iff $B \rightarrow C'[P']$ and $A \rightarrow C[Q]$ iff $B \rightarrow C'[Q']$.*

We conclude this subsection with a useful result stated by Abadi & Fournet [AF01].

Lemma 6. *Static equivalence is closed by structural equivalence.*

B.2 Notations and Definitions

We henceforth assume $(m_1, m'_1), (m_2, m'_2) \in \{(\text{zero}, \text{zero}), (\text{zero}, \text{one}), (\text{one}, \text{zero})\}$.

We introduce the following notations:

$$\begin{aligned}
\text{tally} &= \pi_1(y_1) * \dots * \pi_1(y_n) \\
\text{tally}' &= \pi_2(y_1) * \dots * \pi_2(y_n) \\
\text{result} &= \text{dec}(\text{partial}(sk_T, \text{tally}), \text{tally}) \\
\text{result}' &= \text{dec}(\text{partial}(sk_T, \text{tally}'), \text{tally}') \\
\text{ciph}_i &= \text{penc}(z_{pk}, r_i, x_{\text{vote},i}) \\
\text{ciph}'_i &= \text{penc}(z_{pk}, r'_i, x'_{\text{vote},i}) \\
\text{spk}_i &= \text{spk}(z_{pk}, r_i, x_{\text{vote},i}, \text{ciph}_i) \\
\text{spk}'_i &= \text{spk}(z_{pk}, r'_i, x'_{\text{vote},i}, \text{ciph}'_i) \\
\widehat{\text{spk}}_i &= \text{spk}(z_{pk}, r_i \circ r'_i, x_{\text{vote},i} + x'_{\text{vote},i}, \text{ciph}_i * \text{ciph}'_i) \\
\text{ballot}_i &= (\text{ciph}_i, \text{ciph}'_i, \text{spk}_i, \text{spk}'_i, \widehat{\text{spk}}_i) \\
\sigma_L &= \{m_1/x_{\text{vote},1}, m'_1/x'_{\text{vote},1}\} \\
\sigma_R &= \{m_2/x_{\text{vote},1}, m'_2/x'_{\text{vote},1}\} \\
\tau_L &= \{m_1/x_{\text{vote},1}, m'_1/x'_{\text{vote},1}, m_2/x_{\text{vote},2}, m'_2/x'_{\text{vote},2}\} \\
\tau_R &= \{m_2/x_{\text{vote},1}, m'_2/x'_{\text{vote},1}, m_1/x_{\text{vote},2}, m'_1/x'_{\text{vote},2}\}
\end{aligned}$$

for $n \in \mathbb{N}$, $i \in \{1, 2\}$.

Given N_3, \dots, N_k valid ballots, we define

$$\sigma_{\tilde{N}_k} = \{\text{ballot}_1/y_1, \text{ballot}_2/y_2, N_j/y_j \mid j \in \{3, \dots, k\}\}$$

Given an integer $k \in \mathbb{N}^+$ and a term N , we define N^k (resp. $k.N$) to be $N \circ \dots \circ N$ (resp. $N + \dots + N$) where N is replicated k times.

We associate to the equational theory E a rewriting system \mathcal{R}_E by orienting the Equations E1, E2 and E4 to E8 from left to right. We denote by E' the

equational theory that asserts functions $+$, $*$, \circ are commutative and associative in addition to Equation E3. \mathcal{R}_E modulo E' forms a convergent rewriting system (modulo E'). We denote by $u \rightarrow_E v$ (or often simply $u \rightarrow v$) if u modulo E' can be rewritten to v modulo E' , using \mathcal{R}_E . We denote by $u \downarrow$ a normal form of u modulo E' .

B.3 Some useful lemmas

We prove some useful results about our definitions and notations. First, we remark that tallying valid ballots yield the same result in both worlds.

Lemma 7. *Let N_3, \dots, N_k be valid ballots.*

$$\begin{aligned} \text{result } \sigma_{\tilde{N}_k} \tau_L &= \text{result } \sigma_{\tilde{N}_k} \tau_R, \\ \text{result}' \sigma_{\tilde{N}_k} \tau_L &= \text{result}' \sigma_{\tilde{N}_k} \tau_R \end{aligned}$$

and both are a sum of ones (or zero).

This follows immediately from the equational theory and the definition of valid ballots.

We now show that the encrypted ballots of honest voters and the partial decryptions do not leak any information to the adversary.

Lemma 8. *Let $\phi_6 = \nu sk_T, d, r_1, r'_1, r_2, r'_2 \cdot (\{ballot_1/x_1\} \mid \{ballot_2/x_2\} \mid \{pk(sk_T)/z_{pk}\} \mid \{\text{partial}(sk_T, tally)\}, \text{partial}(sk_T, tally')/x_{\text{partial}}\})$.*

Then $\phi_6 \sigma_{\tilde{N}_k} \tau_L \approx_s \phi_6 \sigma_{\tilde{N}_k} \tau_R$.

Proof. We first slightly decompose ϕ_6 and consider $\phi = \nu \tilde{n} \cdot \theta$ where $\tilde{n} = \{sk_T, d, r_1, r'_1, r_2, r'_2\}$ and $\theta = \{pk(sk_T)/z_{pk}\} \mid \{\text{partial}(sk_T, tally)/x'_{\text{partial}}\} \mid \{\text{partial}(sk_T, tally')/x''_{\text{partial}}\} \mid \{\{ciph_i/x_{ciph_i}\} \mid \{ciph'_i/x_{ciph'_i}\} \mid \{spk_i/x_{spk_i}\} \mid \{spk'_i/x_{spk'_i}\} \mid \{\widehat{spk}_i/x_{\widehat{spk}_i}\} \mid i \in \{1, 2\}\}$ Clearly, $\phi_6 \sigma_{\tilde{N}_k} \tau_L \approx_s \phi_6 \sigma_{\tilde{N}_k} \tau_R$ if and only if $\phi \sigma_{\tilde{N}_k} \tau_L \approx_s \phi \sigma_{\tilde{N}_k} \tau_R$.

An adversary can arbitrarily homomorphically combine ciphertexts $ciph_i$ and $ciph'_i$, possibly adding ciphertext that he has constructed. We enrich the frame ϕ with any such combination of ciphertexts. Formally, for any $\alpha_1, \alpha_2, \alpha_3, \alpha_4 \in \mathbb{N}$, for any terms P, R in normal form such that $(\text{fn}(P) \cup \text{fn}(R)) \cap \tilde{n} = \emptyset$ we define $C_{\alpha_1, \alpha_2, \alpha_3, \alpha_4, P, R} = \text{penc}(pk(sk_T), r_1^{\alpha_1} \circ r_1^{\alpha_2} \circ r_2^{\alpha_3} \circ r_2^{\alpha_4} \circ R, \alpha_1 \cdot x_{\text{vote}, 1} + \alpha_2 \cdot x'_{\text{vote}, 1} + \alpha_3 \cdot x_{\text{vote}, 2} + \alpha_4 \cdot x'_{\text{vote}, 2} + P)$.

Let extended frame $\phi_e = \nu \tilde{n} \cdot (\theta \mid \{C_{\alpha_1, \alpha_2, \alpha_3, \alpha_4, P, R}/x_{\alpha_1, \alpha_2, \alpha_3, \alpha_4, P, R} \mid \alpha_1, \alpha_2, \alpha_3, \alpha_4 \in \mathbb{N} \text{ and terms } P, R \text{ such that } (\text{fn}(P) \cup \text{fn}(R)) \cap \tilde{n} = \emptyset\})$. By Lemma 4, it is sufficient to show $\phi_e \sigma_{\tilde{N}_k} \tau_L \approx_s \phi_e \sigma_{\tilde{N}_k} \tau_R$. Let $\phi'_e = \phi_e \sigma_{\tilde{N}_k}$.

We can then show the following two claims.

