

Election Verifiability in Receipt-free Voting Protocols

Sevdenur Baloglu*, Sergiu Bursuc*, Sjouke Mauw*[†], Jun Pang*[†]

*SnT, University of Luxembourg

[†]DCS, University of Luxembourg

Email: {sevdenur.baloglu, sergiu.bursuc, sjouke.mauw, jun.pang}@uni.lu

Abstract—Electronic voting is a prominent example of conflicting requirements in security protocols, as the triad of privacy, verifiability and usability is essential for their deployment in practice. Receipt-freeness is a particularly strong notion of privacy, stating that it should be preserved even if voters cooperate with the adversary. While there are impossibility results showing we cannot have receipt-freeness and verifiability at the same time, there are several protocols that aim to achieve both, based on carefully devised trust assumptions. To evaluate their security, we propose a general symbolic definition of election verifiability, extending the state of the art to capture the more complex structure of receipt-free protocols. We apply this definition to analyse, using ProVerif, recent protocols with promising practical features: BeleniosRF and several variants of Selene. Against BeleniosRF, we find several attacks showing that verifiability in Belenios does indeed suffer from the attempt to introduce receipt-freeness. On the other hand, Selene satisfies a weaker notion of receipt-freeness, but we show that it satisfies verifiability in stronger corruption scenarios. We introduce a general framework to compare the verifiability of these protocols in various corruption scenarios and conclude with an analysis of SeleneRF, an attempt to get the best of both that we formalise in this paper. In addition to extending the symbolic model, our results point to foundational gaps in current cryptographic models for election verifiability, as they fail to uncover attacks that we do.

I. INTRODUCTION

Electronic voting brings many advantages to the process of running an election, providing more straightforward procedures and reducing the costs and the potential for human errors. On the other hand, it also brings in possible security weaknesses by increasing the scope and abilities of an adversary against the election. Electronic voting protocols aim to address this problem by incorporating a strong cryptographic foundation, ensuring the security and robustness of the system independently of the corruption of software, hardware or any of the parties involved. There are two main security properties such protocols aim to achieve: *voter privacy* [8], [25] and *election verifiability* [20], [28]. Furthermore, privacy should hold even when voters may cooperate with the adversary. This is to prevent, for example, vote buying or voter coercion. This stronger notion of privacy is typically called receipt-freeness. Both receipt-freeness and verifiability should hold with minimal trust assumptions, reflecting a set of maximally admitted corruption abilities for the adversary.

There is a well-known tension between receipt-freeness and verifiability in electronic voting protocols. On the one hand, there are impossibility results like [14], although one

could argue that their underlying assumptions do not apply in all cases. On the other hand, initial protocols attempting to achieve the two, like JCJ/Civitas [15], [27], remained mostly of theoretical interest, due to strong assumptions and weak usability. If we drop the requirement of receipt-freeness, however, there are several deployed practical systems satisfying both privacy and verifiability. Most prominently, Helios [3], [4] and Belenios [2], [21] also enjoy formal security proofs for privacy [16], [22]–[24] and verifiability [6], [17]–[19]. Belenios is an elaboration of Helios aiming to achieve verifiability in stronger corruption scenarios. One step further, BeleniosRF relies on a novel cryptographic primitive (signatures on randomizable ciphertexts) in order to achieve receipt-freeness with trust assumptions similar to Belenios, albeit relying on a less intuitive individual verification procedure [12].

Selene is a recently proposed system that achieves receipt-freeness based on a new mechanism that allows voters to verify their vote directly in the final outcome, while hiding the link between the credentials used for vote casting and the credentials used for vote verification [35]. The idea is that the verification credential (also called tracker) is held in encrypted and committed form next to the voting credential. Using private trapdoor keys associated to each voter, and randomness obtained from trustees after tally, voters can open the commitment and obtain their trackers. Furthermore, this information also allows voters to generate fake randomness in order to resist coercion. Selene has an advantage over BeleniosRF in the usability of its individual verification procedure: voters can directly check their votes associated to trackers in the outcome. The disadvantage is that voters have to generate and register their private trapdoor keys, and receipt-freeness is not immediate like in BeleniosRF. Furthermore, since ballots are posted as such on the bulletin board, the receipt-freeness property is also weaker, as the ballot can be part of a receipt for the adversary.

Formal security proofs have been essential for the development of Helios and Belenios. Many attacks against these systems were found with formal methods and have led to security improvements. In this paper we aim to extend the required formal foundations in order to perform automated verification of election verifiability in receipt-free voting protocols, including BeleniosRF, Selene, and others. A general symbolic definition of election verifiability has been recently proposed in [6] and applied to Helios and Belenios. It considers various adversarial

models and ensures strong end-to-end security guarantees. However, it cannot capture an essential feature of some receipt-free voting protocols, like Selene: the verification credential is not publicly linked to the vote casting credential. This is not a mere technicality and may hide real attacks against election verifiability, for example if the adversary manages to map two voting credentials onto the same verification credential, in a variation of clash attacks [29], [30], [33].

Our contributions. We extend the symbolic definition of election verifiability from [6] to also cover the more complex case of receipt-free voting systems. We also improve the definition in [6] in order to more closely match the general definition of election verifiability used in previous works, e.g. in [17]–[19]. We apply the definition to verify with ProVerif the protocols BeleniosRF, Selene and variants of Selene that improve its efficiency or receipt-freeness. We find new attacks on BeleniosRF and provide new proofs for Selene in strong corruption scenarios. We identify various corruption scenarios that pinpoint the trade-offs between BeleniosRF and Selene, and between receipt-freeness and verifiability. We give more details about these contributions in the following.

General vs symbolic election verifiability. The symbolic definition in [6] is generic and can be applied to many corruption scenarios. It has been used to find real attacks and proofs for electronic voting protocols. However, it has one important drawback, in that it was not proved to imply the general definition of election verifiability used in previous work and first introduced in [19]. Conversely, some attacks found with [6] against Belenios do not represent real attacks against election verifiability, so their definition is too strong for some scenarios. The issue in both these cases lies in a departure that was made in [6] from the notion of corrupted voters towards the notion of corrupted credentials, which creates a gap between what is proved (with respect to credentials) and the real execution of the system (with respect to voters). We fix this problem by introducing a stronger formal connection between voter identities and public credentials.

Symbolic election verifiability vs receipt-freeness. End-to-end election verifiability aims to ensure a correspondence between the final outcome and other events that occur during a run of the voting protocol: registration, corruption, voting, verification, etc. Any of these events can be associated to a particular voting credential. In Selene, each vote in the final outcome is also associated to a verification credential. For receipt-freeness, it is important that verification credentials cannot be publicly linked to their corresponding voting credential. This means that, in order to formalise election verifiability in this context, we need a function to *extract* the verification credential associated to a voting credential. Of course, to preserve receipt-freeness, this function should not be efficiently computable. For verifiability, it is sufficient for it to be well-defined and for its arguments to come from the bulletin board. We provide a framework that ensures the extraction function is well-defined and, when plugged into our updated symbolic verifiability definition, does indeed ensure

end-to-end verifiability.

Attacks vs proofs for BeleniosRF. Our attacks against BeleniosRF may be surprising, considering that it was formally shown to satisfy election verifiability [12], based on a generic transformation proved secure in [19]. In fact, our attacks show the importance of rigorous adversarial models combined with machine-checked proofs. There are two problems with the definition and proofs in [19]. For the case of a corrupt registrar and honest server, the definition in [19] (relying on a cryptographic model) fails to provide the adversary with the ability to control the signing key of voters. Since security is derived following a hand-based proof, the issue of key generation was easy to oversee. A second problem, common to all current cryptographic models of verifiability (including the machine-checked model proposed later for Belenios in [17]), is that, when the voting server is honest, the ballot is posted directly on the bulletin board, assuming in effect that the communication network and all underlying infrastructure on the voting platform is secure. This defeats the purpose of verifiability, that aims to counter precisely attacks against such infrastructure. We follow the approach standard in symbolic definitions [5], [6], [18], [28], which is more realistic in this respect: the only assumption is that the voting server performs ballot verification and authentication, while the adversary is allowed to control the communication network and block ballots on their way to the bulletin board.

Once the adversary is allowed to control the signing key for honest voters and the communication network, the attack against BeleniosRF follows from the fact that voters do not verify the cast ballot directly on the bulletin board. Instead, when a cast ballot b reaches the server, it is rerandomized into a ballot b' (maintaining a signature with respect to the same key as b) that is posted on the bulletin board. The voter verifies b' by ensuring it is signed with the voter's own key, which ensures that b' is indeed a rerandomization of b . However, the following clash attack can be mounted in the considered adversarial scenario. At registration, the adversary assigns the same key to two honest voters id_1 (voting v_1) and id_2 (voting v_2). During voting, the adversary allows the ballot from id_1 to be cast, and blocks the ballot from id_2 , relying on the control of the communication network. Both of these voters will successfully verify the ensuing ballot b' on the bulletin board. For id_2 , this will clearly violate individual verifiability (when v_2 is different from v_1) or the no-clash property (when v_2 is equal to v_1). Overall, end-to-end verifiability is violated since not all votes that have been successfully verified are included in the final outcome.

The case of corrupt voters. We find another weakness of BeleniosRF over Belenios in the case when voter credentials are corrupted. This possibility underlies the stronger notion of verifiability introduced in [6] in order to offer guarantees in the case when voters inadvertently lose their credentials. This case is different from that of a corrupted registrar: on the one hand, it means that all private credentials of a voter are leaked, not only the signing key; on the other hand, the adversary cannot

choose the secret key of voters, preventing the attack presented above. For this scenario, Belenios was proved to be secure in [6], while a simple variation of the attack described above leads to a violation of individual verifiability in BeleniosRF: a voter can successfully verify a ballot b' even if it is the rerandomization of a ballot cast by the adversary.

The case of corrupt platforms. The case of a corrupted voting platform can be countered in Helios and Belenios relying on the classic cast-or-audit mechanism known as the Benaloh challenge [7]. This type of challenge can have a usability problem, as confirmed in a recent study [31]. An interesting aspect of Selene is that such a mechanism is not necessary in this case, since voters verify their votes directly in the outcome. We note, however, that voters still need to ensure their private trapdoor keys are secret, e.g. stored on a device that is more trustworthy than the voting platform. A formal model for this corruption scenario has not been considered before, and it is another contribution of this paper. Considering this model, election verifiability is violated in BeleniosRF (as expected, assuming no Benaloh challenge is performed), while we prove that Selene is indeed secure in this case - satisfying however a weaker notion of security that allows the adversary to cast ballots for honest voters not verifying their votes.

Towards stronger security for Selene. Our findings show stronger verifiability properties in Selene over BeleniosRF, while it is known that BeleniosRF satisfies a stronger notion of receipt-freeness due to signatures on rerandomizable ciphertexts. One may read this as an illustration of the trade-off between the two properties. Yet, as suggested in [35], we can also consider a version of Selene that relies on the same primitive as BeleniosRF to rerandomize ballots before posting. We call this version SeleneRF and we prove with ProVerif that it satisfies the same verifiability properties as Selene. We expect it to satisfy the same receipt-freeness notion as BeleniosRF, but this has to be formally proven in future work. Interestingly, in spite of the fact that the ballots are rerandomized, in SeleneRF we do not find attacks similar to the ones described above for BeleniosRF. This is due to the fact that voters directly verify their votes in the outcome, and not their ballots. An advantage of BeleniosRF is that it thwarts a coercer by design, and the voter does not have to follow any coercion-resistance strategy to achieve receipt-freeness. In Selene and its variants, the voter needs to construct a fake verification credential to resist coercion. Another improvement of Selene that we consider is a recently proposed version called Hyperion, showing a simpler way of generating verification credentials [34]. The system description in [34] is quite informal. We provide a formal specification and we prove its election verifiability property with ProVerif.

Paper structure. Section II contains preliminaries for e-voting protocols and formal verification with ProVerif. Section III introduces receipt-free voting protocols, specifically BeleniosRF, Selene, and Hyperion. A general symbolic verifiability definition revised to include receipt-free voting protocols is presented in Section IV. Finally, a comparative verifiability

analysis of the considered protocols is presented in Section V.

II. PRELIMINARIES

A. Verifiable electronic voting protocols

Bulletin board. Electronic voting protocols utilise a public bulletin board, which we denote by BB , to record information regarding the execution of the protocol for the verifiability and transparency of the election. It can be used by the voters to verify their ballots after the voting procedure (this is also called *individual verifiability*) and by any party to verify the integrity of the election (*universal verifiability*). We can distinguish different parts of the BB by annotations, e.g. $BBkey$ for the public key of the election or $BBreg$ for the registered eligible public credentials.

Setup. A voting protocol assumes a set of eligible voters represented by identities id_1, \dots, id_n . Election authorities generate their corresponding public credentials cr_1, \dots, cr_n , that we assume recorded on $BBreg$. The public key of the election is generated by authorities and posted on $BBkey$. At registration time, voters obtain their (private) id , their public credential cr , and additional private credentials which will be used for authentication and verification.

Voting procedure. During the voting phase, voters use voting platforms to construct a ballot b , encoding their desired vote v . The ballot is sent to the voting server after authentication. The server validates the ballot b and publishes corresponding information on BB . Typically the ballot itself is published on BB along with the public credential of the voter, but other options are possible. For example, BeleniosRF publishes a rerandomization of the ballot.

