

E-cclesia: Universally Composable Self-Tallying Elections

Abstract. The concept of a self-tallying election (STE) scheme was first introduced in [22] and captures electronic voting schemes in which the tallying authorities are the voters of the election themselves. This type of electronic voting is particularly compatible with and suitable for (but not only) Blockchain governance, where governance is expected to be maintained in a fully distributed manner.

In this work, we formalize the requirements for secure STE schemes in the Universal Composability (UC) framework. Our model captures the standard voting properties of *eligibility*, *fairness*, *vote-privacy*, and *one voter-one vote*. We present E-CCLESIA, a new family of STE schemes, and prove that it securely UC-realizes the STE functionality. We propose a first concrete instantiation of E-CCLESIA using RSA accumulators in combination with a Bitcoin-based time-lock encryption scheme. To this end, we provide *en passant* the first UC treatment of dynamic accumulators in the public setting as well as time lock encryption (TLE) schemes, and prove the equivalence of our UC definitions to the standard property-based accumulator definitions.

Keywords: e-voting, accumulators, time-lock encryption

1 Introduction

Voting has been a fundamental component of democratic societies for more than 2500 years. Electronic voting, or e-voting, refers to any voting procedure that involves computer systems to support one (or more) stages, *e.g.* casting and/or tallying. The aspiration being that electronic voting systems could offer elections with higher voter participation and better accuracy, while also providing enhanced security guarantees such as privacy and verifiability, even in the face of dishonest election authorities. With this drive, electronic voting has been of interest to the research community and governments all over the globe for the last three decades. Several cryptographic protocols have been proposed, implemented, and deployed for electronic voting [1,13,21,23,31,29,22,32,20].

The existing electronic voting schemes differ widely, and can be classified per their underlying trust assumptions. On one end, one finds fully centralized schemes which are often vulnerable to large scale attacks on robustness or privacy due to their centralized nature [4,34,16]; on the other end, schemes that are fully decentralized except maybe for setting up the election, where the (tallying) authorities are the voters of the election themselves. This second type has been termed *self-tallying election schemes* [29,22,32,20]. In between the two ends, and

to avoid the difficulties of self-tallying elections (STE) and dangers of centralized elections, a range of e-voting schemes which distribute the trust among a small set of authorities have also been proposed. Such systems achieve security as long as a subset of those authorities is honest. For instance, in mixnet-based schemes at least one authority needs to be honest [1,31], and in threshold schemes at most a fraction ϵ of the authorities can be corrupt and collude [23,21].

Of course, other things being equal, electronic voting schemes at the decentralized end of the spectrum are preferable as they can withstand more powerful adversaries. Self-tallying election schemes are a step forward in this direction, and move us one step closer to direct democracy. With the recent developments of Blockchain technologies and the appeal for on-chain distributed governance, such mechanisms are all the more topical. However, they present major challenges, which explains why full decentralization often needs to be abandoned in favor of achieving all the conflicting requirements that an electronic voting scheme should ensure. The main challenges STEs pose are in guaranteeing that no one (or no coalition of) voter(s) can boycott the election, no intermediate results are being leaked during the casting phase (*fairness*), and no vote can be linked back to the voter that cast it (*vote-privacy*). We refer the reader to [22] for a more in depth discussion on the difficulties STEs present.

Our goal in this paper is to further investigate the self-tallying paradigm, in light of the recent developments in distributed systems that we have been witnessing with the development of Blockchain technologies. Indeed, Nakamoto’s Bitcoin protocol [27] for a decentralized banking system has paved a new way for a decentralized future. With this as our credo too, we develop E-CCLESIA, a new family of self-tallying election schemes that satisfy the standard *eligibility*, *fairness*, *vote-privacy*, and *one voter-one vote* requirements for electronic voting. We formally prove that E-CCLESIA satisfies these properties in the Universal Composability paradigm proposed by Canetti in [8], which is a state-of-the-art framework for specifying and analyzing cryptographic protocols, especially when these are run under concurrent sessions. E-CCLESIA further satisfies *individual* and *universal verifiability* as a direct consequence of its decentralized nature, its reliance on a broadcast channel, and *ballot authentication*. Finally, we provide the first provably secure concrete instance of E-CCLESIA, borrowing ideas from the Blockchain literature. In particular, inspired by Zerocoin [26], we rely on RSA dynamic accumulators for unlinking cast votes from the voters who cast them. For fairness, we rely on a *time-lock encryption* (TLE) scheme, such as the Bitcoin-based scheme proposed by Liu *et al.* in [25].