Claim 1. *Let M be a term such that $\text{fv}(M) \cap (\text{fv}(\phi_e) \setminus \text{dom}(\phi_e)) = \emptyset$ and $\text{fn}(M) \cap \tilde{n} = \emptyset$. If $M \phi'_e \tau \rightarrow U$ for some $\tau \in \{\tau_R, \tau_L\}$ then there exists N such that $U =_{E'} N \phi'_e \tau$ and $M \phi'_e \tau' \rightarrow N \phi'_e \tau'$ for any $\tau' \in \{\tau_R, \tau_L\}$.*

Claim 2. *Let M, N be two terms such that $(\text{fv}(M) \cup \text{fv}(N)) \cap (\text{fv}(\phi_e) \setminus \text{dom}(\phi_e)) = \emptyset$ and $\text{fn}(M, N) \cap \tilde{n} = \emptyset$. If $M\phi'_e\tau =_{E'} N\phi'_e\tau$ for some $\tau \in \{\tau_R, \tau_L\}$ then $M\phi'_e =_{E'} N\phi'_e$.*

We can then conclude the proof. Indeed, let M, N be two terms such that $\text{fn}(M, N) \cap \tilde{n} = \emptyset$ and $M\phi_e\sigma_{\tilde{N}_k}\tau_L =_E N\phi_e\sigma_{\tilde{N}_k}\tau_L$. We assume (possibly by renaming) that $(\text{fv}(M) \cup \text{fv}(N)) \cap (\text{fv}(\phi_e) \setminus \text{dom}(\phi_e)) = \emptyset$. We have $M\phi'_e\tau_L =_E N\phi'_e\tau_L$. Thus $(M\phi'_e\tau_L) \downarrow =_{E'} (N\phi'_e\tau_L) \downarrow$. Applying repeatedly Claim 1, we deduce that there exists M' such that $(M\phi'_e\tau_L) \downarrow = M'\phi'_e\tau_L$ and $M\phi'_e\tau_R \rightarrow^* M'\phi'_e\tau_R$. Similarly, there exists N' such that $(N\phi'_e\tau_L) \downarrow = N'\phi'_e\tau_L$ and $N\phi'_e\tau_R \rightarrow^* N'\phi'_e\tau_R$. From $M'\phi'_e\tau_L =_{E'} N'\phi'_e\tau_L$ and Claim 2, we deduce $M'\phi'_e =_{E'} N'\phi'_e$. Therefore $M'\phi'_e\tau_R =_{E'} N'\phi'_e\tau_R$ and thus $M\phi'_e\tau_R =_E N\phi'_e\tau_R$, that is $M\phi_e\sigma_{\tilde{N}_k}\tau_R =_E N\phi_e\sigma_{\tilde{N}_k}\tau_R$.

Proof of Claim 1: This result is proved by inspection of the rewrite rules, using the fact that the decryption key sk_T is not deducible. More precisely, assume that $M\phi'_e\tau \rightarrow U$ for some $\tau \in \{\tau_R, \tau_L\}$. It means that there exists a rewriting rule $l \rightarrow r \in \mathcal{R}_E$ and a position p such that $M\phi'_e\tau|_p =_{E'} l\theta$ for some θ . p cannot occur below M since $\phi'_e\tau$ is in normal form. If $M|_p = l\theta'$ for some θ' then we conclude that we can rewrite M as expected. The only interesting case is thus when $M|_p$ is not an instance of l but $M\phi'_e\tau|_p$ is. By inspection of the rules, $l \rightarrow r$ can only correspond to one of the three equations E4, E5 or E6. The case of Equation E4 is ruled out by the fact that sk_T is not deducible from $\phi'_e\tau$. For Equation E5, it must be the case that $M\phi_e|_p = \text{result}$ or $M\phi_e|_p = \text{result}'$. Using Lemma 7, we deduce that $M\phi'_e|_p\tau \rightarrow R$ modulo E' where R is a sum of ones and zero. Therefore $M\phi'_e\tau \rightarrow M[R]_p\phi_e\tau$. The last case is when the rule corresponding to Equation E6 is applied. Then it must be the case that $M|_p = x * y$ with x, y variables of $\text{dom}(\phi'_e)$. By construction of ϕ'_e , we have that $(x * y)\phi'_e \rightarrow z\phi'_e$ (applying the rule corresponding to Equation E6), thus the result.

Proof of Claim 2: Assume by contradiction that there exist M, N two terms such that $M\phi'_e\tau =_{E'} N\phi'_e\tau$ for some $\tau \in \{\tau_R, \tau_L\}$ and $M\phi'_e \neq_{E'} N\phi'_e$. Consider M, N two minimal terms that satisfy this property. By case inspection, it must be the case that M and N are both variables. Thus we have $x\phi'_e\tau =_{E'} y\phi'_e\tau$ and $x\phi'_e \neq_{E'} y\phi'_e$ with $x, y \in \text{dom}(\phi_e)$, $x \neq y$. The head symbol of $x\phi'_e\tau$ must be **penc** or **partial**. Assume first that the head symbol of $x\phi'_e\tau$ is **penc**. Then by construction of ϕ_e , τ does not change the randomness used in **penc** and the randomness uniquely determines the variable, which implies $x = y$, contradiction. Assume now that the head symbol of $x\phi'_e\tau$ is **partial**. Then it must be the case that $\text{tally}\sigma_{\tilde{N}_k}\tau =_{E'} \text{tally}'\sigma_{\tilde{N}_k}\tau$ while $\text{tally}\sigma_{\tilde{N}_k} \neq_{E'} \text{tally}'\sigma_{\tilde{N}_k}$. This would require $x_{\text{ciph}_i}\tau = x'_{\text{ciph}_j}\tau$ for some i, j , which is excluded due to the randomness. \square

We further prove that ballots accepted by the bulletin board must have a particular form due to the checks performed by ϕ_{sol} .

Lemma 9. *Let M be term that does not contain the randomness r_1, r'_1, r_2, r'_2 and such that $\text{fv}(M) \subseteq \{z_{\text{pk}}, y_1, y_2\}$. Let $\tau \in \{\tau_L, \tau_R\}$. If $\llbracket \phi_{\text{sol}}\{M/y_{\text{ballot}}\}\{\text{pk}(sk_T)/z_{\text{pk}}, \text{ballot}_1/y_1, \text{ballot}_2/y_2\}\tau \rrbracket = \text{true}$ then M in normal form is a valid ballot.*

Proof. We assume w.l.o.g. that M is in normal form. We know $\llbracket \phi_{\text{sol}}\{M/y_{\text{ballot}}\}\{\text{pk}(sk_T)/z_{\text{pk}}, \text{ballot}_1/y_1, \text{ballot}_2/y_2\}\tau \rrbracket = \text{true}$. Thus it must be the case that $M\sigma$ is of the form $\langle U_1, U_2, U_3, U_4, U_5 \rangle$ with $U_1 = \text{penc}(A_1, B_1, C_1)$ and $U_2 = \text{penc}(A_2, B_2, C_2)$ and $C_1, C_2 \in \{\text{zero}, \text{one}\}$. Due to the disequality tests in ϕ_{sol} , it must be the case that M is of the form $\langle M_1, M_2, M_3, M_4, M_5 \rangle$ and $M_1, M_2 \notin \{\pi_i(y_j) \mid i, j \in \{1, 2\}\}$. Due to the equational theory, it must be the case that $M_1 = \text{penc}(M_1^1, M_1^1, M_1^2) * \pi_1(y_1)^{\alpha_1} * \pi_2(y_1)^{\alpha_2} * \pi_1(y_2)^{\alpha_3} * \pi_2(y_2)^{\alpha_4}$ where each component is optional and $\alpha_i \in \mathbb{N}$. By convention $\alpha_i = 0$ means that the component is skipped. Assume that one of the α_i is not null. Then $B_1 = r \circ B'_1$ with $r \in \{r_1, r_2, r'_1, r'_2\}$. We know $U_3 = \text{spk}(A_1, B_1, C_1, U_1)$. Moreover, U_3 cannot be one of the signatures of proof of knowledge that appear in $\text{ballot}_1\tau_L, \text{ballot}_2\tau_L$. Indeed, due to weeding, U_3 cannot be one of the two first signatures of proof of knowledge that appear in $\text{ballot}_1\tau_L$ (same thing for $\text{ballot}_2\tau_L$). Assume now (by contradiction) that $U_3 = \widehat{\text{spk}}_i$ for some $i \in \{1, 2\}$. Then $B_1 = r_i \circ r'_i$. In that case, let us have a look at U_5 . We know $U_5 = \text{spk}(A_1, B_1 \circ B_2, C_1 + C_2, U_1 * U_2)$. Thus U_5 cannot be one of the signatures of proof of knowledge that appear in $\text{ballot}_1\tau_L, \text{ballot}_2\tau_L$ (the depth of $B_1 \circ B_2$ is too big). Therefore (and due to the equational theory), we must have $M_5 = \text{spk}(M_5^1, M_5^2, M_5^3, M_5^4)$. Since $r_i \circ r'_i$ is not deducible, we cannot have $M_5^2\sigma =_E r_i \circ r'_i \circ B_2$, contradiction.