Individual verification. Voters can perform certain procedures to ensure that the system correctly records their votes. Typically, voters have to check that information they expect to see (e.g. their ballot or their vote) is present on the bulletin board along with tracking information associated to their public credential.

Tally procedure. The ballots to be tallied are recorded alongside eligible public credentials on the bulletin board: $(cr_1, b_1), \dots, (cr_n, b_n)$. Each ballot b_i is stripped of any information that is not needed for vote-counting, like signatures and zero-knowledge proofs, resulting in a ciphertext c_i . The tally phase consists in computing the set of votes $\{v_1, \dots, v_n\}$ corresponding to $\{c_1, \dots, c_n\}$ in an anonymous and verifiable way, typically relying on mixnets or on the homomorphic properties of the encryption scheme. For some voting protocols, ciphertexts are combined with trackers before decryption. In this case, each vote in the outcome is associated to a tracker, like $(tr_1, v_1), \dots, (tr_n, v_n)$.

Universal verification. External auditors can be involved in verifying that the ballots b_1, \dots, b_n are really cast by the corresponding public credentials and satisfy any additional consistency properties expected from the specification of the protocols. For example, they may check that the ballots are signed with the right key, contain valid zero-knowledge proofs,

are the last ballots cast for the corresponding credential, there are no duplicates, etc. In addition, auditors also verify the proofs provided in the tally phase, to ensure that $\{v_1, \dots, v_n\}$ corresponds to $\{c_1, \dots, c_n\}$.

Example 1. Assume three election authorities: a trustee responsible for tallying ballots, a registrar for generating and registering voter keys, and a server for ballot casting.

Setup. The trustee generates a key pair (pk, sk) and publishes pk on BBkey. For each voter id , the registrar generates a signing key pair (spk, ssk) and publishes spk as cr on BBreg. Moreover, the server generates a password pwd to use for authentication of the voter. Each eligible voter obtains (id, pwd, cr, ssk) .

Voting. The voter, relying on the voting platform, generates a ciphertext $c = enc(v, pk, r)$ for some desired vote v and fresh randomness r . Then, the voter signs it with their ssk , obtaining a signature s and constructs related zero-knowledge proof π , resulting in a ballot $b = \langle c, s, \pi \rangle$. The voter authenticates to the server with their pwd , and the ballot b is submitted to the server. After server's ballot validation, (cr, b) is published on BBcast.

Verification. Voters can check their ballot on BBcast anytime during the voting phase, or on BBtally afterwards.

Tally. The last ballot corresponding to cr on BBcast is recorded on BBtally. The result is computed by the trustee by decrypting each ciphertext from ballots in BBtally.

The previous example is a simplified version of Belenios [21], an elaboration of Helios that protects improves election verifiability relying on a registrar. Auditors ensure that all ballots on BB are signed with keys recorded by the registrar. The zero-knowledge proof on the ballot also helps to protect ballot privacy against copy attacks [24]. Even if the adversary tries to copy any ciphertext cast by any voter, they cannot generate the zero-knowledge proof without knowing the randomness used for the ciphertext.

B. Formal verification with ProVerif

ProVerif allows automated verification of security properties for cryptographic protocols [9]. Messages are represented in an equational theory where terms are built from a set of types, a set of constants, a set C of constructors endowed with a set D of destructors. Types are not essential for describing our protocol specifications, so we will not use them in this paper. A constructor is a function symbol f used to build a term $f(M_1, \dots, M_k)$, from terms M_1, \dots, M_k . A destructor is a symbol with associated rewrite rules modeling the properties of cryptographic primitives. Rewrite rules are equations $u = v$ that should be read from left to right. For example, we can define a $proj_i$ destructor to fetch the arguments of f : $proj_i(f(M_1, \dots, M_k)) = M_i$. If a term contains a destructor for which no rewrite rule can be applied, we denote this by $t = \perp$, i.e. the evaluation of the term fails. A term may also be a tuple of the form (M_1, \dots, M_n) , where M_1, \dots, M_n are terms. We consider a special name $pub \in \mathcal{N}$ to represent a public

communication channel. Unless declared *private*, names and functions are known by the adversary. A term M can be stored in a table, e.g. $Tb(M)$. Tables are based on private constructors and are not accessible by the adversary. Terms can be checked with equality or disequality conditions, and such conditions may be combined with logical connectors.

Example 2. The set of constructors $\{pk, spk, enc, pr, sign\}$ can model cryptographic primitives from Example 1, along with the corresponding destructors defined by:

- (1) $dec(enc(x, pk(y), z), y) = x$,
- (2) $check(sign(x, y), spk(y)) = x$,
- (3) $ver(pr(enc(x, y, z), x, z, w), enc(x, y, z), y, w) = ok$.

Equations (1) and (2) specify the standard properties of asymmetric encryption and digital signatures. The proof symbol models a zero-knowledge proof showing that a ciphertext is constructed by a party that knows the corresponding plaintext m and randomness r , and that party associated another message w to this ciphertext. Equation (3) models the corresponding verification procedure for the proof and corresponding arguments.

$M, N ::=$	terms
$a, b, c, \dots, x, y, z, \dots$	names and variables
$(M_1, \dots, M_k), f(M_1, \dots, M_k)$	tuples and functions
$M = N, M <> N$	term equality and disequality
$M \&\& N, M N, \text{not}(M)$	conjunction, disjunction, negation
$P, Q ::=$	processes
0	parallel, replication, restriction
$P Q$	parallel execution of P and Q
$!P$	arbitrary number of parallel executions of P
$\text{new } n; P$	restriction
$\text{in}(N, x); P$	message input and output
$\text{out}(N, M); P$	message output
$\text{if } M \text{ then } P \text{ else } Q$	conditional
$\text{let } x = M \text{ in } P \text{ else } Q$	term evaluation
$\text{insert } Tb(M); P$	table insertion
$\text{get } Tb(x) \text{ suchthat } M$	table extraction
$\text{event}(M_1, \dots, M_k); P$	event

Fig. 1. Terms and processes in ProVerif grammar.

Processes are denoted by P, Q, \dots , and generated with the grammar from Figure 1. Recursively, 0 represents the null process, $P | Q$ the parallel execution of P and Q , and $!P$ an arbitrary number of parallel executions of P . The restriction $\text{new } n; P$ generates the name n , which can represent fresh randomness in the process P . The process $\text{in}(pub, x); P$ models a message received from channel pub and stored in x , while the process $\text{out}(pub, M); P$ models a process sending the message M on channel pub and continuing as P . The conditional $\text{if } M \text{ then } P \text{ else } Q$ runs P if M evaluates to true, runs Q if $M = \text{false}$. If a term M contains a destructor where no rewrite rule can be applied, we say that M fails. The term evaluation $\text{let } x = M \text{ in } P \text{ else } Q$, bounds x to M and executes P , when M does not fail. If M fails, it directly executes the process Q . The table insertion $\text{insert } Tb(M); P$ stores the term M in Tb . The table extraction $\text{get } Tb(x) \text{ suchthat } M \text{ in } P \text{ else } Q$ selects an entry from Tb that evaluates to M . If there is no such entry, it executes Q .

```

let Voting Platform(v, pk) =
  get Cred(id, cr, ssk) in
  get Pwd(= id, pwd) in
  new r; let c = enc(v, pk, r) in
  let s = sign(c, ssk) in
  let pr = proof(c, v, r, cr) in
  let b = (c, s, pr) in
  let a = h(id, pwd, cr, b) in
  event Vote(id, cr, v);
  insert Voted(id, cr, v, b);
  out(pub, (id, cr, b, a)).

let Voting Server(pk) =
  get BBreg(cr) in
  get Pwd(id, pwd) in
  in(pub, (= id, = cr, c, s, pr, a));
  let b = (c, s, pr) in
  if h(id, pwd, cr, b) = a then
  if check(s, cr) = c then
  if ver(pr, c, pk, cr) = ok then
  event BBcast(cr, b);
  insert BBcast(cr, b).

```

Fig. 2. Ballot casting processes for voting platform and server.

Example 3. The processes in Figure 2 model basic operations of a voting protocol. BB is modelled by tables. Voting credentials are retrieved from Cred and Pwd, and a vote v is chosen. A signature s , a proof pr , and authentication information a are computed. Information relevant for individual verifiability is stored in the table Voted(id, cr, v, b), and an event Vote(id, cr, v) is recorded. The tuple (id, cr, b, a) is sent to the public channel. The voting server receives b and publishes it on BB if validation checks are ok.

Security properties. Security properties in this paper are modelled by correspondence assertions, capturing relations between events. For example, the formula $e(M) \Rightarrow e'(M')$ states that if the event $e(M)$ is executed in the trace, then the event $e'(M')$ has been executed before in that trace. Queries can be more complex, combining events with relations between terms inside the events, e.g. we can add $M = M'$ on the right-hand side of the previous query. In general, we consider a set of timepoint variables, denoted by i, j, l, \dots , which will be interpreted over rational numbers. A trace atom is either a term equality $t_1 = t_2$, or a timepoint ordering $i < j$, or a timepoint equality $i = j$, or an event atom $e(M)@i$ - the timepoint can be omitted from the event atom if it is not relevant for the property. A trace formula is a first-order logic formula obtained from trace atoms by applying the usual quantification and logical connectives.

Example 4. Considering the processes from Figure 2 and assuming a destructor open that reveals the vote v inside a ballot b , we can check whether any ballot from BBcast comes from a corresponding Vote event in the voting platform, using formula: $BBcast(cr, b) \Rightarrow Vote(id, cr, v) \wedge open(b) = v$. This formula allows multiple cast ballots per credential, and for all of them it requires a corresponding voting event.

An execution trace τ of a process P is defined as a sequence of events E_1, \dots, E_n that occurred during the execution of P . The satisfaction relation $\tau \models \Phi$, for a trace τ and a trace formula Φ , whose variables are all bounded, is defined recursively as expected, with the following notable case: $\tau \models e(M)@i$ if and only if $e(M)$ occurs on the i -th position of the trace. For a process P , we let $tr(P)$ be the set of traces of P . For trace formulas Ψ, Φ , we let:

$$\begin{aligned}
P \models \Phi & \text{ iff } \forall \tau \in tr(P). \tau \models \Phi, \\
P, \Psi \models \Phi & \text{ iff } \forall \tau \in tr(P). \tau \models \Psi \Rightarrow \Phi.
\end{aligned}$$




Fig. 3. Illustrating revoting and verifiability in various protocols. For Belenios and Selene: b_2 is a revote ballot with desired vote. For BeleniosRF: b_1 contains the desired vote rerandomized to b'_1 .

For verification of security properties, $(P; \Psi)$ is typically a protocol specification and Φ the property to be verified. Having the component Ψ , also called a restriction, in a protocol specification can help to express in a concise way some properties that protocol parties should ensure along an execution trace. If the equational theory \mathcal{E} is not clear from the context, we write $(P, \Psi, \mathcal{E}) \models \Phi$.

III. RECEIPT-FREE VOTING PROTOCOLS

We can distinguish three techniques that are instrumental for receipt-free voting protocols. First, like in JCJ/Civitas [15], [27], voters can be given the ability to construct a fake credential, which can be used to cast a ballot according to the request of the adversary. Second, like in BeleniosRF [12], ballots published on the bulletin board can be derived from cast ballots with advanced cryptography that allows to ensure the published ballots encode the same votes as cast ballots, while hiding the link between them. Finally, like in Selene [35] and Hyperion [34], voters can be given the ability to derive a tracker showing their vote in the final outcome; to resist coercion, they have the ability to derive a fake tracker pointing to a vote desired by the adversary, relying on their private trapdoor keys. In Figure 3, we show how these protocols relate to a classic protocol like Belenios. In Belenios, the ballot is verified on the bulletin board in the same form as it was cast. In BeleniosRF, the ballot is randomized by the voting server before being posted on BB: it cannot be used as a coercion tool. BeleniosRF disallows revoting; otherwise, a corrupt server could choose which of the ballots from a given voter to cast.

In Selene and Hyperion, the vote is verified directly in the outcome after trackers are revealed to voters. Nobody can associate a credential to a tracker on BB, except the voter and dedicated election authorities. In case of coercion, a private trapdoor key allows voters to choose any tracker from the outcome that conforms to the instructions of the coercer. We illustrate this approach with a simple example first. A *trapdoor commitment scheme* consists of two main operations: $com(m, tpk, r)$ allows to commit a message m for trapdoor public key tpk with randomness r ; $open(cm, ts, r)$ allows to obtain the message m from a commitment relying on the corresponding trapdoor secret key and randomness. This key may also allow to produce fake randomness that opens the commitment to a different message.

Example 5. Assume each voter has a trapdoor key pair $\langle \text{tpk}, \text{tsk} \rangle$ for a commitment scheme. Voters publish their tpk next to their credential cr on BB. Election authorities generate a tracker tr for each voter and publish encryption of tr and a commitment to tr for the trapdoor public key of the corresponding voter. After voting, we have the following information on BB for each credential: $\text{cr}_i, \text{enc}(\text{tr}_i, \text{pk}), \text{com}(\text{tr}_i, \text{tpk}_i, r_i), \text{enc}(v_i, \text{pk})$. Votes and trackers are decrypted at tally: $\text{BB} : (\text{tr}_1, v_1), \dots, (\text{tr}_n, v_n)$.