Our contributions and roadmap. In Section 2 we present the necessary background. The subsequent sections are dedicated to the contributions of this work which can be summarized as follows:

Section 3 We provide a UC treatment (i.e. an ideal functionality) of self-tallying elections (STE) and extract its two basic sub-functionalities, namely an *eligibility* and a *vote management* functionality. We then present E-CCLESIA, a family of electronic voting schemes UC-realizing STEs via these two sub-

functionalities. The first provably secure concrete instance of E-CCLESIA can be derived from the next two sections.

Section 4 We provide the first UC definition for dynamic accumulators in the public setting, and prove that any UC-secure dynamic accumulator can be used to realize the eligibility functionality. We further prove that quasi-commutative dynamic accumulators as defined in [7], *e.g.* RSA accumulators are UC-secure dynamic accumulators, and thus provide a concrete realization of the eligibility module of E-CCLESIA.

Section 5 We provide the first UC definition for time-lock encryption (TLE), and show that a hybrid protocol that uses \mathcal{F}_{TLE} as a building block can UC-realize the vote management functionality. Moreover, we discuss the challenges and the solutions of a UC realization of \mathcal{F}_{TLE} by using the constructions of Liu’s *et al.* [25] or Shamir’s *et al.* time lock puzzles [30].

1.1 Related work

Self-tallying. Two self-tallying protocols established in the literature are the ones of Okamoto *et al.* [29] and Kiayias *et al.* [22]. In the former [29], an election authority uses blind signatures [12] in order to authenticate the eligible voters and their ballots. Next, each voter casts the commitment [5] of their ballot signed by the election authority over an anonymous channel. After the casting phase is over, each voter opens their commitment by casting the de-commitment key over the anonymous channel, so that the tally can be produced. One of the main drawbacks of this proposal is leakage of intermediate results, which violates the fairness property. In the latter proposal [22], the voters agree on a table where the sum of each row equals the voting choice of a voter. The link between the identity of a voter and their vote cannot be inferred from the table. In this proposal, the authors attempt to address the fairness problem by introducing a “dummy” party that casts its ballot last so that no intermediate results are leaked. However, this permits a single voter to be able to cause the protocol to abort. Moreover, the inclusion of such a “dummy” party is still not an optimal solution to the fairness problem as it re-introduces a trusted party in the STE setting. While self-tallying protocols such as [32,20] improve these proposals in terms of efficiency, they still share the limitations of [22]. Our work addresses these limitation, while preserving all the standard security requirements.

Time-lock encryption. TLE is a cryptographic primitive that allows a ciphertext to be decrypted only after a specific time period has elapsed. In some proposals, decryption can further be performed without requiring knowledge of any secret information. Previously proposed constructions are based either on witness encryption [18] or symmetric encryption [24]. The authors of these works provide game-based definitions in order to argue about the security of their constructions. Unfortunately, game-based definitions do not capture the variety of adversarial behaviour the UC framework [8] does. Moreover, the task of transferring these definitions to the UC setting is quite challenging due to some incompatibility between the two settings (concrete vs. asymptotic adversary).

Dynamic accumulators. Accumulators [17,14,35,19,33,7,15] are a well studied cryptographic primitive that take as an input a list of objects and output a representative value instead of the whole list. Despite the maturity of the primitive in terms of years of study, the only security arguments that have been provided are game-based. Recently, a work by Baldimtsi *et al.* [3] provides a UC treatment for accumulators with a trusted accumulator manager. The manager can accumulate and delete values from the list by using some trapdoor information. In our UC treatment, this trusted entity is only responsible for deletion but not for accumulation of the values (accumulation is a publicly available procedure), which is vital for our approach.

2 Preliminaries

Throughout the paper, we use λ as the security parameter and \cdot to denote a wildcard character.