Now, since U_3 cannot be one of the signatures of proof of knowledge that appear in $\text{ballot}_1\tau_L, \text{ballot}_2\tau_L$ and due to the equational theory, we must have $M_3 = \text{spk}(M_3^1, M_3^2, M_3^3, M_3^4)$. Since r is not deducible, we cannot have $M_3^2\sigma =_E r \circ B'_1$, contradiction. We therefore deduce that $M_1 = \text{penc}(M_1^1, M_1^1, M_1^2)$. Moreover, $M_1^1\sigma =_E \text{pk}(sk_T)$ implies $M_1^1 = z_{\text{pk}}$ and $M_1^2\sigma =_E \text{zero}$ or one implies $M_1^2 \in \{\text{zero}, \text{one}\}$ due to the equational theory. By symmetry, M_2 has the same form as well. We therefore conclude that M is a valid ballot. \square

B.4 Proof of Theorem 1

We introduce some partial evolutions of the Helios process specification in Figure 6 and define a relation \mathcal{R} between processes in Figure 7. We clearly have that $A_n^\phi[V\{a_1/x_{\text{auth}}\}\tau_L \mid V\{a_2/x_{\text{auth}}\}\tau] \mathcal{R} A_n^\phi[V\{a_1/x_{\text{auth}}\}\tau_R \mid V\{a_2/x_{\text{auth}}\}\sigma]$. We now wish to show that $\mathcal{R} \cup \mathcal{R}^{-1}$ satisfies the three properties of Definition 2. By symmetry we focus on \mathcal{R} .

Static equivalence. We must show for all extended processes A and B , where $A \mathcal{R} B$, that $A \approx_s B$. By Lemma 4, it is sufficient to show $A^7\{N_k/y_k \mid k \in \{3, \dots, n\}\}\tau_L \approx_s A^7\{N_k/y_k \mid k \in \{3, \dots, n\}\}\tau_R$ for any N_i valid ballots. Let $\phi_7 = \nu sk_T, d, r_1, r'_1, r_2, r'_2 \cdot (\{\text{ballot}_1/x_1\} \mid \{\text{ballot}_2/x_2\} \mid \{\text{pk}(sk_T)/z_{\text{pk}}\} \mid \{(\text{partial}(sk_T, \text{tally}), \text{partial}(sk_T, \text{tally}))/x_{\text{partial}}\} \mid \{(\text{result}, \text{result}')/x_{\text{result}}\})$. We have to show $\phi_7\sigma_{N_3, \dots, N_k}\tau_L \approx_s \phi_7\sigma_{N_3, \dots, N_k}\tau_R$. By Lemma 7, we deduce that $(\text{result}, \text{result}')\sigma_{N_3, \dots, N_k}\tau_L = (\text{result}, \text{result}')\sigma_{N_3, \dots, N_k}\tau_R$ and is equal to a constant

Figure 6 Partial evolutions of the Helios process specification

$$\begin{aligned}
A^1 &= \nu sk_T, a_2, d, r_1, r'_1, y_1 \cdot (- \mid \{ballot_1/y_1\} \mid \{pk(sk_T)/z_{pk}\}) \\
A^2 &= A^1[- \mid \{ballot_1/x_1\}] \\
A^3 &= \nu sk_T, d, r_1, r'_1, r'_2, y_1, y_2 \cdot (- \mid \{ballot_1/y_1\} \mid \{ballot_2/y_2\} \mid \\
&\quad \{ballot_1/x_1\} \mid \{pk(sk_T)/z_{pk}\}) \\
A^4 &= A^3[- \mid \{ballot_2/x_2\}] \\
A^5 &= A^4[\nu y_{\text{partial}} \cdot (- \mid \\
&\quad \{\text{partial}(sk_T, \pi_1((tally, tally')), \text{partial}(sk_T, \pi_2((tally, tally'))))\} / y_{\text{partial}} \}}] \\
A^6 &= A^5[- \mid \\
&\quad \{\text{partial}(sk_T, \pi_1((tally, tally')), \text{partial}(sk_T, \pi_2((tally, tally'))))\} / x_{\text{partial}} \}}] \\
A^7 &= A^6[\{\text{dec}(\text{partial}(sk_T, tally), tally), \text{dec}(\text{partial}(sk_T, tally'), tally')\} / x_{\text{result}} \}}] \\
\\
BB_n^1 &= \bar{c}\langle y_1 \rangle \cdot BB_n^2 \\
BB_n^2 &= a_2(y_2) \cdot BB_n^3 \\
BB_n^3 &= \bar{c}\langle y_2 \rangle \cdot BB_{3,n}^3 \\
BB_{j,n} &= a_j(y_j) \cdot \text{if } \phi_{\text{sol}}\{y_j/y_{\text{ballot}}\} \text{ then} \\
&\quad \dots a_n(y_n) \cdot \text{if } \phi_{\text{sol}}\{y_n/y_{\text{ballot}}\} \text{ then} \\
&\quad BB_n^4 \\
BB_{j,n}'' &= \text{if } \phi_{\text{sol}}\{y_j/y_{\text{ballot}}\} \text{ then} \\
&\quad a_{j+1}(y_{j+1}) \cdot \text{if } \phi_{\text{sol}}\{y_{j+1}/y_{\text{ballot}}\} \text{ then} \\
&\quad \dots a_n(y_n) \cdot \text{if } \phi_{\text{sol}}\{y_n/y_{\text{ballot}}\} \text{ then} \\
&\quad BB_n^4 \\
BB_n^4 &= \bar{d}\langle (tally, tally') \rangle \cdot BB_n^5 \\
BB_n^5 &= d(y_{\text{partial}}) \cdot BB_n^6 \\
BB_n^6 &= \bar{c}\langle y_{\text{partial}} \rangle \cdot BB_n^7 \\
BB_n^7 &= \bar{c}\langle (\text{dec}(\pi_1(y_{\text{partial}}), tally), \text{dec}(\pi_2(y_{\text{partial}}), tally')) \rangle \\
\\
T_n^1 &= \bar{d}\langle (\text{partial}(sk_T, \pi_1((tally, tally')), \text{partial}(sk_T, \pi_2((tally, tally')))) \rangle)
\end{aligned}$$

(always deducible) term. Thus by Lemma 3, it is sufficient to show that $\phi_6 \sigma_{N_3, \dots, N_k} \tau_L \approx_s \phi_6 \sigma_{N_3, \dots, N_k} \tau_R$ with ϕ_6 as defined in Lemma 8. We conclude by Lemma 8.

Internal reductions. We must show for all extended processes A and B , where $A \mathcal{R} B$, that if $A \rightarrow A'$ for some A' , then $B \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' . We observe that if $A \equiv A^1[V\{a_2/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\} \mid BB_n^1 \mid T]\sigma_L$ and $B \equiv A^1[V\{a_2/x_{\text{auth}}, m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\} \mid BB_n^1 \mid T]\sigma_R$ – that is, $A \mathcal{R} B$ by (R2) – then there is no extended process A' such that $A \rightarrow A'$; similarly, for (R4), (R5), (R7), (R10), (R11) and (R12). We proceed by case analysis on the remaining cases.