For individual verification, election authorities send the commitment randomness r_i to each voter. Voters obtain their tracker by applying $\text{open}(\text{cm}, \text{tsk}_i, r_i)$, where cm is the commitment recorded next to their credential. For receipt-freeness, voters can choose any tracker tr_j from the outcome and construct fake randomness that opens their commitment to this tracker.

A. BeleniosRF

BeleniosRF [12] is a variant of Belenios [21], which aims for receipt-freeness in addition to end-to-end verifiability. It relies on a cryptographic primitive called signatures on randomizable ciphertexts [10], instantiated from Groth-Shai proofs [26]. Like in Belenios, the public credential of a voter is the public part of the signing key generated by the registrar, and the ballots cast are of the form $b = \langle c, s, \pi \rangle$, where c is an encryption of the desired vote, s is the signature of c , and π is a zero-knowledge proof showing the validity of the ciphertext and associating it to the public credential of the voter. Unlike Belenios, the ballot is rerandomized before being posted next to cr on BBcast. More precisely: the voting server re-encrypts the ciphertext c and, relying on [10], adapts the signature s and the zero-knowledge proof π for the new ciphertext c' . The randomized ballot $b' = \langle c', s', \pi' \rangle$ is now a valid ballot for the same credential cr and is published next to it on BBcast.

The individual verification procedure for a voter with credential cr consists in verifying that there is a ballot b' associated to cr on BBcast, and verifying that b' contains a valid signature with respect to cr . Intuitively, this should be sufficient to ensure election verifiability if either the registrar or the server is not corrupt. If the server is honest, it will log the identity of voters together with their public credentials, so even if a voter's secret key is corrupted, the adversary should not be able to cast a ballot for the credential of an honest voter. Symmetrically, if the registrar is honest, the only party who can construct valid ballots for a public credential is the corresponding voter - verifiability follows even for a corrupt server since revoting is not allowed. In addition, BeleniosRF satisfies receipt-freeness: a coerced voter cannot prove how they voted, as there is no way for the adversary to determine whether any ballot provided by the voter corresponds to the one cast on the bulletin board - relying on the assumption that the voting server is honest.

B. Selene

The main feature of Selene [35] is that it enables intuitive end-to-end verifiability by allowing voters to verify their vote

directly in the outcome of the election. For this purpose voters are provided with trackers that point to their votes in the outcome. For receipt-freeness, trackers can be obtained by voters only at the end of the election: relying on trapdoor commitments like in Example 5, election authorities send voters the information that is required to open their tracker from a commitment, and also to construct a fake tracker for coercion-resistance. Selene relies on ElGamal encryption, a digital signature scheme, non-interactive zero-knowledge proofs and Pedersen trapdoor commitments. Mixnets are used for anonymising ciphertexts containing votes and trackers before decryption.

Encrypted trackers $\text{enc}(\text{tr}_1, \text{pk}), \dots, \text{enc}(\text{tr}_n, \text{pk})$ are published on BB, where pk is the public key of the election (with the secret key held by election trustees) and $\text{tr}_1, \dots, \text{tr}_n$ is a random permutation obtained by encrypting each element of a public list of trackers and passing it through a re-encryption mixnet [13], [32]. Like in Belenios, the public credential of each voter is a public signing key. In addition, voters generate a trapdoor key pair $\langle \text{tpk}, \text{tsk} \rangle$, and register its public part next to their credential: $\text{BB} : (\text{cr}_1, \text{tpk}_1), \dots, (\text{cr}_n, \text{tpk}_n)$. Trustees associate each credential with an encrypted tracker. In addition, based on the homomorphic properties of ElGamal and on its interaction with the Pedersen commitment scheme, the trustees can compute a commitment to the tracker for the public trapdoor key associated to the corresponding voter. The commitment along with a zero knowledge proof of correctness π_{ptr}^i is recorded for each credential cr_i on the bulletin board:

$$\overbrace{\text{BBptr} : \text{cr}_i, \text{tpk}_i, \text{enc}(\text{tr}_i, \text{pk}), \text{com}(\text{tr}_i, \text{tpk}_i, r_i), \pi_{\text{ptr}}^i, \text{pk}}^{\text{tracking data}}$$

Ballots cast by the voter are of the form $b = \langle c, s, \pi \rangle$, which is the same as in Belenios, and are posted alongside the public credential on BB. Before the tally, zero-knowledge proofs and signatures are verified. For computing the outcome, the pairs of ciphertexts $(\text{enc}(\text{tr}_1, \text{pk}), c_1), \dots, (\text{enc}(\text{tr}_n, \text{pk}), c_n)$ corresponding to each voter are passed through a parallel re-encryption mixnet. Such a mixnet will produce a list of pairs that encode the same information as the input list of pairs, where each element is unlinkable to the corresponding element in the input. Each pair is verifiably decrypted, resulting in the outcome $\text{BBres} : (\text{tr}_1, v_1), \dots, (\text{tr}_n, v_n), \pi_{\text{res}}$, where π_{res} is an overall proof of re-encryption and decryption, ensuring that the outcome is correctly computed from the list of input ciphertext pairs. After the outcome is computed, each voter receives the randomness r for their corresponding commitment from trustees. Using their trapdoor secret key tsk and the randomness r , they can open the commitment and obtain their trackers for verification. If voters are coerced, they can construct fake randomness to open their commitment to a different tracker. Suppose the coercer has forced the voter to vote for v' . The voter finds a tracker-vote pair (tr', v') from the outcome on BB, and computes the fake randomness r' using their tsk , tr and r that will open their commitment to tr' . The trapdoor signing key tsk of the voter and fake randomness r'

can be provided to the coercer, which opens their commitment to a different tracker tr' , showing that the voter voted as the coercer desired.

We also consider SeleneRF, that adopts the primitive for signatures on rerandomizable ciphertexts from BeleniosRF. The protocol is exactly the same as Selene, the only difference being that the signature and proofs on the ballot are performed with the primitives from [12]. The tally and verification procedures remain the same as in Selene, as a rerandomized ballot b' can be tallied in the same way as b and the trackers can be extracted in the same way.

C. Hyperion

Hyperion improves the tracker management in Selene by deriving the tracker for each voter directly from their public trapdoor key [34]. This removes the need to generate, encrypt and anonymise trackers in advance. Moreover, computing pre-tracking information and the trackers in the final outcome is also more efficient. The main idea in Hyperion is that, instead of encrypted and committed trackers like in Selene, the pre-tracking information is simply tpk raised to a power r , where tpk is the trapdoor public key of the voter, and r is randomness chosen by trustees, different for each voter. We call tpk^r a pre-tracker. Once the voting phase has ended, each pre-tracker tpk^r is paired with the ciphertexts c encoding the vote of the corresponding voter. The list of all such pairs $(tpk_1^{r_1}, c_1), \dots, (tpk_n^{r_n}, c_n)$ is passed through a parallel mixnet where the second component of each pair is re-encrypted to a different ciphertext, as in a classical mixnet, while the first component is exponentiated to a power s , where s is the same for all the elements of the list. The mixnet for the first component is called an exponentiation mix. After decrypting the ciphertexts in the resulting outcome, we obtain:

$$\text{BBres} : (tr_1, v_1), \dots, (tr_n, v_n), \pi_{res} \\ \text{where } tr_1 = tpk_1^{r_1 s}, \dots, tr_n = tpk_n^{r_n s}$$

The trustees know g^{r_i} and s , from which they compute $g^{r_i s}$ and send it to the corresponding voter. Furthermore, we should note that $tpk_i = g^{spk_i}$, where spk_i is the trapdoor secret key of the voter. Therefore, by exponentiating $g^{r_i s}$ to the power spk_i , voters can obtain their trackers $tr_i = tpk_i^{r_i s}$. As for Selene, we gather all data relevant for tracking on BBptr :

$$\overbrace{\text{BBptr} : cr_i, tpk_i, tpk_i^{r_i}}^{\text{tracking data}}, \pi$$

The proof $\pi = \text{prmix}(M_{in}, s, M_{out})$ represents the (exponentiation) mixnet proof showing that the final set of trackers M_{out} is obtained by exponentiating each element in the set of pre-trackers M_{in} to the same secret power s . A further enhancement proposed in [34] is to derive a separate bulletin board for each voter. This improves receipt-freeness, since any fake tracker that the voter may derive cannot clash with trackers owned by the coercer. This is especially pertinent in elections with a small number of voters. We do not consider this option in this paper, but it should be straightforward to extend our analysis for this

case, as the individual bulletin board for each voter is derived from a further exponentiation mix, with a separate s_i for each voter, and the corresponding tracker is $tpk_i^{r_i s s_i}$. We need to record g^{s_i} on BB and update the tracker extraction function accordingly.

IV. ELECTION VERIFIABILITY DEFINITIONS

Before introducing our symbolic definition for election verifiability, we recall the definition of election verifiability from [19], that has been adapted both to computational [17] and symbolic models [18]. Furthermore, we introduce some notation that will be useful for bridging the gap from the general definition to the symbolic model, and we also extend the definition from [19], in order to have a generic notion that can cover stronger or weaker guarantees of election verifiability depending on considered voting protocols or trust assumptions.

A. General definition of election verifiability

Definition 1. A voting specification is a tuple $(\mathcal{P}, \Psi, \mathcal{E}, \Omega)$, for a process \mathcal{P} , a restriction Ψ , an equational theory \mathcal{E} and a trace formula Ω called the revote policy. Furthermore, we assume \mathcal{P} relies on fact symbols Vote , BBres , Verified , Corr to record the following events:

- $\text{Vote}(id, v)$: vote v is cast by the voter with identity id ;
- $\text{BBres}(tr, v)$: vote v is in the result next to a counter tr ;
- $\text{Verified}(id, cr, v)$: vote v is verified by a voter with identity id and public credential cr ;
- $\text{Corr}(id)$: the voter with identity id is corrupt.

The process \mathcal{P} and restrictions Ψ model the procedures of an electronic voting protocol: setup, voting, tally, individual and public verification; as well as the abilities of the adversary to corrupt parties and control the network. The revote policy $\Omega(id, v)$ should say under what conditions the vote v that was cast by id should be counted. For example, it may say that v is the first or the last vote cast by the respective voter. Some protocols like Helios, Belenios, and BeleniosRF do not have counters (or trackers) associated to votes in the outcome. For such systems, the result of the election is considered to be the multiset of votes encoded by the ballots recorded on a special portion of the bulletin board called BBtally , where each ballot b is recorded next to a public credential cr , resulting in events $\text{BBtally}(cr, b)$. The votes inside these ballots are determined by an extraction function open , as in e.g. [6], [18]. To apply Definition 1 to such protocols, we can therefore take $tr = cr$ and define: $\text{BBres}(cr, v) \equiv \text{BBtally}(cr, b) \wedge \text{open}(b) = v$.

The definition of election verifiability from [19] ensures that votes in the final result can be partitioned into three multisets V_1, V_2, V_3 , where

- 1) V_1 are votes cast by honest voters who verified their votes;
- 2) V_2 is a subset of the votes from honest voters who did not verify their votes; it contains at most one vote for each such voter, but some votes may be dropped under the influence of the adversary;
- 3) V_3 represents the votes cast by the adversary, and its size should be bounded by the number of corrupt voters.

Point 1) can be strengthened by requiring that all of the verified votes should be present in the outcome, and not only those of honest voters. On the other hand, point 3) can be weakened by allowing the adversary to also cast votes for honest voters who did not verify their vote. This motivates our definition of the following sets associated to a trace:

- Ver^* - the multiset of verified votes satisfying the revoke policy;
- Ver° - as above, but only including votes of honest voters;
- Adv^* - the set of corrupt voter identities;
- Adv° - the set of corrupt voter identities for which the verification procedure has not been performed.
- $\text{Vote}^{\nabla, \diamond}$ - the set containing all subsets of the votes from *non-adversarial voters* who have not verified their votes; the definition of this set varies according to the scenario $(\nabla, \diamond) \in \{\bullet, \circ\} \times \{\bullet, \circ\}$ which is used for instantiating the notions of non-adversarial voters and verified votes.
- Res - the multiset of votes in the result of the election.