2.1 The Universal Composability model

Overview. The security analysis in this work follows the *Universal Composability (UC)* paradigm introduced by Canetti in [8], which is the state-of-the-art cryptographic model for arguing about the security of protocols when run under concurrent sessions. In the UC framework, the parties engage in a protocol session (labeled by a unique session ID, sid) modeled as interactive Turing Machines (ITMs) that communicate in the presence of an *adversary* ITM \mathcal{A} that may control some of the parties. The protocol execution is scheduled by an *environment* ITM \mathcal{Z} that provides parties with inputs and may interact arbitrarily with \mathcal{A} . The intuition here is that (i) \mathcal{Z} captures the external “observer” that aims to break security by interacting with the protocol interface during session sid , while (ii) \mathcal{A} plays the role of the “insider” that helps \mathcal{Z} via any possible information it can obtain by engaging in the session in the back-end of the current execution.

The UC security of a protocol Π follows the *real-world/ideal-world indistinguishability* approach. Namely, security is captured via a special *ideal protocol* that has the same interface as Π that \mathcal{Z} interacts with, but now the parties are “dummy”, in the sense that they only forward their inputs provided by \mathcal{Z} to an *ideal functionality* \mathcal{F} , which is in the center of the back-end (i.e., the ideal protocol has a star topology) and *does not interact* with \mathcal{Z} directly. The ideal functionality \mathcal{F} formalizes a trusted party carrying out the task that Π intends to realize (e.g., secure communication, key agreement, authentication, etc.). The functionality \mathcal{F} interacts with the adversary present in the ideal protocol, usually called a *simulator* \mathcal{S} , and this interaction results in a “minimum leakage of information” that determines the ideal level of security that *any protocol* realizing said task should satisfy (not only Π). E.g., if \mathcal{F} formalizes an ideal secure channel, then the minimum leakage could be the ciphertext length. In case that \mathcal{Z} gives an input to a corrupted party P in the ideal world, the functionality \mathcal{F} pass that messages to \mathcal{S} and returns back to P whatever received from \mathcal{S} .

The real-world protocol is UC-secure if no environment \mathcal{Z} can distinguish its execution from the one of the ideal protocol managed by \mathcal{F} .

More formally, let $\text{EXEC}_{\mathcal{Z},\mathcal{A}}^{\Pi}$ denote an execution of a real-world protocol Π in the presence of the adversary \mathcal{A} scheduled by an environment \mathcal{Z} , and $\text{EXEC}_{\mathcal{Z},\mathcal{S}}^{\mathcal{F}}$ denote an execution of the ideal protocol managed by \mathcal{F} in the presence of a simulator \mathcal{S} , again scheduled by \mathcal{Z} . The UC security of Π is defined as follows.

Definition 1 (UC realization [8]). *The protocol Π is said to UC-realize the ideal functionality \mathcal{F} if for any PPT adversary \mathcal{A} , there exists a PPT simulator \mathcal{S} such that for any PPT environment \mathcal{Z} , the random variables $\text{EXEC}_{\mathcal{Z},\mathcal{A}}^{\Pi}$ and $\text{EXEC}_{\mathcal{Z},\mathcal{S}}^{\mathcal{F}}$ are computationally indistinguishable.*

Composition and modularity. Perhaps the most prominent feature of the UC paradigm – which is at the heart of the E-CCLESIA design – is the preservation of security of a protocol that runs concurrently with other protocol instances, or as a subroutine of another (often more complex) execution. In particular, assume a protocol Π that UC-realizes an ideal functionality \mathcal{F} according to Definition 1, and is used as a subroutine of a “larger” protocol $\tilde{\Pi}$. Then, UC guarantees that if we replace any instance of Π with \mathcal{F} , we obtain a “hybrid” protocol, denoted by $\tilde{\Pi}^{\Pi \rightarrow \mathcal{F}}$, that enjoys the same security as $\tilde{\Pi}$. Namely, if $\tilde{\Pi}$ UC-realizes some ideal functionality $\tilde{\mathcal{F}}$, then so does $\tilde{\Pi}^{\Pi \rightarrow \mathcal{F}}$.

The power of composition facilitates the design of complicated cryptographic schemes with a *high-degree of modularity*. Namely, the scheme’s formal description can be over the composition of ideal modules that are concurrently executed as subroutines. When a protocol Π using the functionalities $\mathcal{F}_1, \dots, \mathcal{F}_k$ UC-realizes a functionality \mathcal{F} , we say that it does so in the $\{\mathcal{F}_1, \dots, \mathcal{F}_k\}$ -*hybrid model* and we write $\Pi^{\mathcal{F}_1, \dots, \mathcal{F}_k}$ to clearly denote the hybrid functionalities.