(R1) We have $A \equiv A_n^{\phi_{\text{sol}}}[V\{a_1/x_{\text{auth}}, m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\} \mid V\{a_2/x_{\text{auth}}, m_2/x_{\text{vote}},$

Figure 7 Definition of the relation \mathcal{R}

Given integer $n \geq 2$, for all $3 \leq j \leq n$, terms M_j , terms N_j such that their normal form are valid ballots, and distinct variables $x_{\text{partial}}, x_{\text{result}}, x_1, x_2$ such that $\text{fv}(M_j) \cup \text{fv}(N_j) \subseteq \text{dom}(A^4)$ and $(\text{fn}(M_j) \cup \text{fn}(N_j)) \cap \text{bn}(A^4) = \emptyset$, consider the smallest relation \mathcal{R} which is closed under structural equivalence and includes the following pairs of extended processes.

$$\begin{aligned} & A_n^{\phi_{\text{sol}}} [V\{a_1/x_{\text{auth}}, m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\} \mid V\{a_2/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\}], \\ & A_n^{\phi_{\text{sol}}} [V\{a_1/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\} \mid V\{a_2/x_{\text{auth}}, m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\}] \end{aligned} \quad (\text{R1})$$

$$\begin{aligned} & A^1 [V\{a_2/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\} \mid BB_n^1 \mid T] \sigma_L, \\ & A^1 [V\{a_2/x_{\text{auth}}, m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\} \mid BB_n^1 \mid T] \sigma_R \end{aligned} \quad (\text{R2})$$

$$\begin{aligned} & A^2 [V\{a_2/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\} \mid BB_n^2 \mid T] \sigma_L, \\ & A^2 [V\{a_2/x_{\text{auth}}, m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\} \mid BB_n^2 \mid T] \sigma_R \end{aligned} \quad (\text{R3})$$

$$A^3 [BB_n^3 \mid T] \tau_L, \quad A^3 [BB_n^3 \mid T] \tau_R \quad (\text{R4})$$

$$\begin{aligned} & A^4 [BB'_{j,n} \{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\} \mid T] \tau_L, \\ & A^4 [BB'_{j,n} \{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\} \mid T] \tau_R \end{aligned} \quad (\text{R5})$$

$$\begin{aligned} & A^4 [BB''_{j,n} \{N_k/y_k \mid k \in \{3, \dots, j-1\}\} \{M_j/y_j\} \mid T] \tau_L, \\ & A^4 [BB''_{j,n} \{N_k/y_k \mid k \in \{3, \dots, j-1\}\} \{M_j/y_j\} \mid T] \tau_R \end{aligned} \quad (\text{R6})$$

$$A^4 [0 \mid T] \tau_L, \quad A^4 [0 \mid T] \tau_R \quad (\text{R7})$$

$$\begin{aligned} & A^4 [BB_n^4 \{N_k/y_k \mid k \in \{3, \dots, n\}\} \mid T] \tau_L, \\ & A^4 [BB_n^4 \{N_k/y_k \mid k \in \{3, \dots, n\}\} \mid T] \tau_R \end{aligned} \quad (\text{R8})$$

$$\begin{aligned} & A^4 [BB_n^5 \{N_k/y_k \mid k \in \{3, \dots, n\}\} \mid T_n^1] \tau_L, \\ & A^4 [BB_n^5 \{N_k/y_k \mid k \in \{3, \dots, n\}\} \mid T_n^1] \tau_R \end{aligned} \quad (\text{R9})$$

$$\begin{aligned} & A^5 [BB_n^6 \{N_k/y_k \mid k \in \{3, \dots, n\}\} \tau_L, \\ & A^5 [BB_n^6 \{N_k/y_k \mid k \in \{3, \dots, n\}\} \tau_R \end{aligned} \quad (\text{R10})$$

$$\begin{aligned} & A^6 [BB_n^7 \{N_k/y_k \mid k \in \{3, \dots, n\}\} \tau_L, \\ & A^6 [BB_n^7 \{N_k/y_k \mid k \in \{3, \dots, n\}\} \tau_R \end{aligned} \quad (\text{R11})$$

$$A^7 \{N_k/y_k \mid k \in \{3, \dots, n\}\} \tau_L, \quad A^7 \{N_k/y_k \mid k \in \{3, \dots, n\}\} \tau_R \quad (\text{R12})$$

$m'_2/x'_{\text{vote}}\}$] and $B \equiv A_n^{\phi_{\text{sol}}}[V\{a_1/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\} \mid V\{a_2/x_{\text{auth}}, m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\}]$. If $A \rightarrow A'$, then it must be the case that $A \equiv C[\bar{a}_1\langle y_1 \rangle.0 \mid a_1(y_1).BB_n^1]\sigma_L$ and $A' \equiv C[0 \mid BB_n^1]\sigma_L$, where $C[-] = A^1[\nu a_1.(- \mid V\{a_2/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\} \mid T)]$. It follows from $B \equiv C[\bar{a}_1\langle y_1 \rangle.0 \mid a_1(y_1).BB_n^1]\sigma_R$, that $B \rightarrow B'$, where $B' = A^1[V\{a_2/x_{\text{auth}}, m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\} \mid BB_n^1 \mid T]\sigma_R$. Since $A^1[V\{a_2/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\} \mid BB_n^1 \mid T]\sigma_L \mathcal{R} B'$ and $A' \equiv A^1[V\{a_2/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\} \mid BB_n^1 \mid T]\sigma_L$, we derive $A' \mathcal{R} B'$ by the closure of \mathcal{R} under structural equivalence.

(R3) This case is similar to (R1). We have $A \equiv A^2[V\{a_2/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\} \mid BB_n^2 \mid T]\sigma_L$ and $B \equiv A^2[V\{a_2/x_{\text{auth}}, m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\} \mid BB_n^2 \mid T]\sigma_R$. If $A \rightarrow A'$, then it must be the case that $A \equiv C[\bar{a}_2\langle y_2 \rangle.0 \mid a_2(y_2).BB_n^3]\tau_L$ and $A' \equiv C[0 \mid BB_n^3]\tau_L$, where $C[-] = A^3[\nu a_2.(- \mid T)]$. It follows from $B \equiv C[\bar{a}_2\langle y_2 \rangle.0 \mid a_2(y_2).BB_n^3]\tau_R$, that $B \rightarrow B'$, where $B' = A^3[BB_n^3 \mid T]\tau_R$. Since $A^3[BB_n^3 \mid T]\tau_L \mathcal{R} B'$ and $A' \equiv A^3[BB_n^3 \mid T]\tau_L$, we derive $A' \mathcal{R} B'$ by the closure of \mathcal{R} under structural equivalence.

(R6) We have $A \equiv A^4[BB''_{j,n}\{N_k/y_k \mid k \in \{3, \dots, j-1\}\}\{M_j/y_j\} \mid T]\tau_L$ and $B \equiv A^4[BB''_{j,n}\{N_k/y_k \mid k \in \{3, \dots, j-1\}\}\{M_j/y_j\} \mid T]\tau_R$ for some integer $j \in \{3, \dots, n\}$, valid ballots N_3, \dots, N_{j-1} and term M_j such that $\text{fv}(M_j) \cup \bigcup_{3 \leq i \leq j-1} \text{fv}(N_i) \subseteq \text{dom}(A^4)$ and $(\text{fn}(M_j) \cup \bigcup_{3 \leq i \leq j-1} \text{fn}(N_i)) \cap \text{bn}(A^4) = \emptyset$. If $A \rightarrow A'$, then it must be the case that $A \equiv C[\text{if } \phi_{\text{sol}}\{M_j/y_{\text{ballot}}, \text{ballot}_1/y_1, \text{ballot}_2/y_2\} \text{ then } P \text{ else } 0]\tau_L$, where $C[-] = A^4[- \mid T]$, and if $j < n$, then $P = BB'_{j+1,n}\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\}\{M_j/y_j\}$, otherwise $P = BB_n^4\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\}\{M_j/y_j\}$. We also have $B \equiv C[\text{if } \phi_{\text{sol}}\{M_j/y_{\text{ballot}}, \text{ballot}_1/y_1, \text{ballot}_2/y_2\} \text{ then } P \text{ else } 0]\tau_R$. We have $\text{ballot}_1\tau_R$ is syntactically equal to $x_1\sigma'$ and $\text{ballot}_2\tau_R$ is syntactically equal to $x_2\sigma'$, where $\sigma' = \{\text{ballot}_1\tau_R/x_1, \text{ballot}_2\tau_R/x_2, \text{pk}(sk_T)/z_{\text{pk}}\}$; it follows from $\varphi(C[\text{if } \phi_{\text{sol}}\{M_j/y_{\text{ballot}}, \text{ballot}_1/y_1, \text{ballot}_2/y_2\} \text{ then } P \text{ else } 0]\tau_R) = \nu sk_T, d, r_1, r'_1, r_2, r'_2, y_1, y_2.\sigma'$ that $B \equiv C[\text{if } \phi_{\text{sol}}\{M_j/y_{\text{ballot}}, x_1/y_1, x_2/y_2\} \text{ then } P \text{ else } 0]\tau_R$. We proceed by case analysis on the structure of A' :