Definition 2. Let τ be a trace produced by a voting specification having Ω as revoke policy. We consider the formulas:

$$\begin{aligned}\text{Ver}^*(\text{id}, \text{cr}, v) &\equiv \text{Verified}(\text{id}, \text{cr}, v) \wedge \Omega(\text{id}, v) \\ \text{Ver}^\circ(\text{id}, \text{cr}, v) &\equiv \text{Verified}(\text{id}, \text{cr}, v) \wedge \Omega(\text{id}, v) \wedge \neg \text{Corr}(\text{id}) \\ \text{Adv}^*(\text{id}) &\equiv \text{Corr}(\text{id}) \\ \text{Adv}^\circ(\text{id}) &\equiv \text{Corr}(\text{id}) \vee \neg \text{Verified}(\text{id}, \text{cr}, v')\end{aligned}$$

and define the following sets, for $(\nabla, \diamond) \in \{\circ, \bullet\} \times \{\circ, \bullet\}$:

$$\begin{aligned}\text{Ver}^\nabla(\tau) &= \{\{v \mid \exists \text{id}, \text{cr}. \tau \models \text{Ver}^\nabla(\text{id}, \text{cr}, v) \wedge v \neq \perp\}\} \\ \text{Ver}_{\text{id}}^\nabla(\tau) &= \{\text{id} \mid \exists \text{id}', \text{cr}, v. \tau \models \text{Ver}^\nabla(\text{id}', \text{cr}, v) \\ &\quad \wedge \text{Reg}(\text{id}, \text{cr}) \wedge v \neq \perp\} \\ \text{Adv}^\diamond(\tau) &= \{\text{id} \mid \tau \models \text{Adv}^\diamond(\text{id})\} \\ \text{Res}(\tau) &= \{\{v \mid \exists \text{tr}. \tau \models \text{BBres}(\text{tr}, v) \wedge v \neq \perp\}\}\end{aligned}$$

Additionally, we let $\text{Vote}^{\nabla, \diamond}(\tau)$ to be the set of multisets $\{\{v_1, \dots, v_s\} \mid \text{s.t. there are distinct } \text{id}_1, \dots, \text{id}_s \text{ and for all } i \in \{1, \dots, s\} \text{ we have:}$

$$\tau \models \text{Vote}(\text{id}_i, v_i) \wedge \text{id}_i \notin (\text{Ver}_{\text{id}}^\nabla(\tau) \cup \text{Adv}^\diamond(\tau)).$$

Definition 3. Let $(\nabla, \diamond) \in \{\circ, \bullet\} \times \{\circ, \bullet\}$. A trace τ of a voting specification S satisfies $[\text{iv}_\nabla, \text{res}_\diamond]$ -election verifiability iff there exist multisets V_1, V_2, V_3 of votes such that $\text{Res}(\tau) = V_1 \uplus V_2 \uplus V_3$ and

- (1) $V_1 = \text{Ver}^\nabla(\tau)$,
- (2) $V_2 \in \text{Vote}^{\nabla, \diamond}(\tau)$,
- (3) $|V_3| \leq |\text{Adv}^\diamond(\tau) \setminus \text{Ver}_{\text{id}}^\nabla(\tau)|$.

We denote this by $\tau \models \text{E2E}[\text{iv}_\nabla, \text{res}_\diamond]$. If all traces of S satisfy it, we have $S \models \text{E2E}[\text{iv}_\nabla, \text{res}_\diamond]$.

Definition 3 gives four different notions of verifiability:

- $\text{E2E}[\text{iv}_\bullet, \text{res}_\bullet]$ is the strongest notion, ensuring that verified votes of corrupt voters are also included in the final tally (this is strong individual verifiability).
- $\text{E2E}[\text{iv}_\circ, \text{res}_\bullet]$ corresponds to the definition from [19], providing individual verifiability only for honest voters.

- $\text{E2E}[\text{iv}_\bullet, \text{res}_\circ]$: in addition to corrupt voters, this notion allows the adversary to cast votes for voters who have not successfully performed the verification procedure (this is sometimes called ballot stuffing). However, for corrupt voters who successfully verified their votes, we are certain these votes are correctly counted.
- $\text{E2E}[\text{iv}_\circ, \text{res}_\circ]$ is the weakest notion, only allowing individual verifiability for honest voters and allows ballot stuffing for honest voters who did not verify their votes.

B. Symbolic definition of election verifiability

We extend the symbolic verifiability definition from [6] by introducing two symbolic notions: one to make a stronger link between voter identities and their public credentials, and a second one to make a link between the public credentials and trackers used for verification. The former helps making a stronger connection to the general notion of election verifiability from Definition 3 and [6]. The latter helps modelling receipt-free voting systems like Selene.

The registration events. Universal verifiability of the tally guarantees that all the votes in the final outcome correspond to ballots recorded next to public credentials before the tally phase, i.e. we have:

$$\text{BBtally} : (\text{cr}_1, b_1), \dots, (\text{cr}_n, b_n) \quad \text{BBres} : (\text{tr}_1, v_1), \dots, (\text{tr}_n, v_n)$$

and we can ensure that v_1, \dots, v_n is a permutation of votes encoded in b_1, \dots, b_n . However, how can we ensure that b_1, \dots, b_n are ballots cast by eligible voters, i.e. that the adversary cannot do ballot stuffing? The election verifiability definition should precisely limit the influence of the adversary over the ballots in the final tally. The definition in [6] considers a notion of *corrupt credentials* in order to specify which ballots among b_1, \dots, b_n should be allowed to be cast by the adversary. This notion turns out to be distinct from the notion of *corrupt voters*, which is the more natural notion used in the general definition of election verifiability. As a consequence, the symbolic definition in [6] cannot be used to prove election verifiability. On the other hand, although it can find real attacks on verifiability in voting protocols, it can also detect false attacks. For example, if there is one corrupt voter, one corrupt registrar and the result contains one vote cast by the adversary, this should not be counted as an attack on verifiability, since the number of adversarial votes is bounded by the number of corrupt voters. However, the definition in [6] finds an attack against Belenios in this scenario, due to the fact that the corrupt registrar associates a so-called honest credential to the corrupt voter. We solve this problem by introducing events $\text{Reg}(\text{id}, \text{cr})$ that make a formal link between voter identities and the corresponding public credentials, and adding a consistency requirement for these events in the security definition.

Depending on the trust assumptions in the protocol and the corruption scenario for which verification is performed, the placement of the registration events may vary. For example, Belenios and BeleniosRF are supposed to provide verifiability if either the registrar or the voting server is honest. In the former case, we place $\text{Reg}(\text{id}, \text{cr})$ in the specification of the

registrar at registration time, and in the latter case we place it in the specification of the voting server when it receives a ballot from that voter. This way we can obtain a general symbolic definition that does not have to vary according to the corruption scenario, but only requires to place the registration facts according to it. In a previous framework for symbolic election verifiability [18], distinct definitions and assumptions were devised for distinct corruption scenarios in Belenios (corrupt server or corrupt registrar). We are able to obtain security proofs of election verifiability for both of these scenarios using the same definition and assumptions. Finally, we note that in general $\text{Reg}(\text{id}, \text{cr})$ may be recorded as an event without requiring trust assumptions, but this would require additional cryptographic techniques to hide the link between id and cr .

The linking events. To account for the fact that voters may verify their votes directly in the final tally, we need to further link their public credentials to the credentials used for verification, i.e. to their corresponding trackers. Since the tally process is universally verifiable, this link can be established without any trust assumptions, relying only on the guarantees provided by the zero-knowledge proofs. To formalise this link, we require the existence of a function that can extract a tracker from the public data that is associated to a credential if the corresponding zero-knowledge proofs are valid. Then, if u_1, \dots, u_n is all the public data associated to a credential cr and $\text{tr} = \text{extract}(u_1, \dots, u_n)$, we add an event $\text{Link}(\text{cr}, \text{tr})$ to record the fact that tr should be considered as associated to cr . The symbolic definition of election verifiability will ensure that this link is consistent, i.e. any credential has a unique associated tracker, and conversely. Definition 4 introduces the additional events that are required for registration and linking of credentials.

Definition 4. A voting specification S is enhanced if its rules record the following events:

- $\text{Reg}(\text{id}, \text{cr})$: a voter with identity id is associated with the public credential cr ;
- $\text{BBreg}(\text{cr})$: the public credential cr is recorded on BB as corresponding to an eligible voter;
- $\text{BBptr}(\text{cr}, u_1, \dots, u_n)$: records auxiliary information relevant for tracking votes for cr .

The event $\text{BBptr}(\text{cr}, u_1, \dots, u_n)$ puts together any pre-tracking data and zero-knowledge proofs related to a given credential cr . This allows for a simpler specification of the tracker extraction function in the following. For a sequence of terms u_1, \dots, u_n , we denote $\text{ver}(u_1, \dots, u_n) = \text{ok}$ if, for any term in u_1, \dots, u_n that represents a zero-knowledge proof, the corresponding verification equation evaluates to true. We assume, implicitly, that all arguments required for the verification of proofs are also part of u_1, \dots, u_n . If all proofs related to a credential are valid, an admissible extraction function should always succeed in computing a corresponding tracker from pre-tracking data, as formalised by the following definition.

Definition 5. An enhanced voting specification $(P, \Psi, \mathcal{E}, \Omega)$ has an admissible extraction function if there is a destructor extract defined in \mathcal{E} that, for any terms $\text{cr}, u_1, \dots, u_n$, satisfies:

$$[(P, \Psi, \mathcal{E}, \Omega) \models \text{BBptr}(\text{cr}, u_1, \dots, u_n) \wedge \text{ver}(u_1, \dots, u_n) = \text{ok}] \implies \text{extract}(u_1, \dots, u_n) \neq \perp.$$

The satisfaction of Definition 5 can be checked automatically with ProVerif, as we do for Selene and Hyperion. For Belenios and BeleniosRF, the extraction function is trivial, as it needs to associate cr to itself.

Example 6. In Selene, we have $\text{BBptr}(\text{cr}, \langle \text{tpk}, \text{etr}, \text{cm}, \pi_{\text{ptr}}, \text{pk} \rangle)$ where π_{ptr} shows, in particular, that etr is of the form $\text{enc}(\text{tr}, \text{pk}, r)$, for some terms tr and r . Therefore, we can define the extraction function by $\text{extract}(x_1, \text{enc}(x, y, z), x_2, x_3, y) = x$.

For Selene, the tracker is a term encoded inside a ciphertext in pre-tracking information, and the extraction function simply returns this term. In Hyperion, the extraction function needs to be more involved to account for the computation that trustees should perform in order to derive trackers.

Example 7. In Hyperion, we have $\text{BBptr}(\text{cr}, \langle \text{rtpk}, \pi \rangle)$, where $\pi = \text{prmix}(\text{M}_{\text{in}}, s, \text{M}_{\text{out}})$ shows that the set of final trackers is equal to the set of pre-trackers where each element rtpk is raised to the same power s . Therefore, we can define the extraction function by $\text{extract}(x, \text{prmix}(x_1, y, x_2)) = x^y$.

One may wonder if, in some sense, a wrong extraction function could be defined, that is still admissible for a protocol. Our symbolic definition of verifiability, if satisfied by a protocol and extraction function specification, will ensure that the extraction function determines a bijection between the set of credentials and the set of trackers, and furthermore that voters derive trackers that are consistent with this function. We show that these two properties (together with others from our definition) are sufficient to ensure end-to-end election verifiability - regardless of the actual definition of the extraction function. Another observation is that the extraction function is not efficiently computable in general, as can be seen from the previous examples. Indeed, as the corresponding tracker points directly to the vote in Selene and Hyperion, this would violate privacy. In Belenios and BeleniosRF this is not a problem, since the tracker points to the ballot, and not to the vote.

Definition 6. Let S be an enhanced voting specification with an admissible extraction function. We define the event $\text{Link}(\text{cr}, \text{tr})$ to record that tr can be extracted from pre-tracking information associated to cr . Formally, for a trace τ of S , we have $\tau \models \text{BBptr}(\text{cr}, u_1, \dots, u_n) \wedge \text{ver}(u_1, \dots, u_n) = \text{ok} \wedge \text{extract}(u_1, \dots, u_n) = \text{tr} \implies \tau \models \text{Link}(\text{cr}, \text{tr})$.

In ProVerif, the event Link satisfying the definition can be specified in a process that applies extract and adds Link as soon as the required information from BBptr is available. Definition 6 formalises the association between credentials and trackers that is ensured by any admissible extraction function. As discussed above, our symbolic definition of verifiability,

relying on trace formulas from Figure 4, will ensure additional constraints for the event Link.

Basic properties (for $\nabla, \diamond \in \{\circ, \bullet\}$)	
Φ_{iv1}^∇	$: \text{Ver}^\nabla(\text{id}, \text{cr}, v) \wedge \text{Link}(\text{cr}, \text{tr}) \wedge \text{BBres}(\text{tr}, v') \Rightarrow v = v'$
Φ_{iv2}^∇	$: \text{Ver}^\nabla(\text{id}, \text{cr}, v) \wedge \text{Ver}^\nabla(\text{id}', \text{cr}, v') \Rightarrow \text{id} = \text{id}'$
Φ_{iv3}^∇	$: \text{Ver}^\nabla(\text{id}, \text{cr}, v) \Rightarrow \text{Link}(\text{cr}, \text{tr}) \wedge \text{BBreg}(\text{cr})$
Φ_{eli}	$: \text{BBres}(\text{tr}, v)$ $\Rightarrow \text{Link}(\text{cr}, \text{tr}) \wedge \text{BBreg}(\text{cr}) \wedge (\text{Reg}(\text{id}, \text{cr}) \vee v = \perp)$
Φ_{res}^\diamond	$: \text{BBres}(\text{tr}, v) \wedge \text{Link}(\text{cr}, \text{tr}) \wedge \text{Reg}(\text{id}, \text{cr})$ $\Rightarrow \text{Vote}(\text{id}, v) \vee \text{Adv}^\diamond(\text{id}) \vee v = \perp$
(formulas Ver^∇ and Adv^\diamond introduced in Definition 2)	
Consistency properties	
Φ_{reg1}	$: \text{Reg}(\text{id}, \text{cr}) \wedge \text{Reg}(\text{id}', \text{cr}) \Rightarrow \text{id} = \text{id}'$
Φ_{reg2}	$: \text{Reg}(\text{id}, \text{cr}) \wedge \text{Reg}(\text{id}, \text{cr}') \Rightarrow \text{cr} = \text{cr}'$
Φ_{link1}	$: \text{Link}(\text{cr}, \text{tr}) \wedge \text{Link}(\text{cr}', \text{tr}) \Rightarrow \text{cr} = \text{cr}'$
Φ_{link2}	$: \text{Link}(\text{cr}, \text{tr}) \wedge \text{Link}(\text{cr}, \text{tr}') \Rightarrow \text{tr} = \text{tr}'$
Φ_{one}	$: \text{BBres}(\text{tr}, v) @ i \wedge \text{BBres}(\text{tr}, v') @ j \Rightarrow i = j$

Fig. 4. Formulas for symbolic election verifiability (Definition 7).