A major advantage of the modular design in UC is that we can describe the cryptographic scheme in a partially abstract black-box way by simply using the hybrid functionalities instead of the associated real-world subroutines, and then prove security in the corresponding hybrid model. Henceforth, whenever the designer wishes to improve the scheme by implementing a more efficient version of some subroutine, they just need to prove the security of the new subroutine w.r.t. the associated hybrid functionality, which will directly imply the overall security of the improved scheme.

2.2 Setup functionalities

In Subsection 2.1, we summarized the UC model and showed how provably secure design can be done in a modular way by utilizing hybrid functionalities. We stress that in the UC literature, hybrid functionalities do not only play the role of abstracting some UC-secure real-world subroutine (e.g., a secure channel), but also formalize possible setup assumptions that are required to prove security when this is not done (and in many cases even impossible to achieve) in the “standard model”. For example, this type of setup functionalities may capture the concept of a trusted source of randomness, a clock, or a Public Key Infrastructure (PKI). Moreover, these setup functionalities can be *global*, i.e. they act

as shared states across multiple protocol instances and they can be accessed by other functionalities and even the environment that is external to the current session (recall that standard ideal functionalities do not directly interact with the environment). The extension of the UC framework in the presence of global setups has been introduced by Canetti *et al.* in [10]. Below, we present the setup functionalities that we consider across the E-CCLESIA design.

The global clock functionality $\mathcal{G}_{\text{clock}}$. In Fig. 1, we provide the definition of a *global clock* functionality $\mathcal{G}_{\text{clock}}$ similar to [2]. Time advances only when the environment has allowed all involved parties to advance.

The Global Clock functionality $\mathcal{G}_{\text{clock}}(\mathbf{P}, \mathbf{F})$.

For each session sid , the functionality initializes the global clock variable $\text{Cl} \leftarrow 0$ and the set of advanced parties per round as $L_{\text{adv}} \leftarrow \emptyset$.

- Upon receiving $(\text{sid}, \text{ADVANCE_CLOCK})$ from $P \in \mathbf{P}$, if $P \notin L_{\text{adv}}$, then it adds P to L_{adv} , sends the message $(\text{sid}, \text{ADVANCE_CLOCK})$ to P and notifies \mathcal{A} by forwarding $(\text{sid}, \text{ADVANCE_CLOCK}, P)$. If $L_{\text{adv}} = \mathbf{P} \cup \mathbf{F}$, then it updates as $\text{Cl} \leftarrow \text{Cl} + 1$ and resets $L_{\text{adv}} \leftarrow \emptyset$.

- Upon receiving $(\text{sid}, \text{ADVANCE_CLOCK})$ from $\mathcal{F} \in \mathbf{F}$, if $\mathcal{F} \notin L_{\text{adv}}$, then it adds \mathcal{F} to L_{adv} and sends the message $(\text{sid}, \text{ADVANCE_CLOCK})$ to \mathcal{F} . If $L_{\text{adv}} = \mathbf{P} \cup \mathbf{F}$, then it updates as $\text{Cl} \leftarrow \text{Cl} + 1$ and resets $L_{\text{adv}} \leftarrow \emptyset$.

- Upon receiving $(\text{sid}, \text{READ_CLOCK})$ from $X \in \mathbf{P} \cup \mathbf{F} \cup \{\mathcal{Z}, \mathcal{A}\}$, then it sends $(\text{sid}, \text{READ_CLOCK}, \text{Cl})$ to X .

Fig. 1. The global clock functionality $\mathcal{G}_{\text{clock}}(\mathbf{P}, \mathbf{F})$ interacting with the parties in \mathbf{P} , the functionalities in \mathbf{F} , the environment \mathcal{Z} and the adversary \mathcal{A} .

The common reference string functionality \mathcal{F}_{CRS} . Another setup assumption is the *common random string* model, where a single random string is drawn from a uniform distribution over strings. Fig. 2 formally defines \mathcal{F}_{CRS} as given by [8]. Note that \mathcal{F}_{CRS} requests permission from the simulator \mathcal{S} before returning the value, disclosing the identity of the requesting party. This action is often called (*public*) *delayed output* in the literature. In Section 4, when we refer to e.g. $\mathcal{F}_{\text{CRS}}^{\text{gen}}$, we mean the output distribution of the function **Gen**.