- If $A' \equiv C[P]\tau_L$, then by closure of internal reduction under structural equivalence we have $C[\text{if } \phi_{\text{sol}}\{M_j/y_{\text{ballot}}, x_1/y_1, x_2/y_2\} \text{ then } P \text{ else } 0]\tau_L \rightarrow C[P]\tau_L$ because $\text{ballot}_1\tau_L$ is syntactically equal to $x_1\sigma$ and $\text{ballot}_2\tau_L$ is syntactically equal to $x_2\sigma$, where $\sigma = \{\text{ballot}_1\tau_L/x_1, \text{ballot}_2\tau_L/x_2, \text{pk}(sk_T)/z_{\text{pk}}\}$ and $\varphi(C[\text{if } \phi_{\text{sol}}\{M_j/y_{\text{ballot}}, x_1/y_1, x_2/y_2\} \text{ then } P \text{ else } 0]\tau_L) = \nu sk_T, d, r_1, r'_1, r_2, r'_2, y_1, y_2.\sigma$. By Corollary 1 we have $B \rightarrow B' = A^4[P \mid T]\tau_R$, and it remains to show $A' \mathcal{R} B'$. We know $\llbracket \phi_{\text{sol}}\{M_j/y_{\text{ballot}}, x_1/y_1, x_2/y_2\}\sigma \rrbracket = \text{true}$ and $\llbracket \phi_{\text{sol}}\{M_j/y_{\text{ballot}}, x_1/y_1, x_2/y_2\}\sigma' \rrbracket = \text{true}$; it follows, for $\tau \in \{\tau_L, \tau_R\}$, that $\llbracket \phi_{\text{sol}}\{M_j/y_{\text{ballot}}\}\{\text{pk}(sk_T)/z_{\text{pk}}, \text{ballot}_1/y_1, \text{ballot}_2/y_2\}\tau \rrbracket = \text{true}$ and by Lemma 9 we know that M_j in normal form is a valid ballot. We continue by case analysis on the structure of P :

1. If $P = BB'_{j+1,n}\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\}\{M_j/y_j\}$, then we have $j < n$. Let $j' = j + 1$ and $N_j = M_j$, observe $P =$

$BB'_{j',n}\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j\}\}$ and $A^4[BB'_{j',n}\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j\}\} \mid T]_{\tau_L} \mathcal{R} B'$. The result $A' \mathcal{R} B'$ follows by closure of \mathcal{R} under structural equivalence.

2. If $P = BB_n^4\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\}\{M_j/y_j\}$, then it must be the case that $j = n$. Let $N_j = M_j$ and hence $P = BB_n^4\{N_k/y_k \mid k \in \{3, \dots, n\}\}$. Since $A^4[BB_n^4\{N_k/y_k \mid k \in \{3, \dots, n\}\} \mid T]_{\tau_L} \mathcal{R} B'$ and $A' \equiv A^4[BB_n^4\{N_k/y_k \mid k \in \{3, \dots, n\}\} \mid T]_{\tau_L}$, we derive $A' \mathcal{R} B'$ by the closure of \mathcal{R} under structural equivalence.

– $A' \equiv C[0]_{\tau_L}$, then similarly to above we have $C[\text{if } \phi_{\text{sol}}\{M_j/y_{\text{ballot}}, x_1/y_1, x_2/y_2\} \text{ then } P \text{ else } 0]_{\tau_L} \rightarrow C[0]_{\tau_L}$ and it follows by Corollary 1 that $B \rightarrow B' = C[0]_{\tau_R}$. Since $A^4[0 \mid T] \mathcal{R} B'$ and $A' \equiv A^4[0 \mid T]$, we derive $A' \mathcal{R} B'$ by the closure of \mathcal{R} under structural equivalence.

(R8) We have $A \equiv A^4[BB_n^4\{N_k/y_k \mid k \in \{3, \dots, n\}\} \mid T]_{\tau_L}$ and $B \equiv A^4[BB_n^4\{N_k/y_k \mid k \in \{3, \dots, n\}\} \mid T]_{\tau_R}$, where $BB_n^4 = \bar{d}(\langle \pi_1(y_1) * \dots * \pi_1(y_n), \pi_2(y_1) * \dots * \pi_2(y_n) \rangle) \cdot BB_n^5$ and $T = d(y_{\text{tally}}) \cdot \bar{d}(\langle \text{partial}(sk_T, \pi_1(y_{\text{tally}})), \text{partial}(sk_T, \pi_2(y_{\text{tally}})) \rangle)$. If $A \rightarrow A'$, then it must be the case that $A' \equiv A^4[BB_n^5\{N_k/y_k \mid k \in \{3, \dots, n\}\} \mid T_n^1]_{\tau_L}$. It follows immediately that $B \rightarrow B'$, where $B' = A^4[BB_n^5\{N_k/y_k \mid k \in \{3, \dots, n\}\} \mid T_n^1]_{\tau_R}$. We derive $A' \mathcal{R} B'$ by the closure of \mathcal{R} under structural equivalence.

(R9) We have $A \equiv A^4[BB_n^5\{N_k/y_k \mid k \in \{3, \dots, n\}\} \mid T_n^1]_{\tau_L}$ and $B \equiv A^4[BB_n^5\{N_k/y_k \mid k \in \{3, \dots, n\}\} \mid T_n^1]_{\tau_R}$. If $A \rightarrow A'$, then it must be the case that $A \equiv A^5[\bar{d}(y_{\text{partial}}) \mid d(y_{\text{partial}}) \cdot BB_n^6]\{N_k/y_k \mid k \in \{3, \dots, n\}\}_{\tau_L}$ and $A' \equiv A^5[0 \mid BB_n^6]\{N_k/y_k \mid k \in \{3, \dots, n\}\}_{\tau_L}$. It follows from $B \equiv A^5[d(y_{\text{partial}}) \cdot BB_n^6 \mid \bar{d}(y_{\text{partial}})]\{N_k/y_k \mid k \in \{3, \dots, n\}\}_{\tau_R}$ that $B \rightarrow B'$, where $B' = A^5[BB_n^6]\{N_k/y_k \mid k \in \{3, \dots, n\}\}_{\tau_R}$. We derive $A' \mathcal{R} B'$ by the closure of \mathcal{R} under structural equivalence.

Labelled reductions. We must show for all extended processes A and B , where $A \mathcal{R} B$, that if $A \xrightarrow{\alpha} A'$ such that $\text{fv}(\alpha) \subseteq \text{dom}(A)$ and $\text{bn}(\alpha) \cap \text{fn}(B) = \emptyset$, then $B \rightarrow^* \overset{\alpha}{\rightarrow} B'$ and $A' \mathcal{R} B'$ for some B' . We observe cases (R1), (R3), (R6), (R7), (R8), (R9) and (R12) cannot be reduced by labelled reductions and proceed by case analysis on the remaining cases.

(R2) We have $A \equiv A^1[V\{a_2/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\} \mid BB_n^1 \mid T]_{\sigma_L}$ and $B \equiv A^1[V\{a_2/x_{\text{auth}}, m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\} \mid BB_n^1 \mid T]_{\sigma_R}$. If $A \xrightarrow{\alpha} A'$ such that $\text{fv}(\alpha) \subseteq \text{dom}(A)$ and $\text{bn}(\alpha) \cap \text{fn}(B) = \emptyset$, then it must be the case that $A \equiv A^1[V\{a_2/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\} \mid \bar{c}(y_1) \cdot BB_n^2 \mid T]_{\sigma_L}$ and $A' \equiv A^2[V\{a_2/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\} \mid BB_n^2 \mid T]_{\sigma_L}$ for some variable x_1 where $\alpha = \nu x_1. \bar{c}(x_1)$ and $x_1 \neq z_{\text{pk}}$. It follows from $B \equiv A^1[V\{a_2/x_{\text{auth}}, m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\} \mid \bar{c}(y_1) \cdot BB_n^2 \mid T]_{\sigma_R}$, that $B \xrightarrow{\alpha} B'$ where $B' = A^2[V\{a_2/x_{\text{auth}}, m_1/x_{\text{vote}}, m'_1/x'_{\text{vote}}\} \mid BB_n^2 \mid T]_{\sigma_R}$. We have $A^2[V\{a_2/x_{\text{auth}}, m_2/x_{\text{vote}}, m'_2/x'_{\text{vote}}\} \mid BB_n^2 \mid T]_{\sigma_L} \mathcal{R} B'$ and by closure of \mathcal{R} under structural equivalence $A' \mathcal{R} B'$.