Definition 7. Consider the formulas from Figure 4. For $(\nabla, \diamond) \in \{\circ, \bullet\} \times \{\circ, \bullet\}$, we let

$$\text{SE2E}[\text{iv}_\nabla, \text{res}_\diamond] = \Phi_{iv1}^\nabla \wedge \Phi_{iv2}^\nabla \wedge \Phi_{iv3}^\nabla \wedge \Phi_{eli} \wedge \Phi_{res}^\diamond \wedge \Phi_{reg1} \wedge \Phi_{reg2} \wedge \Phi_{link1} \wedge \Phi_{link2} \wedge \Phi_{one}$$

An enhanced e-voting specification S satisfies $[\text{iv}_\nabla, \text{res}_\diamond]$ -symbolic election verifiability iff it has an admissible extraction function and $S \models \text{SE2E}[\text{iv}_\nabla, \text{res}_\diamond]$.

The formula Φ_{iv1}^∇ represents individual verifiability, ensuring that all verified votes should be part of the final outcome, i.e. the vote recorded for that tracker is the expected one. The formula Φ_{iv2}^∇ requires that no clash should occur on the public credentials of voters who verified their vote. The formula Φ_{iv3}^∇ makes a link from the voter credential to a vote verification tracker tr . The formula Φ_{eli} ensures that trackers on BBres are connected through Link to registered eligible credentials. The formula Φ_{res}^\diamond ensures that there is no ballot stuffing. It relies on Link and Reg events to determine what is the voter identity associated to each tracker in the outcome. Our constraints related to Link and Reg are completed by the consistency properties, ensuring that each credential has a unique associated tracker and voter identity, and conversely. Finally, the property Φ_{one} ensures that there is at most one vote counted for each tracker, and therefore for each credential and each voter.

C. Symbolic E2E verifiability implies general E2E verifiability

In this section we prove that, if an enhanced voting specification satisfies $\text{SE2E}[\text{iv}_\nabla, \text{res}_\diamond]$, then it also satisfies $\text{E2E}[\text{iv}_\nabla, \text{res}_\diamond]$. We recall the general flow of electronic voting protocols that we consider in this paper: a set of public credentials is determined at registration and published on

BBreg ; ballots from corresponding voters are collected, and the ballot to be tallied for each credential is recorded on BBtally ; the final result is computed from the ballots in BBtally . Assuming there are n public credentials, a snapshot of the end-to-end run of the election can then be represented by:

$$\begin{aligned} \text{BBreg} &: \quad \text{cr}_1 \quad \dots \quad \text{cr}_n \\ \text{BBtally} &: (\text{cr}_1, \text{b}_1) \quad \dots \quad (\text{cr}_n, \text{b}_n) \\ \text{BBres} &: (\text{tr}_1, \text{v}_1) \quad \dots \quad (\text{tr}_n, \text{v}_n) \end{aligned}$$

where $\text{b}_i = \perp$ and $\text{v}_i = \perp$ in case of abstention, invalid ballot casting or ballot blocking by the adversary. Recall that in Helios, Belenios and BeleniosRF, BBres is defined implicitly from BBtally taking $\text{tr}_i = \text{cr}_i$ and $\text{v}_i = \text{open}(\text{b}_i)$. In order to prove our result, we will make one additional assumption about the execution trace of a voting system, namely that it contains the final state of the bulletin board. According to our observations above, the fact that a trace contains the final state of the bulletin board can be formalised by the formula

$$\Psi_{\text{E2E}} : |\{\text{cr} \mid \tau \models \text{BBreg}(\text{cr})\}| = |\{\text{tr} \mid \exists v. \tau \models \text{BBres}(\text{tr}, v)\}|$$

stating that the number of trackers (or credentials) recorded on the final result bulletin board is equal to the number of registered credentials. A trace that satisfied Ψ_{E2E} is called an *end-to-end voting trace*. We note that the assumption of a complete trace is inherent to the property of election verifiability, which cannot be evaluated until the election authorities commit to the final result.

Theorem 1. For every voting specification S , symbolic election verifiability implies general election verifiability for end-to-end voting traces. That is, for any trace τ of S , we have

$$\tau \models \text{SE2E}[\text{iv}_\nabla, \text{res}_\diamond] \wedge \tau \models \Psi_{\text{E2E}} \Rightarrow \tau \models \text{E2E}[\text{iv}_\nabla, \text{res}_\diamond]$$

The proof is given in Section A of the appendix.

V. VERIFIABILITY ANALYSIS OF BELENIOSRF AND SELENE

In this section, we present the verifiability analysis with ProVerif of BeleniosRF, Selene, SeleneRF, and Hyperion, which are introduced in Section III, with respect to the definition introduced in Section IV. We start by introducing the common components used for the formalisation of these protocols (voting server, registrar, voting platform, etc.) and the corruption models that we consider. We also present the equational theories for cryptographic primitives. Then, we provide comparative verification results for these systems obtained with ProVerif, and explain the attacks and proofs found. We show some parts from the specification of each protocol and refer to the ProVerif code attached as supplementary material for full details.

A. Common protocol structure

We consider the following parties common to all protocols:

- *Administrator A* determines the list of candidates $\text{v}_1, \dots, \text{v}_k$ and voters $\text{id}_1, \dots, \text{id}_n$.
- *Trustees T* generate the election keys (pk, sk) and publish pk on BB .
- *Registrar VR* generates a signing key pair (spk, ssk) for each

voter and publishes $cr = spk$ on BBreg. While Selene and its variants do not explicitly mention a registrar, they assume that voters are registered with public keys before the start of the protocol. We consider that this registration is performed relying on a registrar. The case when voters register themselves is equivalent to the case of an honest registrar (if they use an honest device) or a corrupt registrar (if they use a corrupt one).

- *Voting server* VS is responsible for publishing authenticated valid ballots on BBcast in addition to generating a password pwd for each voter. The ballots to be tallied are put on BBtally.
- *Voters* V use their voting platform VP to construct and cast a ballot; they verify their ballot or vote on BB.
- *Voting platform* VP constructs a ballot; encrypting the desired vote, signing it and attaching a proof, and sends it to VS with authentication information.

Setup. Parties generate election parameters as described above, with public elements recorded in corresponding bulletin board tables: BBkey for the public key of the election and BBreg for public voter credentials. Each voter records the credential tuple (id, pwd, cr, ssk) .

Voting. Ballots are constructed by VP and posted to VS . VP takes as input the vote v and voter credentials. The ballot b has three parts:

$$\begin{aligned} c &= \text{enc}(v, pk, r) \quad (\text{encryption of } v), \\ s &= \text{sign}(c, ssk) \quad (\text{signature}), \\ \pi &= \text{pr}(c, v, r, \ell) \quad (\text{zk proof of knowledge of } v, r). \end{aligned}$$

We denote the operation of generating r and constructing b by $\text{Gen}(v, pk, ssk)$. In BeleniosRF, the label ℓ in π is equal to cr and b is rerandomized before being published; in Selene and Hyperion, ℓ is empty. Authentication is modelled relying on a hash function and a term $h(id, pwd, cr, b)$, which allows VS to verify that b was sent by a voter knowing the pwd for id . We note that this is an abstraction of any underlying authentication process, which allows the server to validate the identity of the voter sending the ballot.

Tally and verification. The zk proofs and signatures are verified before tally, which is done by verifiably decrypting the ballots from BB, with additional procedures for deriving trackers in Selene and Hyperion. Individual verification procedures are specific in each protocol, as shown later.

B. Corruption models

Other than the administrator, any party may be corrupted:

- Corrupt trustees: \mathcal{A} can choose the secret key and perform any operations assigned to trustees in the corresponding protocol.
- Corrupt registrar: \mathcal{A} is allowed to choose the public and private credentials for each voter.
- Corrupt server: \mathcal{A} chooses voter passwords and accepts any ballot without authentication checks.
- Corrupt voter: all credentials are leaked to \mathcal{A} .
- Corrupt platform: reveals to \mathcal{A} the voter credentials (except the trapdoor secret key in Selene and Hyperion), and \mathcal{A} chooses the ballot that the voter will cast.

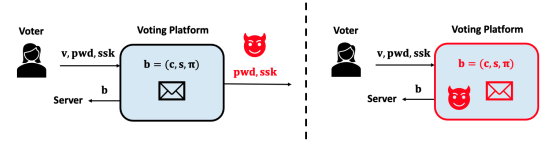


Fig. 5. Corrupt voter (left) vs corrupt platform (right).

```

let VotingPlatform(v, pk) =
  get Cred(id, cr, ssk) in
  get Pwd(= id, pwd) in
  let b = Gen(v, pk, ssk) in
  let a = h(id, pwd, cr, b) in
  event Vote(id, cr, v);
  insert Voted(id, cr, v);
  out(pub, (id, cr, b, a)).

let VotingServer(pk) =
  get BBreg(cr) in
  get Pwd(id, pwd) in
  in(pub, (= id, = cr, b, a);
  if h(id, pwd, cr, b) = a then
    if Verify(b, cr) = ok then
      let b' = Rand(b, pk) in
      insert BBcast(cr, b').

let VoterVer =
  get Voted(id, cr, v); BBtally(= cr, b') in
  if Verify(b', cr) = ok then event Verified(id, cr, v).

```

Fig. 6. Voting and verification processes for BeleniosRF.

Figure 5 illustrates the difference between a corrupt V and a corrupt VP . In the latter, we let the adversary construct the ballot as it wishes, providing \mathcal{A} with credentials required to do so. Our notion of a corrupt VP is much stronger than the one considered in [6]: their model only allows \mathcal{A} to choose the randomness used for the encryption; our model allows \mathcal{A} to construct the whole ballot. Note that, in addition to voting credentials (id, pwd, cr, ssk) , a voting protocol may provide voters with separate secret material required for verification. This is the case in Selene and Hyperion, where voters have an additional trapdoor secret key required to derive trackers. In the case of a corrupt VP , contrary to that of a corrupt V , we assume these credentials are kept secret. In practice, this means that the verification device should be different from the voting platform, or that trapdoor keys are generated and stored on special trusted hardware. A corruption model (or scenario) is defined by a subset of the corruption abilities described above. In Table I and Table II, we list all the considered corruption models and verification results found with ProVerif. In all cases, the communication network and voters can be corrupt. The ProVerif code is available online [1] and the running time for each scenario is given in its corresponding file. For BeleniosRF, we consider the trivial extraction function that associates cr to itself; for Selene and Hyperion, we consider extraction functions like in Example 6 and Example 7. For all cases, we consider the revote policy that counts the last vote cast by each voter.

C. BeleniosRF specification and verification

To model the cryptographic primitives of BeleniosRF, in addition to equations from Example 2, we consider the equations (4)-(6) below to model the properties that allow ciphertext rerandomization (4) and the adaptation of the corresponding signature (5) and zk proof (6). Equation (7) models the verification of a proof of correct decryption, provided by




Fig. 7. Clash attack by corrupt registrar in BeleniosRF.

trustees in all protocols. Part of the specification for BeleniosRF is presented in Figure 6. It is similar to Example 1, the main difference being the randomization of the ballot b into $b' = \langle c', s', \pi' \rangle$ before casting on BB. We denote this operation, together with the generation of r' , by $\text{Rand}(b, pk)$. The voter verifies b' on BB by verifying the signature s' and the zero-knowledge proof π' ; we denote the verification result by $\text{Verify}(b, cr)$.

$$\begin{aligned} c' &= \text{renc}(c, pk, r') && (\text{re-encryption}), \\ s' &= \text{resign}(s, pk, r') && (\text{adaptation of signature}), \\ \pi' &= \text{repr}(pr, pk, r', cr) && (\text{adaptation of proof}). \end{aligned}$$

- (4) $\text{renc}(\text{enc}(x, y, z_1), y, z_2) = \text{enc}(x, y, z_2)$
- (5) $\text{resign}(\text{sign}(\text{enc}(x, y, z_1), w), y, z_2) = \text{sign}(\text{enc}(x, y, z_2), w)$
- (6) $\text{repr}(pr(\text{enc}(x, y, z_1), x, z_1, w), y, z_2, w)) = pr(\text{enc}(x, y, z_2), x, z_2, w)$
- (7) $\text{ver}'(pr'(\text{enc}(x, pk(y), z), x, y), \text{enc}(x, pk(y), z), x, pk(y))) = \text{ok}$

Attacks. We find an attack against individual verifiability in the case of a corrupt registrar, i.e. the scenario \mathcal{A}_3 . The weakness comes from the fact that individual verification does not directly check the ballot posted by voter, but its randomization. The properties Φ_{iv1}^∇ and Φ_{iv2}^∇ in $\text{E2E}[iv_\nabla, res_o]$ are violated for $(\nabla, \diamond) \in \{\circ, \bullet\} \times \{\circ, \bullet\}$, because of a clash attack as illustrated in Figure 7. It shows a corrupt registrar communicating the same signature key pair $\langle cr, ssk \rangle$ to two honest voters id_1 and id_2 . Assume the two voters cast their ballots separately relying on their credential cr . A ballot b from one of the voters is accepted, randomized, and published on BB. For the second voter, we assume the ballot is dropped by the network, before the voter noticing it. When verifying their ballots (possibly later when the voting phase is finished), both voters can verify that, for the tuple (cr, b') on BB, b' is the rerandomization of a ballot submitted with their credential. They both verify the attached signature and zk proof since they share the same credentials. However, only one vote will be counted for these two voters. A similar attack is possible in the case of a corrupt voter, even when all other parties are honest (these are the two lines filled with \times for the case of iv_\circ in Table I). It is sufficient for the communication network to block the ballot of that voter, and for the adversary to cast a different ballot with the credential of that voter. In Belenios this does not lead to an attack, but in BeleniosRF it does.