The anonymous broadcast functionality $\mathcal{F}_{\text{an.BC}}$. To guarantee the privacy of our STE scheme we assume that voters communicate via an *anonymous broadcast channel*. This communication interface is formalized via the functionality $\mathcal{F}_{\text{an.BC}}$ described in Fig. 3.

Providing a provably UC-secure realization of $\mathcal{F}_{\text{an.BC}}$ is out of the scope of this work. However, intuitively, one can instantiate an anonymous broadcast channel by deploying a blockchain (or any transaction ledger) where the users access the blockchain via an anonymous communications channel such as Tor or any mix-network routing protocol.

The common reference string functionality $\mathcal{F}_{\text{CRS}}(\mathbf{P}, D)$.

The functionality initializes a waiting list $L_{\text{wait}} \leftarrow \emptyset$.

- Upon receiving (sid, CRS) from $P_i \in \mathbf{P}$, if no value r is recorded, it picks $r \xleftarrow{\$} D$, adds P_i to L_{wait} and sends $(\text{sid}, \text{ALLOW}, P_i)$ to \mathcal{S} .
- Upon receiving $(\text{sid}, \text{ALLOWED}, P_i)$ from \mathcal{S} , if $P_i \in L_{\text{wait}}$, it sends $(\text{sid}, \text{CRS}, r)$ to P_i and \mathcal{S} and removes P_i from L_{wait} .

Fig. 2. The CRS functionality \mathcal{F}_{CRS} interacting with the parties in \mathbf{P} and the simulator \mathcal{S} , parameterized by distribution D .

The Anonymous Broadcast functionality $\mathcal{F}_{\text{an.BC}}(\mathbf{P})$.

The functionality initializes a list $L_{\text{pend}} \leftarrow \emptyset$ of messages pending to be broadcast.

- Upon receiving $(\text{sid}, \text{BROADCAST}, M)$ from $P_i \in \mathbf{P}$, it adds (M, P_i) to L_{pend} and sends $(\text{sid}, \text{ALLOW}, M)$ to \mathcal{S} .
- Upon receiving $(\text{sid}, \text{ALLOWED}, M)$ from \mathcal{S} , if $(M, P_i) \in L_{\text{pend}}$, then it sends $(\text{sid}, \text{BROADCAST}, M)$ to P_1, \dots, P_n , and \mathcal{S} . Then, it removes (M, P_i) from L_{pend} .

Fig. 3. The anonymous broadcast functionality $\mathcal{F}_{\text{an.BC}}$ interacting with the parties in $\mathbf{P} = \{P_1, \dots, P_n\}$ and the simulator \mathcal{S} .

The certification functionality $\mathcal{F}_{\text{cert}}(P, \mathbf{V})$.

- Upon receiving $(\text{sid}, \text{CORRUPT}, \mathbf{V}_{\text{corr}})$ from \mathcal{S} , if $\mathbf{V}_{\text{corr}} \subseteq \mathbf{V} \cup P$, it fixes \mathbf{V}_{corr} as the set of corrupted parties.
- Upon receiving $(\text{sid}, \text{SETUP})$ from P , if $\text{sid} = (P, \text{sid}')$ for some sid' , it sends $(\text{sid}, \text{SETUP})$ to \mathcal{S} .
- Upon receiving $(\text{sid}, \text{ALGORITHMS}, \text{Verify}, \text{Sign})$ from \mathcal{S} , it stores the algorithms and sends $(\text{sid}, \text{SETUP})$ to P .
- Upon receiving $(\text{sid}, \text{SIGN}, m)$ from P , it sets $\sigma \leftarrow \text{Sign}(m)$. If $\text{Verify}(m, \sigma) \neq 1$, it aborts. Otherwise it records (m, σ) and sends $(\text{sid}, \text{SIGNATURE}, m, \sigma)$ to P .
- Upon receiving $(\text{sid}, \text{VERIFY}, m, \sigma)$ from $V \in \mathbf{V}$, if $\text{Verify}(m, \sigma) = 1$, the signer is not corrupted and no entry (m, σ') for any σ' is recorded, it aborts. Otherwise it sends $(\text{sid}, \text{VERIFIED}, m, \text{Verify}(m, \sigma))$ to V .