- (R4) We have $A \equiv A^3[BB_n^3 \mid T]\tau_L$ and $B \equiv A^3[BB_n^3 \mid T]\tau_R$. If $A \xrightarrow{\alpha} A'$ such that $\text{fv}(\alpha) \subseteq \text{dom}(A)$ and $\text{bn}(\alpha) \cap \text{fn}(B) = \emptyset$, then it must be the case that $A \equiv A^3[\bar{c}\langle y_2 \rangle . BB'_{3,n} \mid T]\tau_L$ and $A' \equiv A^4[BB'_{3,n} \mid T]\tau_L$ for some variable x_2 , where $\alpha = \nu x_2 . \bar{c}\langle x_2 \rangle$ and $x_2 \notin \{x_1, z_{\text{pk}}\}$. It follows from $B \equiv A^3[\bar{c}\langle y_2 \rangle . BB'_{3,n} \mid T]\tau_R$, that $B \xrightarrow{\alpha} B'$, where $B' = A^4[BB'_{3,n} \mid T]\tau_R$. Since $A^4[BB'_{3,n} \mid T]\tau_L = A^4[BB'_{j,n}\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\} \mid T]\tau_L$ and $A^4[BB'_{3,n} \mid T]\tau_R = A^4[BB'_{j,n}\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\} \mid T]\tau_R$ when $j = 3$, we have $A^4[BB'_{3,n} \mid T]\tau_L \mathcal{R} B'$ and derive $A' \mathcal{R} B'$ by closure of \mathcal{R} under structural equivalence $A' \mathcal{R} B'$.
- (R5) We have $A \equiv A^4[BB'_{j,n}\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\} \mid T]\tau_L$ and $B \equiv A^4[BB'_{j,n}\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\} \mid T]\tau_R$ for some integer $j \in \{3, \dots, n\}$ and terms N_3, \dots, N_{j-1} , where $\bigcup_{3 \leq i \leq j-1} \text{fv}(N_i) \subseteq \text{dom}(A^4)$ and $\text{bn}(A^4) \cap \bigcup_{3 \leq i \leq j-1} \text{fn}(N_i) = \emptyset$. If $A \xrightarrow{\alpha} A'$ such that $\text{fv}(\alpha) \subseteq \text{dom}(A)$ and $\text{bn}(\alpha) \cap \text{fn}(B) = \emptyset$, then it must be the case that $A \equiv A^4[a_j(y_j) . BB''_{j,n}\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\} \mid T]\tau_L$ and $A' \equiv A^4[BB''_{j,n}\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\}\{M_j/y_j\} \mid T]\tau_L$, where $\alpha = c(M_j)$ for some term M_j such that $\text{fn}(\alpha) \cap \text{bn}(A^4) = \emptyset$. It follows from $B \equiv A^4[a_j(y_j) . BB''_{j,n}\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\} \mid T]\tau_R$, that $B \xrightarrow{\alpha} B'$, where $B' = A^4[BB''_{j,n}\{N_k/y_k \mid j > 3 \wedge k \in \{3, \dots, j-1\}\}\{M_j/y_j\} \mid T]\tau_R$. We have $A^4[BB''_{j,n}\{N_k/y_k \mid k \in \{3, \dots, j-1\}\}\{M_j/y_j\} \mid T]\tau_L \mathcal{R} B'$, and derive $A' \mathcal{R} B'$ by closure of \mathcal{R} under structural equivalence.
- (R10) We have $A \equiv A^5[BB_n^6]\{N_k/y_k \mid k \in \{3, \dots, n\}\}\tau_L$ and $B \equiv A^5[BB_n^6]\{N_k/y_k \mid k \in \{3, \dots, n\}\}\tau_R$. If $A \xrightarrow{\alpha} A'$ such that $\text{fv}(\alpha) \subseteq \text{dom}(A)$ and $\text{bn}(\alpha) \cap \text{fn}(B) = \emptyset$, then it must be the case that $A' \equiv A^6[BB_n^7]\{N_k/y_k \mid k \in \{3, \dots, n\}\}\tau_L$ for some variable x_{partial} , where $\alpha = \nu x_{\text{partial}} . \bar{c}\langle x_{\text{partial}} \rangle$ and $x_{\text{partial}} \notin \{x_1, x_2, z_{\text{pk}}\}$. It follows immediately that $B \xrightarrow{\alpha} B'$, where $B' = A^6[BB_n^7]\{N_k/y_k \mid k \in \{3, \dots, n\}\}\tau_R$. We have $A^6[BB_n^7]\{N_k/y_k \mid k \in \{3, \dots, n\}\}\tau_L \mathcal{R} B'$ and by closure of \mathcal{R} under structural equivalence $A' \mathcal{R} B'$.
- (R11) This case is similar to (R10). We have $A \equiv A^6[BB_n^7]\{N_k/y_k \mid k \in \{3, \dots, n\}\}\tau_L$ and $B \equiv A^6[BB_n^7]\{N_k/y_k \mid k \in \{3, \dots, n\}\}\tau_R$. If $A \xrightarrow{\alpha} A'$ such that $\text{fv}(\alpha) \subseteq \text{dom}(A)$ and $\text{bn}(\alpha) \cap \text{fn}(B) = \emptyset$, then it must be the case that $A' \equiv A^7\{N_k/y_k \mid k \in \{3, \dots, n\}\}\tau_L$ for some variable x_{result} , where $\alpha = \nu x_{\text{result}} . \bar{c}\langle x_{\text{result}} \rangle$ and $x_{\text{result}} \notin \{x_1, x_2, x_{\text{partial}}, z_{\text{pk}}\}$. It follows immediately that $B \xrightarrow{\alpha} B'$, where $B' = A^7\{N_k/y_k \mid k \in \{3, \dots, n\}\}\tau_R$. We have $A^7\{N_k/y_k \mid k \in \{3, \dots, n\}\}\tau_L \mathcal{R} B'$ and by closure of \mathcal{R} under structural equivalence $A' \mathcal{R} B'$.

References

- [Adi06] Ben Adida. *Advances in Cryptographic Voting Systems*. PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 2006.
- [Adi08] Ben Adida. Helios: Web-based Open-Audit Voting. In *USENIX Security'08: 17th USENIX Security Symposium*, pages 335–348. USENIX Association, 2008.
- [Adi10] Ben Adida. Attacks and Defenses. Helios documentation, <http://documentation.heliosvoting.org/attacks-and-defenses>, 2010.
- [AF01] Martín Abadi and Cédric Fournet. Mobile values, new names, and secure communication. In *POPL'01: 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 104–115. ACM Press, 2001.
- [AMPQ09] Ben Adida, Olivier de Marneffe, Olivier Pereira, and Jean-Jacques Quisquater. Electing a University President Using Open-Audit Voting: Analysis of Real-World Use of Helios. In *EVT/WOTE'09: Electronic Voting Technology Workshop/Workshop on Trustworthy Elections*. USENIX Association, 2009.
- [AP10] Ben Adida and Olivier Pereira. Private email communication, November 2010.
- [AR00] Martín Abadi and Phillip Rogaway. Reconciling Two Views of Cryptography (The Computational Soundness of Formal Encryption). In *IFIP TCS'00: 1st International Conference on Theoretical Computer Science*, volume 1872 of *LNCS*, pages 3–22. Springer, 2000.
- [AR02] Martín Abadi and Phillip Rogaway. Reconciling Two Views of Cryptography (The Computational Soundness of Formal Encryption). *Journal of Cryptology*, 15(2):103–127, 2002.
- [BAF08] Bruno Blanchet, Martín Abadi, and Cédric Fournet. Automated verification of selected equivalences for security protocols. *Journal of Logic and Algebraic Programming*, 75(1):3–51, February–March 2008.
- [BCP⁺11] David Bernhard, Véronique Cortier, Olivier Pereira, Ben Smyth, and Bogdan Warinschi. Adapting Helios for provable ballot secrecy. In *ESORICS'11: 16th European Symposium on Research in Computer Security*, LNCS. Springer, 2011. To appear.
- [Ben06] Josh Benaloh. Simple Verifiable Elections. In *EVT'06: Electronic Voting Technology Workshop*. USENIX Association, 2006.