Related attacks. BeleniosVS [18] is a variant of Belenios that uses the same ballot rerandomization feature as BeleniosRF, in order to achieve better privacy in the case of a dishonest voting platform. In consequence, a similar attack as ours applies to

TABLE I
VERIFICATION RESULTS FOR BELENIOS AND BELENIOSRF

Corruption Models		\mathcal{A}_1	\mathcal{A}_2	\mathcal{A}_3	\mathcal{A}_4
Trustees		H	C	C	C
Registrar		H	H	C	H
Server		H	H	H	C
Voting Platform		H	H	H	H
Belenios*	$\text{E2E}[iv_\bullet, res_\bullet]$	✓	✓	✓	✓
	$\text{E2E}[iv_\circ, res_\bullet]$	✓	✓	✓	✓
	$\text{E2E}[iv_\bullet, res_\circ]$	✓	✓	✓	✓
	$\text{E2E}[iv_\circ, res_\circ]$	✓	✓	✓	✓
BeleniosRF†	$\text{E2E}[iv_\bullet, res_\bullet]$	✗	✗	✗	✗
	$\text{E2E}[iv_\circ, res_\bullet]$	✓	✓	✗	✓
	$\text{E2E}[iv_\bullet, res_\circ]$	✗	✗	✗	✗
	$\text{E2E}[iv_\circ, res_\circ]$	✓	✓	✗	✓

*: no revoting or voter verification in the tally phase

†: no revoting

BeleniosVS in the case of a corrupted registrar, and the authors of [18] have already observed this. As a fix, they propose that the voting server sends an acknowledgement to the voter when the ballot is successfully cast. In an attack scenario similar to the one above (from Figure 7), this should counter the dropping of the second ballot (b_2) by the network, since the voter id_2 would notice that no acknowledgement was sent. We plan to investigate in future work if a similar fix could help in BeleniosRF, and if it still protects against a stronger adversarial model in BeleniosVS. Indeed, in our analysis, we consider a strong attacker model that controls the secure channel that is established between the voting platform and the voting server. A simple acknowledgement would not help against such an attacker, since the attacker could impersonate the voting server and spoof the acknowledgement.

Proofs. BeleniosRF satisfies election verifiability, i.e. $\text{E2E}[iv_\diamond, res_\diamond]$ for $\diamond \in \{\circ, \bullet\}$, when the registrar, the voter and the voting platform are honest. In this case, the trustees and the voting server can be fully corrupt. All attacks and proofs for Belenios and BeleniosRF are found automatically with ProVerif within several seconds.

D. Selene and Hyperion specification and verification

To model the cryptographic primitives of Selene, in addition to equations from Example 2, and equations (4) and (7) from above, we consider equations for additional zk proofs, for trapdoor commitments and for deriving a commitment from a ciphertext. We omit from the description the equations for zk proofs, as they are similar to (3) and (7), where terms are modified accordingly to account for the desired relation. To model trapdoor commitments, we use a theory previously considered in [11], [25] for verifying privacy-type properties. The main property is that a commitment can be opened with the corresponding trapdoor secret key, as in equation (8) below. Furthermore, equation (9) models the fact that a trapdoor secret key along with the randomness from the commitment can be used to generate randomness that opens that commitment to

Selene	Hyperion
Voter(id, cr) generates : Td(id, tpk, ts)	Voter(id, cr) generates : Td(id, tpk, ts) where $tpk = g^{tsk}$
Trustees generate : BBptr(cr, tpk, enc(tr, pk, r), com(tr, tpk, r_T), pr_{ptr} , pk) BBres(tr, v), Com(cr, r_T)	Trustees generate : BBptr(cr, tpk', pr_{exp}) BBres(tpk ^s , v), Exp(cr, g ^s)
let VoterVerification = get Voted(id, cr, v) BBptr(= cr, x_1, x_2, y, x_3, x_4) BBres(tr, = v), Com(= cr, r_T) Td(id, tpk, ts) in if open(y, ts, r_T) = tr then event Verified(id, cr, v).	let VoterVerification = get Voted(id, cr, v) BBres(tr, = v) Exp(= cr, g ^s) Td(id, tpk, ts) in if (g ^s) ^{tsk} = tr then event Verified(id, cr, v).

Fig. 8. Tracker generation and voter verification in Selene and Hyperion.

any desired value. As shown in [11], two more equations need to be added in order to make this theory convergent for use in ProVerif, which we will omit here. Equation (10) models a relation between ElGamal and Pedersen commitments, which allows trustees to compute a commitment to a tracker from an encryption of that tracker, without decrypting it. This operation relies on homomorphically combining the encryption of a tracker x with the encryption of a term w^z , where w will play the role of the public trapdoor key, and z will play the role of the randomness in the commitment. The corresponding secret decryption key is needed to obtain the commitment.

- (8) $\text{open}(\text{com}(x, \text{pk}(y), z), y, z) = x$
- (9) $\text{com}(x_2, \text{pk}(y), \text{fake}(x_1, x_2, y, z)) = \text{com}(x_1, \text{pk}(y), z)$
- (10) $\text{ecom}(\text{enc}(x, \text{pk}(y), r), \text{enc}(w^z, \text{pk}(y), r'), y) = \text{com}(x, w, z)$

TABLE II
VERIFICATION RESULTS FOR SELENE, SELENERF AND HYPERION

Corruption Models	\mathcal{A}'_1	\mathcal{A}'_2	\mathcal{A}'_3	\mathcal{A}'_4	\mathcal{A}'_5	\mathcal{A}'_6	\mathcal{A}'_7
Trustees	H	C	C	C	C	C	C
Registrar	H	H	C	H	H	C	H
Server	H	H	H	C	H	H	C
Voting Platform	H	H	H	H	C	C	C
E2E[iv _o , res _o]	✓	✗	✗	✗	✗	✗	✗
E2E[iv _o , res _o]	✓	✓	✓	✓	✗	✗	✗
E2E[iv _o , res _o]	✓	✗	✗	✗	✗	✗	✗
E2E[iv _o , res _o]	✓	✓	✓	✓	✓	✓	✓

Part of the formal specification of parties in Selene is given in Figure 8. We use tables for storing voter credentials and keys, and also for the private channel Com used by the trustees to communicate commitment randomness to voters. Trackers are generated by trustees, encrypted with the public key of the election and transformed to derive the corresponding commitments for the public trapdoor key of voters, relying on equation (10). All this pre-tracking information is then stored in BBptr. Along with the ballot, it is used by trustees for tally and by voters to derive their trackers, after receiving the corresponding randomness from trustees. We assume trustees publish zk proofs of correctness for their computation and that these are verified for all credentials and ballots tallied.

Attacks. A trust assumption in Selene is that the voters performing verification are honest. On the other hand, the trustees can be fully corrupt. If both the verifying voter and the trustees are corrupt, there is a trivial attack against election verifiability. The main idea is that trustees can use the trapdoor secret key of the voter to compute fake randomness that opens their commitment to a different tracker. If that tracker points to the same vote as chosen by the voter, this results in a clash attack. This scenario is formally represented by the properties iv_{\bullet} (corrupt voters) and attackers $\mathcal{A}'_2 - \mathcal{A}'_7$ (corrupt trustees) in Table II. Another notable attack is ballot stuffing (i.e. attack against iv_{\bullet}) when the voting platform is corrupt. Indeed, in that case, the adversary can replace the vote of the voter with any other vote, as soon as the voter does not verify their vote. This shows that the tracking mechanism in Selene protects voters against corrupt platforms only if they verify their votes, as expected.

Proofs. Notable cases are the positive results in corruption scenarios $\mathcal{A}'_3, \mathcal{A}'_4$ and $\mathcal{A}'_6, \mathcal{A}'_7$. Assuming that the voting platform is honest, the former show that the (standard) verifiability property $E2E[iv_{\bullet}, res_{\bullet}]$ is true if either the registrar or the voting server is honest, like in Belenios. Furthermore, if the voting platform is corrupt, the results for $\mathcal{A}'_6, \mathcal{A}'_7$ show that the weaker notion $E2E[iv_{\bullet}, res_{\bullet}]$ still holds, meaning that ballot stuffing is possible only if voters have not verified their votes.

E. Hyperion and SeleneRF

To model the cryptographic primitives of Hyperion, one set of new equations that needs to be added are for handling the specific zk proofs, related to exponentiation instead of commitment and encryption. We also need to consider a model of exponentiation that is more general than the usual model used in ProVerif, which only allows to capture Diffie-Hellman-like interactions with two exponentiations. We need three exponentiations, since trackers for voters are generated through: $tr_i = ((g^{tsk_i})^{r_i})^s$. For this, we consider the equations (11) and (12) from below. We note that this theory is still incomplete, as in general one needs to cover any number of exponentiations in order to reason about security, and not only to be able to execute the protocol.

$$\begin{aligned}
 (11) \quad (g^x)^y &= (g^y)^x & (13) \quad \text{verify}(\text{prexp}(x, y, x^y), x, x^y) &= \text{ok} \\
 (12) \quad (g^x)^y)^z &= (g^y)^z)^x & (14) \quad \text{extract}(x, \text{prexp}(x, y, z)) &= x^y
 \end{aligned}$$

In order to add the event Link(cr, tr), according to Definition 6, Definition 5 and Example 7, we need to verify an exponentiation mixnet proof. Since we consider an unbounded number of voters, the verification of such proofs cannot be directly modelled in ProVerif. Instead, we consider an abstraction of mixnets that takes each tracker individually, exponentiates it to a secret power s , and outputs the resulting tracker together with a proof of exponentiation, whose verification is modelled by equation (13). Then, we can rely on the equation (14) for extraction. This abstraction of mixnets is sufficient in our case since we are not considering privacy properties. In fact, this model represents a weaker guarantee than a mixnet, which is still sufficient for election verifiability.

of Selene and Hyperion. This (malicious) mixnet is allowed to drop some of its inputs (these inputs will be considered as dropped ballots by the definition). What is important for the verifiability proof, is that the elements that do get output by the mixnet are in the prescribed relation with respect to the corresponding input elements. This is ensured by equation (13) in case of an exponentiation mixnet, and a similar equation for the case of a re-encryption mixnet. We also specify SeleneRF, that combines the models for ballot randomization from BeleniosRF and the tracking mechanism in Selene. For these models of Hyperion and SeleneRF, we obtain the same analysis results as for Selene.

VI. CONCLUSION AND FUTURE WORK

We have made several contributions related to automated verification of election verifiability in electronic voting protocols: we have proposed a new symbolic definition improving on the state-of-the-art, we have applied it to verify several protocols that were not verified before, and we have found new attacks and security proofs as a result. Our symbolic definition allows to cover many corruption scenarios for various protocols in a generic way, can be applied to receipt-free protocols and we have proved that it implies a general, realistic notion of election verifiability. We have found new attacks on BeleniosRF in scenarios that were previously thought to be secure and performed automated security proofs for Selene and its variants.

In future work, we plan to complement our analysis of receipt-free systems by verifying the property of receipt-freeness, aiming for a global security analysis including both verifiability and receipt-freeness. We have seen that one of our attacks on BeleniosRF is very similar to an attack against BeleniosVS that was discussed [18], which also propose a simple countermeasure based on an acknowledgement message from the voting server. We plan to apply our framework for verifying if a refinement of this countermeasure would be sufficient against the stronger network attacker that we consider in this paper.

ACKNOWLEDGMENT

This work was supported by the Luxembourg National Research Fund (FNR) and the Research Council of Norway for the joint INTER project SURCVS, ref. 11747298.