Fig. 4. The certification functionality $\mathcal{F}_{\text{cert}}$ interacting with a prover P , a set of verifiers \mathbf{V} and the simulator \mathcal{S} .

The certification functionality $\mathcal{F}_{\text{cert}}$. The registration process for voting traditionally makes use of a channel through which the voters can identify themselves to the election authority. It would be natural to utilize a form of public key infrastructure, but modelling it in UC is not straightforward, and turns out to not be necessary for the purposes of the protocol that will be introduced in Section 3. Instead, we will use a certification scheme which provides signatures not bound to keys, but to *identities*. $\mathcal{F}_{\text{cert}}$, shown in Fig. 4 as defined by [8] (2005), provides commands for signature generation and verification, and is tied to a single party (so each party requires a separate instance). It can be realized as in [9] by an EUF-CMA secure signature scheme combined with a party acting as a trusted certificate authority.

3 A framework for STE schemes and the E-ccllesia family

In this section, we formalize STE schemes and provide an abstract description of the E-CCLLESIA family along the lines of modular UC design described in Subsection 2.1. The entities involved in an STE execution include the *voters* V_1, \dots, V_n and a *setup authority* SA that is active only prior to the voting and tally period and specifies the election parameters (e.g., list of eligible voters and the period of each election phase). In our threat model, the voters are statically corrupted while SA remains honest. Intuitively, the security properties that an STE scheme should satisfy are the following:

1. *Eligibility*: only eligible voters can vote.
2. *Fairness*: before the tally phase begins, no party can learn some partial result.
3. *Voter Privacy*: the voters' identities cannot be linked to their votes.
4. *One voter-one vote*: only one vote per (eligible) voter can be included in the tally.

Our approach to formally describing the E-CCLLESIA family is to first capture eligibility, fairness, and voter privacy via two main ideal modules:

(i) A *vote management functionality* \mathcal{F}_{vm} that handles ballot encryption, anonymous broadcast and opening.

(ii) An *eligibility functionality* $\mathcal{F}_{\text{elig}}$ that is responsible for generating anonymous credentials and authenticating the ballots of the eligible voters. The functionality can also link two ballots originating from the same (corrupted) voter.

We stress that ballot encryption and opening run by \mathcal{F}_{vm} ensure fairness, while credential generation and ballot authentication run by $\mathcal{F}_{\text{elig}}$ guarantee eligibility. Moreover, both functionalities combined safeguard voter privacy by implicitly featuring anonymous ballot authentication and casting. In addition, one voter-one vote is achieved by using $\mathcal{F}_{\text{elig}}$ for discarding multiple ballots that are linked to the same (corrupted) voter.

Given the aforementioned approach, the major technical challenge in designing E-CCLLESIA instantiations is to devise real-world protocols that UC-realize \mathcal{F}_{vm} and $\mathcal{F}_{\text{elig}}$. Upon completion of this task, the step towards full STE security is merely a careful composition in terms of interface by specifying how the STE

entities interact with \mathcal{F}_{vm} and $\mathcal{F}_{\text{elig}}$, and check the information they obtain from their engagement in the overall execution, so that the four security properties are preserved.

Therefore, in order to avoid repetition and excessive formalization, this section is organized as follows: in Subsection 3.1, we present the general *ideal STE functionality* \mathcal{F}_{STE} that captures all four aforementioned properties. To provide intuition, we describe \mathcal{F}_{STE} in a complete yet not strictly formal manner. Next, the functionalities $\mathcal{F}_{\text{elig}}$ and \mathcal{F}_{vm} are defined formally in Subsections 3.2 and 3.3, respectively, with precise references to \mathcal{F}_{STE} 's steps that are essentially handled by $\mathcal{F}_{\text{elig}}$ and \mathcal{F}_{vm} . Finally, in Subsection 3.4, we present the E-CCLLESIA family as a protocol $\Pi_{\text{E-CCLLESIA}}^{\mathcal{F}_{\text{elig}}, \mathcal{F}_{\text{vm}}, \mathcal{G}_{\text{clock}}}$ that UC-realizes \mathcal{F}_{STE} in the $\{\mathcal{F}_{\text{elig}}, \mathcal{F}_{\text{vm}}, \mathcal{G}_{\text{clock}}\}$ -hybrid model in a straightforward manner. We note that in the subsequent Sections 4 and 5 we construct two instantiations of $\mathcal{F}_{\text{elig}}$ and \mathcal{F}_{vm} , respectively, that specify the first version of E-CCLLESIA as a provably secure STE scheme.