- [Ben07] Josh Benaloh. Ballot Casting Assurance via Voter-Initiated Poll Station Auditing. In *EVT'07: Electronic Voting Technology Workshop*. USENIX Association, 2007.
- [BHM08] Michael Backes, Cătălin Hrițcu, and Matteo Maffei. Automated Verification of Remote Electronic Voting Protocols in the Applied Pi-calculus. In *CSF'08: 21st Computer Security Foundations Symposium*, pages 195–209. IEEE Computer Society, 2008.
- [BVQ10] Josh Benaloh, Serge Vaudenay, and Jean-Jacques Quisquater. Final Report of IACR Electronic Voting Committee. International Association for Cryptologic Research. http://www.iacr.org/elections/eVoting/finalReportHelios_2010-09-27.html, Sept 2010.
- [BY86] Josh Cohen Benaloh and Moti Yung. Distributing the Power of a Government to Enhance the Privacy of Voters. In *PODC'86: 5th Principles of Distributed Computing Symposium*, pages 52–62. ACM Press, 1986.
- [CCM07] Michael R. Clarkson, Stephen Chong, and Andrew C. Myers. Civitas: Toward a Secure Voting System. Technical Report 2007-2081, Cornell University, May 2007. Revised March 2008.
- [CCM08] Michael R. Clarkson, Stephen Chong, and Andrew C. Myers. Civitas: Toward a Secure Voting System. In *SEP'08: 29th Security and Privacy Symposium*, pages 354–368. IEEE Computer Society, 2008.
- [CDS94] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of Partial Knowledge and Simplified Design of Witness Hiding Protocols. In *CRYPTO'94: 14th International Cryptology Conference*, volume 839 of *LNCS*, pages 174–187. Springer, 1994.
- [CEG88] David Chaum, Jan-Hendrik Evertse, and Jeroen van de Graaf. An Improved Protocol for Demonstrating Possession of Discrete Logarithms and Some Generalizations. In *EUROCRYPT'87: 4th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 304 of *LNCS*, pages 127–141. Springer, 1988.
- [CEGP87] David Chaum, Jan-Hendrik Evertse, Jeroen van de Graaf, and René Peralta. Demonstrating Possession of a Discrete Logarithm Without Revealing It. In *CRYPTO'86: 6th International Cryptology Conference*, volume 263 of *LNCS*, pages 200–212. Springer, 1987.
- [CFSY96] Ronald Cramer, Matthew K. Franklin, Berry Schoenmakers, and Moti Yung. Multi-Authority Secret-Ballot Elections with Linear Work. In *EUROCRYPT'96: 15th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 1070 of *LNCS*, pages 72–83. Springer, 1996.

- [CGMA85] Benny Chor, Shafi Goldwasser, Silvio Micali, and Baruch Awerbuch. Verifiable Secret Sharing and Achieving Simultaneity in the Presence of Faults. In *FOCS'85: 26th Foundations of Computer Science Symposium*, pages 383–395. IEEE Computer Society, 1985.
- [CGS97] Ronald Cramer, Rosario Gennaro, and Berry Schoenmakers. A Secure and Optimally Efficient Multi-Authority Election Scheme. In *EUROCRYPT'97: 16th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 1233 of *LNCS*, pages 103–118. Springer, 1997.
- [CP93] David Chaum and Torben P. Pedersen. Wallet Databases with Observers. In *CRYPTO'92: 12th International Cryptology Conference*, volume 740 of *LNCS*, pages 89–105. Springer, 1993.
- [CR87] Benny Chor and Michael O. Rabin. Achieving Independence in Logarithmic Number of Rounds. In *PODC'87: 6th Principles of Distributed Computing Symposium*, pages 260–268. ACM Press, 1987.
- [CRS05] David Chaum, Peter Y. A. Ryan, and Steve Schneider. A Practical Voter-Verifiable Election Scheme. In *ESORICS'05: 10th European Symposium On Research In Computer Security*, volume 3679 of *LNCS*, pages 118–139. Springer, 2005.
- [CS11] Véronique Cortier and Ben Smyth. Attacking and fixing Helios: An analysis of ballot secrecy. In *CSF'11: 24th Computer Security Foundations Symposium*. IEEE Computer Society, 2011. To appear.
- [Dag07] Participants of the Dagstuhl Conference on Frontiers of E-Voting. *Dagstuhl Accord*, 2007. <http://www.dagstuhlaccord.org/>.
- [DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-Malleable Cryptography. In *STOC'91: 23rd Theory of computing Symposium*, pages 542–552. ACM Press, 1991.
- [DDN00] Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable Cryptography. *Journal on Computing*, 30(2):391–437, 2000.
- [DKR06] Stéphanie Delaune, Steve Kremer, and Mark Ryan. Coercion-Resistance and Receipt-Freeness in Electronic Voting. In *CSFW'06: 19th Computer Security Foundations Workshop*, pages 28–42. IEEE Computer Society, 2006.
- [DKR09] Stéphanie Delaune, Steve Kremer, and Mark D. Ryan. Verifying privacy-type properties of electronic voting protocols. *Journal of Computer Security*, 17(4):435–487, July 2009.
- [DKR10] Stéphanie Delaune, Steve Kremer, and Mark D. Ryan. Verifying Privacy-Type Properties of Electronic Voting Protocols: A Taster. In David Chaum, Markus Jakobsson, Ronald L. Rivest, and Peter Y.

- A. Ryan, editors, *Towards Trustworthy Elections: New Directions in Electronic Voting*, volume 6000 of *LNCS*, pages 289–309. Springer, 2010.
- [DRS08] Stéphanie Delaune, Mark D. Ryan, and Ben Smyth. Automatic verification of privacy properties in the applied pi-calculus. In *IFIPTM'08: 2nd Joint iTrust and PST Conferences on Privacy, Trust Management and Security*, volume 263 of *International Federation for Information Processing (IFIP)*, pages 263–278. Springer, 2008.
- [ED10] Saghar Estehghari and Yvo Desmedt. Exploiting the Client Vulnerabilities in Internet E-voting Systems: Hacking Helios 2.0 as an Example. In *EVT/WOTE'10: Electronic Voting Technology Workshop/Workshop on Trustworthy Elections*. USENIX Association, 2010.
- [ElG85] Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4):469–472, 1985.
- [Est07] Est Républicain. June, 18th 2007. Meurthe-et-Moselle edition (Daily French Newspaper).
- [FOO92] Atsushi Fujioka, Tatsuaki Okamoto, and Kazuo Ohta. A Practical Secret Voting Scheme for Large Scale Elections. In *AUSCRYPT'92: Workshop on the Theory and Application of Cryptographic Techniques*, volume 718 of *LNCS*, pages 244–251. Springer, 1992.
- [Fr] *Article L65 of the French electoral code*. <http://www.legifrance.gouv.fr/>.
- [Fr10] *Résultat par bureau du premier tour des élections régionales*, 2010. http://www.monaulnay.com/wp-content/uploads/2010/03/resultat_regionale_par_bureau.pdf.
- [Gen95] Rosario Gennaro. Achieving independence efficiently and securely. In *PODC'95: 14th Principles of Distributed Computing Symposium*, pages 130–136. ACM Press, 1995.
- [Gen00] Rosario Gennaro. A Protocol to Achieve Independence in Constant Rounds. *IEEE Transactions on Parallel and Distributed Systems*, 11(7):636–647, 2000.
- [HBH10] Stuart Haber, Josh Benaloh, and Shai Halevi. The Helios e-Voting Demo for the IACR. International Association for Cryptologic Research. <http://www.iacr.org/elections/eVoting/heliosDemo.pdf>, May 2010.