REFERENCES

- [1] Additional material: specifications in ProVerif. <https://github.com/sbaloglu/proverif-codes>.
- [2] Belenios - Verifiable online voting system. <https://belenios.org/>.
- [3] Helios - Verifiable online elections. <https://heliosvoting.org/>.
- [4] Ben Adida. Helios: Web-based open-audit voting. In *17th USENIX Security Symposium, July 28-August 1, 2008, San Jose, CA, USA*, pages 335–348, 2008.
- [5] Michael Backes, Cătălin Hrițcu, and Matteo Maffei. Automated verification of remote electronic voting protocols in the applied pi-calculus. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium, Pittsburgh, Pennsylvania, USA, 23-25 June 2008*, pages 195–209. IEEE Computer Society, 2008.
- [6] Sevdnur Baloglu, Sergiu Bursuc, Sjouke Mauw, and Jun Pang. Election verifiability revisited: Automated security proofs and attacks on Helios and Belenios. In *34th IEEE Computer Security Foundations Symposium, CSF 2021, Dubrovnik, Croatia, June 21-25, 2021*, pages 1–15. IEEE, 2021.
- [7] Josh Benaloh. Simple verifiable elections. In Dan S. Wallach and Ronald L. Rivest, editors, *2006 USENIX/ACCURATE Electronic Voting Technology Workshop, EVT'06, Vancouver, BC, Canada, August 1, 2006*. USENIX Association, 2006.
- [8] David Bernhard, Véronique Cortier, David Galindo, Olivier Pereira, and Bogdan Warinschi. SoK: A comprehensive analysis of game-based ballot privacy definitions. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, pages 499–516. IEEE Computer Society, 2015.
- [9] Bruno Blanchet. Modeling and verifying security protocols with the applied pi calculus and ProVerif. *Foundations and Trends in Privacy and Security*, 1(1-2):1–135, 2016.
- [10] Olivier Blazy, Georg Fuchsbauer, David Pointcheval, and Damien Vergnaud. Signatures on randomizable ciphertexts. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *Public Key Cryptography - PKC 2011 - 14th International Conference on Practice and Theory in Public Key Cryptography, Taormina, Italy, March 6-9, 2011. Proceedings*, volume 6571 of *Lecture Notes in Computer Science*, pages 403–422. Springer, 2011.
- [11] Alessandro Bruni, Eva Drewsen, and Carsten Schürmann. Towards a mechanized proof of Selene receipt-freeness and vote-privacy. In Robert Krimmer, Melanie Volkamer, Nadja Braun Binder, Norbert Kersting, Olivier Pereira, and Carsten Schürmann, editors, *Electronic Voting - Second International Joint Conference, E-Vote-ID 2017, Bregenz, Austria, October 24-27, 2017, Proceedings*, volume 10615 of *Lecture Notes in Computer Science*, pages 110–126. Springer, 2017.
- [12] Pyrros Chaidos, Véronique Cortier, Georg Fuchsbauer, and David Galindo. BeleniosRF: A non-interactive receipt-free electronic voting scheme. In *23rd ACM Conference on Computer and Communications Security (CCS'16)*, pages 1614–1625, Vienna, Austria, October 2016. ACM.
- [13] David L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2):84–90, feb 1981.
- [14] Benoît Chevallier-Mames, Pierre-Alain Fouque, David Pointcheval, Julien Stern, and Jacques Traoré. On some incompatible properties of voting schemes. In David Chaum, Markus Jakobsson, Ronald L. Rivest, Peter Y. A. Ryan, Josh Benaloh, Mirosław Kutylowski, and Ben Adida, editors, *Towards Trustworthy Elections, New Directions in Electronic Voting*, volume 6000 of *Lecture Notes in Computer Science*, pages 191–199. Springer, 2010.
- [15] Michael R. Clarkson, Stephen Chong, and Andrew C. Myers. Civitas: Toward a secure voting system. In *2008 IEEE Symposium on Security and Privacy (S&P 2008), 18-21 May 2008, Oakland, California, USA*, pages 354–368, 2008.
- [16] Véronique Cortier, Constantin Cătălin Drăgan, François Dupressoir, Benedikt Schmidt, Pierre-Yves Strub, and Bogdan Warinschi. Machine-checked proofs of privacy for electronic voting protocols. In *IEEE Symposium on Security and Privacy*, pages 993–1008, 2017.
- [17] Véronique Cortier, Constantin Cătălin Drăgan, François Dupressoir, and Bogdan Warinschi. Machine-checked proofs for electronic voting: Privacy and verifiability for Belenios. In *31st IEEE Computer Security Foundations Symposium, Oxford, United Kingdom, July 9-12, 2018*, pages 298–312. IEEE Computer Society, 2018.
- [18] Véronique Cortier, Alicia Filipiak, and Joseph Lallemand. BeleniosVS: Secrecy and verifiability against a corrupted voting device. In *32nd IEEE Computer Security Foundations Symposium, Hoboken, NJ, USA, June 25-28, 2019*, pages 367–381, 2019.
- [19] Véronique Cortier, David Galindo, Stéphane Glondou, and Malika Izabachene. Election verifiability for Helios under weaker trust assumptions. In *ESORICS - European Symposium on Research in Computer Security*, pages 327–344. Springer International Publishing, 2014.
- [20] Véronique Cortier, David Galindo, Ralf Küsters, Johannes Müller, and Tomasz Truderung. SoK: Verifiability notions for e-voting protocols. In *IEEE Symposium on Security and Privacy, SP 2016, San Jose, CA, USA, May 22-26, 2016*, pages 779–798. IEEE Computer Society, 2016.
- [21] Véronique Cortier, Pierrick Gaudry, and Stéphane Glondou. Belenios: A simple private and verifiable electronic voting system. In Joshua D. Guttman, Carl E. Landwehr, José Meseguer, and Dusko Pavlovic, editors, *Foundations of Security, Protocols, and Equational Reasoning - Essays*

Dedicated to Catherine A. Meadows, volume 11565 of *Lecture Notes in Computer Science*, pages 214–238. Springer, 2019.

- [22] Véronique Cortier, Niklas Grimm, Joseph Lallemand, and Matteo Maffei. A type system for privacy properties. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, pages 409–423. ACM, 2017.
- [23] Véronique Cortier, Joseph Lallemand, and Bogdan Warinschi. Fifty shades of ballot privacy: Privacy against a malicious board. In *33rd IEEE Computer Security Foundations Symposium, CSF 2020, Boston, MA, USA, June 22-26, 2020*, pages 17–32. IEEE, 2020.
- [24] Véronique Cortier and Ben Smyth. Attacking and fixing Helios: An analysis of ballot secrecy. In *Proceedings of the 24th IEEE Computer Security Foundations Symposium, CSF 2011, Cernay-la-Ville, France, 27-29 June, 2011*, pages 297–311. IEEE Computer Society, 2011.
- [25] Stéphanie Delaune, Steve Kremer, and Mark Ryan. Verifying privacy-type properties of electronic voting protocols. *J. Comput. Secur.*, 17(4):435–487, 2009.
- [26] Jens Groth and Amit Sahai. Efficient non-interactive proof systems for bilinear groups. In Nigel P. Smart, editor, *Advances in Cryptology - EUROCRYPT 2008, 27th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Istanbul, Turkey, April 13-17, 2008. Proceedings*, volume 4965 of *Lecture Notes in Computer Science*, pages 415–432. Springer, 2008.
- [27] Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-resistant electronic elections. In *Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society, WPES 2005, Alexandria, VA, USA, November 7, 2005*, pages 61–70, 2005.
- [28] Steve Kremer, Mark Ryan, and Ben Smyth. Election verifiability in electronic voting protocols. In Dimitris Gritzalis, Bart Preneel, and Marianthi Theoharidou, editors, *Computer Security - ESORICS 2010, 15th European Symposium on Research in Computer Security, Athens, Greece, September 20-22, 2010. Proceedings*, volume 6345 of *Lecture Notes in Computer Science*, pages 389–404. Springer, 2010.
- [29] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Verifiability, privacy, and coercion-resistance: New insights from a case study. In *32nd IEEE Symposium on Security and Privacy, S&P 2011, 22-25 May 2011, Berkeley, California, USA*, pages 538–553. IEEE Computer Society, 2011.
- [30] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Clash attacks on the verifiability of e-voting systems. In *Proceedings of the 33rd IEEE Symposium on Security and Privacy, 21-23 May 2012, San Francisco, California, USA*, pages 395–409. IEEE Computer Society, 2012.
- [31] Karola Marky, Oksana Kulyk, Karen Renaud, and Melanie Volkamer. *What did I really vote for? On the usability of verifiable e-voting schemes*, page 1–13. Association for Computing Machinery, New York, NY, USA, 2018.
- [32] C. Andrew Neff. A verifiable secret shuffle and its application to e-voting. In *CCS 2001, ACM Conference on Computer and Communications Security*, pages 116–125, 2001.
- [33] Olivier Pereira and Dan S. Wallach. Clash attacks and the STAR-Vote system. In Robert Krimmer, Melanie Volkamer, Nadja Braun Binder, Norbert Kersting, Olivier Pereira, and Carsten Schürmann, editors, *Electronic Voting - Second International Joint Conference, E-Vote-ID 2017, Bregenz, Austria, October 24-27, 2017. Proceedings*, volume 10615 of *Lecture Notes in Computer Science*, pages 228–247. Springer, 2017.
- [34] Peter Y. A. Ryan, Peter B. Roenne, and Simon Rastkian. Hyperion: An enhanced version of the Selene end-to-end verifiable voting scheme. In *Sixth International Joint Conference on Electronic Voting E-Vote-ID*. University of Tartu Press, 2021.
- [35] Peter Y. A. Ryan, Peter B. Rønne, and Vincenzo Iovino. Selene: Voting with transparent verifiability and coercion-mitigation. In *Financial Cryptography and Data Security - FC 2016 International Workshops, BITCOIN, VOTING, and WAHC, Christ Church, Barbados, February 26, 2016, Revised Selected Papers*, volume 9604 of *Lecture Notes in Computer Science*, pages 176–192. Springer, 2016.

APPENDIX A

PROOF OF THEOREM 1

Theorem 1. *For every voting specification S , symbolic election verifiability implies general election verifiability for end-*

to-end voting traces. That is, for any trace τ of S , we have

$$\tau \models \text{SE2E}[\text{iv}_{\nabla}, \text{res}_o] \wedge \tau \models \Psi_{\text{E2E}} \implies \tau \models \text{E2E}[\text{iv}_{\nabla}, \text{res}_o]$$

Proof. Assume $\tau \models \text{SE2E}[\text{iv}_{\nabla}, \text{res}_o]$ and $\tau \models \Psi_{\text{E2E}}$ for an end-to-end voting trace τ of a voting specification.

(1) Let $\{\text{cr} \mid \tau \models \text{BBreg}(\text{cr})\}$ be the set of all registered public credentials. Assume there are n such credentials in total. By $\tau \models \Psi_{\text{E2E}}$, we have $n = |\{\text{tr} \mid \exists v. \tau \models \text{BBres}(\text{tr}, v)\}|$, implying that there are precisely n trackers $\{\text{tr}_1, \dots, \text{tr}_n\}$ for which $\tau \models \text{BBres}(\text{tr}_i, v)$ is true for some v . Then, by $\tau \models \Phi_{\text{eli}}$, each tracker $\text{tr}_i \in \{\text{tr}_1, \dots, \text{tr}_n\}$ is linked to a credential $\text{cr}_i \in \text{BBreg}$, i.e. we have $\tau \models \text{Link}(\text{cr}_i, \text{tr}_i)$. From the consistency properties $\tau \models \Phi_{\text{link1}} \wedge \Phi_{\text{link2}}$, for any $\tau \models \text{Link}(\text{cr}_i, \text{tr}_i) \wedge \text{Link}(\text{cr}_j, \text{tr}_j)$, we have $\text{tr}_i \neq \text{tr}_j$ and $\text{cr}_i \neq \text{cr}_j$ for any $i \neq j$. Furthermore, from Φ_{one} , we know that there is at most one occurrence of the event $\tau \models \text{BBres}(\text{tr}, v)$ for each tracker. Therefore, we can deduce that there is one-to-one correspondence between the list of trackers in the outcome, i.e. $\text{tr}_1, \dots, \text{tr}_n$, and the credentials that are linked to those. Let us call these credentials $\text{cr}_1, \dots, \text{cr}_n$.

(2) Without loss of generality, let $\{\text{tr}_1, \dots, \text{tr}_k\}$ be the subset of trackers for which $v_i \neq \perp$ for all $1 \leq i \leq k$. By $\tau \models \Phi_{\text{eli}}$, for each tracker tr_i in this subset, we have $\tau \models \text{Link}(\text{cr}_i, \text{tr}_i) \wedge \text{Reg}(\text{id}_i, \text{cr}_i)$. From the consistency properties $\tau \models \Phi_{\text{reg1}} \wedge \Phi_{\text{reg2}}$, for any $\tau \models \text{Reg}(\text{id}_i, \text{cr}_i) \wedge \text{Reg}(\text{id}_j, \text{cr}_j)$, we have $\text{id}_i \neq \text{id}_j$ and $\text{cr}_i \neq \text{cr}_j$, for any $i \neq j$. This implies that there are exactly k -distinct voter identities $\text{id}_1, \dots, \text{id}_k$ registered with the credentials $\text{cr}_1, \dots, \text{cr}_k$ for which we have $\tau \models \text{Reg}(\text{id}_i, \text{cr}_i)$ for $1 \leq i \leq k$. By (1), we can conclude that there is one-to-one correspondence between the set of distinct voter identities $\text{id}_1, \dots, \text{id}_k$, and the trackers $\text{tr}_1, \dots, \text{tr}_k$ for which $v_i \neq \perp$ on BBres .