- [JCJ02] Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-Resistant Electronic Elections. *Cryptology ePrint Archive*, Report 2002/165, 2002.
- [JCJ05] Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-Resistant Electronic Elections. In *WPES'05: 4th Workshop on Privacy in the Electronic Society*, pages 61–70. ACM Press, 2005. See also <http://www.rsa.com/rsalabs/node.asp?id=2860>.
- [JCJ10] Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-Resistant Electronic Elections. In David Chaum, Markus Jakobsson, Ronald L. Rivest, and Peter Y. A. Ryan, editors, *Towards Trustworthy Elections: New Directions in Electronic Voting*, volume 6000 of *LNCS*, pages 37–63. Springer, 2010.
- [KR05] Steve Kremer and Mark D. Ryan. Analysis of an Electronic Voting Protocol in the Applied Pi Calculus. In *ESOP'05: 14th European Symposium on Programming*, volume 3444 of *LNCS*, pages 186–200. Springer, 2005.
- [KRS10] Steve Kremer, Mark D. Ryan, and Ben Smyth. Election verifiability in electronic voting protocols. In *ESORICS'10: 15th European Symposium on Research in Computer Security*, volume 6345 of *LNCS*, pages 389–404. Springer, 2010.
- [KT09] Ralf Küsters and Tomasz Truderung. An Epistemic Approach to Coercion-Resistance for Electronic Voting Protocols. In *S&P'09: 30th IEEE Symposium on Security and Privacy*, pages 251–266. IEEE Computer Society, 2009.
- [Lan10] Lucie Langer. *Privacy and Verifiability in Electronic Voting*. PhD thesis, Fachbereich Informatik, Technischen Universität Darmstadt, 2010.
- [LBD⁺04] Byoungcheon Lee, Colin Boyd, Ed Dawson, Kwangjo Kim, Jeongmo Yang, and Seungjae Yoo. Providing Receipt-Freeness in Mixnet-Based Voting Protocols. In *ICISC'03: 6th International Conference on Information Security and Cryptology*, volume 2971 of *LNCS*, pages 245–258. Springer, 2004.
- [LCM08] Adrian Leung, Liqun Chen, and Chris J. Mitchell. On a Possible Privacy Flaw in Direct Anonymous Attestation (DAA). In *Trust'08: 1st International Conference on Trusted Computing and Trust in Information Technologies*, number 4968 in *LNCS*, pages 179–190. Springer, 2008.
- [LL90] Arjen K. Lenstra and Hendrik W. Lenstra Jr. Algorithms in Number Theory. In Jan van Leeuwen, editor, *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity*, chapter 12, pages 673–716. MIT Press, 1990.

- [LSB⁺10] Lucie Langer, Axel Schmidt, Johannes Buchmann, Melanie Volkamer, and Alexander Stolfik. Towards a Framework on the Security Requirements for Electronic Voting Protocols. In *Re-Vote'09: First International Workshop on Requirements Engineering for E-Voting Systems*, pages 61–68. IEEE Computer Society, 2010.
- [LSBV10] Lucie Langer, Axel Schmidt, Johannes Buchmann, and Melanie Volkamer. A Taxonomy Refining the Security Requirements for Electronic Voting: Analyzing Helios as a Proof of Concept. In *ARES'10: 5th International Conference on Availability, Reliability and Security*, pages 475–480. IEEE Computer Society, 2010.
- [Oka98] Tatsuaki Okamoto. Receipt-Free Electronic Voting Schemes for Large Scale Elections. In *SP'97: 5th International Workshop on Security Protocols*, volume 1361 of *LNCS*, pages 25–35. Springer, 1998.
- [PAM10] Olivier Pereira, Ben Adida, and Olivier de Marneffe. Bringing open audit elections into practice: Real world uses of helios. Swiss e-voting workshop, https://www.e-voting-cc.ch/images/sevot10/slides/helios_20100906.pdf. See also <http://www.uclouvain.be/crypto/electionmonitor/>, 2010.
- [Ped91] Torben P. Pedersen. A Threshold Cryptosystem without a Trusted Party. In *EUROCRYPT'91: 10th International Conference on the Theory and Applications of Cryptographic Techniques*, number 547 in *LNCS*, pages 522–526. Springer, 1991.
- [Pfi94] Birgit Pfitzmann. Breaking Efficient Anonymous Channel. In *EUROCRYPT'94: 11th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 950 of *LNCS*, pages 332–340. Springer, 1994.
- [Pri10] Princeton University. *Princeton Election Server*, 2010. <https://princeton-helios.appspot.com/>.
- [RS98] Peter Y. A. Ryan and Steve A. Schneider. An Attack on a Recursive Authentication Protocol. A Cautionary Tale. *Information Processing Letters*, 65(1):7–10, 1998.
- [RS11] Mark D. Ryan and Ben Smyth. Applied pi calculus. In Véronique Cortier and Steve Kremer, editors, *Formal Models and Techniques for Analyzing Security Protocols*, chapter 6. IOS Press, 2011.
- [Rud07] Carsten Rudolph. Covert Identity Information in Direct Anonymous Attestation (DAA). In *SEC'07: 22nd International Information Security Conference*, volume 232 of *International Federation for Information Processing (IFIP)*, pages 443–448. Springer, 2007.

- [SC10] Ben Smyth and Véronique Cortier. Attacking ballot secrecy in Helios. YouTube video, linked from <http://www.bensmyth.com/publications/10-attacking-helios/>, 2010.
- [SC11] Ben Smyth and Véronique Cortier. Replay attacks that violate privacy in electronic voting schemes. Technical report, Loria, CNRS & INRIA Nancy Grand Est, June 2011.
- [Sch90] Claus-Peter Schnorr. Efficient Identification and Signatures for Smart Cards. In *CRYPTO'89: 9th International Cryptology Conference*, volume 435 of *LNCS*, pages 239–252. Springer, 1990.
- [Sch99] Berry Schoenmakers. A simple publicly verifiable secret sharing scheme and its application to electronic voting. In *CRYPTO'99: 19th International Cryptology Conference*, volume 1666 of *LNCS*, pages 148–164. Springer, 1999.
- [Sch09] Berry Schoenmakers. Voting Schemes. In Mikhail J. Atallah and Marina Blanton, editors, *Algorithms and Theory of Computation Handbook, Second Edition, Volume 2: Special Topics and Techniques*, chapter 15. CRC Press, 2009.
- [Sha71] Daniel Shanks. Class number, a theory of factorization and genera. In *Number Theory Institute*, volume 20 of *Symposia in Pure Mathematics*, pages 415–440. American Mathematical Society, 1971.
- [SK94] Kazue Sako and Joe Kilian. Secure Voting Using Partially Compatible Homomorphisms. In *CRYPTO'94: 14th International Cryptology Conference*, volume 839 of *LNCS*, pages 411–424. Springer, 1994.
- [Smy11] Ben Smyth. *Formal verification of cryptographic protocols with automated reasoning*. PhD thesis, School of Computer Science, University of Birmingham, 2011.
- [SRKK10] Ben Smyth, Mark D. Ryan, Steve Kremer, and Mounira Kourjeh. Towards automatic analysis of election verifiability properties. In *ARSPA-WITS'10: Joint Workshop on Automated Reasoning for Security Protocol Analysis and Issues in the Theory of Security*, volume 6186 of *LNCS*, pages 165–182. Springer, 2010.
- [VG10] Melanie Volkamer and Rüdiger Grimm. Determine the Resilience of Evaluated Internet Voting Systems. In *Re-Vote'09: First International Workshop on Requirements Engineering for E-Voting Systems*, pages 47–54. IEEE Computer Society, 2010.
- [Vol09] Melanie Volkamer. *Evaluation of Electronic Voting: Requirements and Evaluation Procedures to Support Responsible Election Authorities*, volume 30 of *Lecture Notes in Business Information Processing*. Springer, 2009.

- [War03] Bogdan Warinschi. A Computational Analysis of the Needham-Schröder-(Lowe) Protocol. In *CSFW'03: 16th Computer Security Foundations Workshop*, pages 248–262. IEEE Computer Society, 2003.
- [War05] Bogdan Warinschi. A computational analysis of the Needham-Schroeder-(Lowe) protocol. *Journal of Computer Security*, 13(3):565–591, 2005.
- [Wik06] Douglas Wikström. Simplified Submission of Inputs to Protocols. Cryptology ePrint Archive, Report 2006/259, 2006.
- [Wik08] Douglas Wikström. Simplified Submission of Inputs to Protocols. In *SCN'08: 6th International Conference on Security and Cryptography for Networks*, volume 5229 of *LNCS*, pages 293–308. Springer, 2008.