(3) According to the one-to-one correspondence between trackers and voter identities established at point (2), we can define the following set:

$$\text{Res}_{\text{ID}, \text{V}}(\tau) = \{(\text{id}, v) \mid \tau \models \text{BBres}(\text{tr}, v) \wedge \text{Link}(\text{cr}, \text{tr}) \wedge \text{Reg}(\text{id}, \text{cr}) \wedge v \neq \perp\}$$

which associates to each voter identity the corresponding vote in the final result. Note that $\text{Res}(\tau) = \{v \mid (\text{id}, v) \in \text{Res}_{\text{ID}, \text{V}}(\tau)\}$. By $\tau \models \Phi_{\text{res}}^\circ$, for each $(\text{id}, v) \in \text{Res}_{\text{ID}, \text{V}}(\tau)$, either $\tau \models \text{Vote}(\text{id}, v)$ or $\tau \models \text{Adv}^\circ(\text{id})$, and these events are not mutually exclusive. If we define the following multisets,

$$\begin{aligned} \text{Vote}_{\text{ID}, \text{V}}(\tau) &= \{(\text{id}, v) \mid \tau \models \text{Vote}(\text{id}, v)\} \\ \text{Adv}_{\text{ID}, \text{V}}(\tau) &= \{(\text{id}, v) \mid \text{id} \in \text{Adv}^\circ(\tau)\} \end{aligned}$$

and sets:

$$\begin{aligned} R'_2 &= (\text{Res}_{\text{ID}, \text{V}}(\tau) \cap \text{Vote}_{\text{ID}, \text{V}}(\tau)) \setminus \text{Adv}_{\text{ID}, \text{V}}(\tau) \\ R'_3 &= \text{Res}_{\text{ID}, \text{V}}(\tau) \cap \text{Adv}_{\text{ID}, \text{V}}(\tau) \end{aligned}$$

then, each $(\text{id}, v) \in \text{Res}_{\text{ID}, \text{V}}(\tau)$ is either in R'_2 or in R'_3 , where $R'_2 \cap R'_3 = \emptyset$. Intuitively, R'_2 and R'_3 represent the partition of the final result according to honest or adversarial votes. Now, let us define:

$$\text{Ver}_{\text{ID}, \text{V}}(\tau) = \{(\text{id}, v) \mid \exists \text{id}', \text{cr}. \tau \models \text{Ver}^\nabla(\text{id}', \text{cr}, v) \wedge \text{Reg}(\text{id}, \text{cr})\}$$

and

$$\begin{aligned} R_1 &= \text{Res}_{\text{ID},V}(\tau) \cap \text{Ver}_{\text{ID},V}(\tau), \\ R_2 &= R'_2 \setminus \text{Ver}_{\text{ID},V}(\tau), \\ R_3 &= R'_3 \setminus \text{Ver}_{\text{ID},V}(\tau). \end{aligned}$$

By definition, $R_1 \cap R_2 = R_1 \cap R_3 = R_2 \cap R_3 = \emptyset$. Thus, we can conclude $\text{Res}_{\text{ID},V}(\tau) = R_1 \uplus R_2 \uplus R_3$. It follows that $\text{Res}(\tau) = V_1 \uplus V_2 \uplus V_3$, where $V_i = \{\nu \mid \exists \text{id}. (\text{id}, \nu) \in R_i\}$ for $i = 1, 2, 3$. We show that the multisets V_1, V_2, V_3 satisfy the requirements of Definition 3, respectively. That is, we show:

$$\begin{aligned} (1) \quad & V_1 = \text{Ver}^\nabla(\tau) \quad (2) \quad V_2 \in \text{Vote}^{\nabla, \diamond}(\tau), \\ (3) \quad & |V_3| \leq |\text{Adv}^\diamond(\tau) \setminus \text{Ver}_{\text{id}}^\nabla(\tau)|. \end{aligned}$$

(4.1) By definition, we have $V_1 \subseteq \text{Ver}^\nabla(\tau)$. We have to show $\text{Ver}^\nabla(\tau) \subseteq V_1$. Let $\text{id}_1, \dots, \text{id}_q$ be the set of (mutually distinct) identities of the voters who verified votes for some credential, and for which the votes are different from \perp . This means that there are non-empty sets of votes A_1, \dots, A_q such that

$$\forall i \in \{1, \dots, q\}, \forall \nu \in A_i, \exists \text{cr}. \tau \models \text{Ver}^\nabla(\text{id}_i, \text{cr}, \nu).$$

Intuitively, for each id_i , A_i represents the set of votes that are verified by that voter and which satisfy the revote policy. By definition, we have $\text{Ver}^\nabla(\tau) = A_1 \uplus \dots \uplus A_q$.

For each $i \in \{1, \dots, q\}$, let C_i be the set of credentials for which $\tau \models \text{Ver}^\nabla(\text{id}_i, \text{cr}, \nu)$. From $\tau \models \Phi_{\text{iv}2}^\nabla$, we deduce that $C_i \cap C_j = \emptyset$ for any $i \neq j$. Furthermore, assume $C_i = \{\text{cr}_1, \dots, \text{cr}_s\}$ for some $i \in \{1, \dots, q\}$. From $\tau \models \Phi_{\text{iv}3}^\nabla \wedge \Phi_{\text{link}2}$, we deduce

$$\begin{aligned} \tau &\models \text{Link}(\text{cr}_1, \text{tr}_1) \wedge \dots \wedge \text{Link}(\text{cr}_s, \text{tr}_s) \\ \tau &\models \text{BBreg}(\text{cr}_1) \wedge \dots \wedge \text{BBreg}(\text{cr}_s) \end{aligned}$$

for mutually distinct $\text{tr}_1, \dots, \text{tr}_s$. Let T_i be the set of trackers linked to the credentials in C_i , i.e. $T_i = \{\text{tr}_1, \dots, \text{tr}_s\}$. Then, from $\tau \models \Phi_{\text{link}1}$ and $C_i \cap C_j = \emptyset$, we conclude $T_i \cap T_j = \emptyset$ for any $i \neq j$.

On the other hand, as shown at point (I), each cr on BBreg is associated to a unique tr on BBres with $\tau \models \text{Link}(\text{cr}, \text{tr})$. From the consistency properties $\Phi_{\text{link}1}$ and $\Phi_{\text{link}2}$, the same credential cannot be associated to two different trackers. Then, we can deduce that each set T_i of trackers, that are linked to the credentials in C_i , are included in the final result. Moreover, by $\tau \models \Phi_{\text{one}}$, there is at most one vote per tracker in the result. Therefore, we have $\tau \models \text{BBres}(\text{tr}_1, \nu'_1) \wedge \dots \wedge \text{BBres}(\text{tr}_s, \nu'_s)$, for some votes ν'_1, \dots, ν'_s .

For each $i \in \{1, \dots, q\}$, let A'_i be the set of votes in the result that correspond to the trackers in T_i . Since we have established a one-to-one correspondence between C_i (associated to the set A_i) and T_i (associated to the set A'_i), we know that there is a one-to-one correspondence between the sets A_i and A'_i . We will show that these sets are actually equal. Indeed, for any $\text{cr} \in C_i, \nu \in A_i$ and associated $\text{tr} \in T_i, \nu' \in A'_i$, we have $\tau \models \text{Ver}^\nabla(\text{id}_i, \text{cr}, \nu) \wedge \text{Link}(\text{cr}, \text{tr}) \wedge \text{BBres}(\text{tr}, \nu')$. From $\tau \models \Phi_{\text{iv}1}^\nabla$, we deduce $\nu = \nu'$, and we can conclude that $A_i = A'_i$. Therefore, we have obtained:

$$\text{Ver}^\nabla(\tau) = A'_1 \uplus \dots \uplus A'_q = \{\nu'_1, \dots, \nu'_p\} \quad (\star)$$

for some votes ν'_1, \dots, ν'_p and some $p \geq q$ (every voter verified at least one vote).

On the other hand, by definition of the set R_1 , it consists of all the pairs $(\text{id}, \nu) \in \text{Res}_{\text{ID},V}(\tau)$ such that $\nu \neq \perp$ and some id' verified ν . Thus, we deduce

$$\begin{aligned} \{(\text{id}_1, \nu'_1), \dots, (\text{id}_p, \nu'_p)\} &\subseteq R_1, \text{ and therefore} \\ \{\nu'_1, \dots, \nu'_p\} &\subseteq V_1 \end{aligned} \quad (\star\star)$$

From (\star) and $(\star\star)$, we deduce $\text{Ver}^\nabla(\tau) \subseteq V_1$ and we can conclude $V_1 = \text{Ver}^\nabla(\tau)$ as required. Furthermore, we have also established

$$|\text{Res}_{\text{ID},V}(\tau) \cap \text{Ver}_{\text{ID},V}(\tau)| \geq |\text{Ver}_{\text{id}}^\nabla(\tau)| \quad (*)$$

i.e. the number of verified votes in the final result is greater than the number of voters for whom their votes are verified.

(4.2) We show that $V_2 \in \text{Vote}^{\nabla, \diamond}(\tau)$. By definition, we have

$$R_2 = (\text{Res}_{\text{ID},V}(\tau) \cap \text{Vote}_{\text{ID},V}(\tau)) \setminus (\text{Ver}_{\text{ID},V}(\tau) \cup \text{Adv}_{\text{ID},V}(\tau))$$

and $V_2 = \{\nu \mid \exists \text{id}. (\text{id}, \nu) \in R_2\}$. Furthermore, from $\Psi_{\text{res}2}$ and the one-to-one correspondence between voter identities and trackers established at point (2), we know that

$$R_2 = \{(\text{id}_1, \nu_1), \dots, (\text{id}_s, \nu_s)\} \text{ and } V_2 = \{\nu_1, \dots, \nu_s\}$$

for some mutually distinct $\text{id}_1, \dots, \text{id}_s$. Thus, for all $i \in \{1, \dots, s\}$ we have:

$$\tau \models \text{Vote}(\text{id}_i, \nu_i) \wedge \text{id}_i \notin (\text{Ver}_{\text{id}}^\nabla(\tau) \cup \text{Adv}^\diamond(\tau))$$

Therefore, using the definition of $\text{Vote}^{\nabla, \diamond}(\tau)$, we can conclude that $V_2 \in \text{Vote}^{\nabla, \diamond}(\tau)$.

(4.3) We show that $|V_3| \leq |\text{Adv}^\diamond(\tau) \setminus \text{Ver}_{\text{id}}^\nabla(\tau)|$. By definition, we have

$$R_3 = (\text{Res}_{\text{ID},V}(\tau) \cap \text{Adv}_{\text{ID},V}(\tau)) \setminus \text{Ver}_{\text{ID},V}(\tau)$$

and $V_3 = \{\nu \mid \exists \text{id}. (\text{id}, \nu) \in R_3\}$. Relying on Φ_{one} and the one-to-one correspondence between voter identities and trackers, we have that

$$\text{Res}_{\text{ID},V}(\tau) \cap \text{Adv}_{\text{ID},V}(\tau) = \{(\text{id}_1, \nu_1), \dots, (\text{id}_t, \nu_t)\}$$

for some mutually distinct $\text{id}_1, \dots, \text{id}_t$. We can also deduce that

$$|\text{Res}_{\text{ID},V}(\tau) \cap \text{Adv}_{\text{ID},V}(\tau)| \leq |\text{Adv}^\diamond(\tau)| \quad (\dagger)$$

From $(*)$, since by Φ_{one} there is one and only one vote in the result for each tracker (and therefore for each voter id according to our established correspondence), we deduce

$$|\text{Res}_{\text{ID},V}(\tau) \cap \text{Ver}_{\text{ID},V}(\tau) \cap \text{Adv}_{\text{ID},V}(\tau)| \geq |\text{Ver}_{\text{id}}^\nabla(\tau) \cap \text{Adv}^\diamond(\tau)| \quad (\dagger\dagger)$$

Thus, we can conclude that

$$\begin{aligned} |R_3| &= |(\text{Res}_{\text{ID},V}(\tau) \cap \text{Adv}_{\text{ID},V}(\tau)) \setminus \text{Ver}_{\text{ID},V}(\tau)| \\ &\leq |(\text{Res}_{\text{ID},V}(\tau) \cap \text{Adv}_{\text{ID},V}(\tau)) \setminus (\text{Res}_{\text{ID},V}(\tau) \cap \text{Ver}_{\text{ID},V}(\tau))| \\ &\leq |\text{Adv}^\diamond(\tau) \setminus \text{Ver}_{\text{id}}^\nabla(\tau)| \end{aligned}$$

where, to deduce the second inequality, we rely on $(\dagger), (\dagger\dagger)$ and the general set theory fact that for all sets A_1, A_2, B_1, B_2 :

$$\begin{aligned} |A_1| &\leq |A_2| \wedge |A_1 \cap B_1| \geq |A_2 \cap B_2| \\ &\implies |A_1 \setminus B_1| \leq |A_2 \setminus B_2| \end{aligned}$$

Hence, $|V_3| = |R_3| \leq |\text{Adv}^\diamond(\tau) \setminus \text{Ver}_{\text{id}}^\nabla(\tau)|$, as required. \square