3.1 The STE functionality \mathcal{F}_{STE}

The functionality \mathcal{F}_{STE} interacts with the setup authority SA, the voters in the set $\mathbf{V} = \{V_1, \dots, V_n\}$ and the simulator \mathcal{S} . In addition, it forwards all (sid, ADVANCE_CLOCK) and (sid, READ_CLOCK) commands to $\mathcal{G}_{\text{clock}}$ (cf. Fig. 1) on behalf of the (dummy) parties, while it also reads the time CI from $\mathcal{G}_{\text{clock}}$ whenever necessary. In the spirit of [8], we allow \mathcal{S} to provide \mathcal{F}_{STE} with the election algorithms, noting that this is *only for consistency* of the output format that \mathcal{F}_{STE} sends to the parties. As it will be clear in the description, the main security properties of eligibility, fairness, voter privacy and one voter-one vote are preserved by \mathcal{F}_{STE} per se, independently of the security of the algorithms provided by \mathcal{S} . The functionality consists of the following phases:

Setup. The functionality initializes as empty the lists of: eligible voters' credentials L_{elig} , ready voters L_{ready} , generated ballots L_{gball} , cast ballots L_{cast} , algorithms L_{vm} , state St_{fin} as 0, and ballots included in the tally set L_{tally} . It also stores the set $\mathbf{V}_{\text{corr}} \subseteq \mathbf{V}$ of corrupted voters provided by \mathcal{S} . Then, it operates as follows:

- Upon receiving (sid, ELECTION_INFO, $\mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}}$) from SA, where (i) \mathbf{V}_{elig} is the list of eligible voters and (ii) $t_{\text{cast}} < t_{\text{open}}$ are moments that determine the beginning of the **Cast** and **Tally** phases (see below), it acts as follows: In some cases we need more distinct time points than t_{cast} and t_{open} . However, these points should be able to be derived from t_{cast} and t_{open} , because these are the only times that are provided by \mathcal{Z} . For example, a time point can be the t_{wait} where it defines the period in which a voter can cast their ballot (from time t_{cast} until t_{wait}) and the period that everyone must wait before the tally phase (from time t_{wait} until t_{open}). The time point t_{wait} can be essential in some protocols that they need a period of synchronization. For that reason we define the function `define_time` which it takes as input the times t_{cast} and t_{open} and outputs a vector \mathbf{t} . This vector can be just the pair $t_{\text{cast}}, t_{\text{open}}$ (in that case `define_time` is the identity function) or a vector with more distinct time points (e.g $\mathbf{t} = (t_{\text{cast}}, t_{\text{open}}, t_{\text{wait}})$). It

depends on the specification of the protocol we want to realize how the function `define_time` is defined. Formally, \mathcal{F}_{STE} computes $\mathbf{t} \leftarrow \text{define_time}(t_{\text{cast}}, t_{\text{open}})$ and if $\mathbf{t} \neq \perp$, it sends $(\text{sid}, \text{SETUP_OK}, \mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}}, \mathbf{t})$ to SA after the permission of \mathcal{S} via a public delayed output.

■ Upon receiving $(\text{sid}, \text{ELIGIBLE})$ from SA, it informs \mathcal{S} , which replies with the *eligibility algorithms* `GenCred, AuthBallot, VrfyBallot, UpState` and an *initial credential state* St_{gen} . Then it provides $\langle V_i \rangle_{i \in [n]}$ and \mathcal{S} with the *election parameters* $\text{elec.par} := (\{\mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}}, \mathbf{t}\}, St_{\text{gen}})$.

■ Upon receiving $(\text{sid}, \text{START})$ from a voter $V \in \mathbf{V}_{\text{elig}} \setminus \mathbf{V}_{\text{corr}}$, if $V \notin L_{\text{ready}}$:

1. If $L_{\text{vm}} \neq \emptyset$, it (a) adds V to L_{ready} and (b) returns $(\text{sid}, \text{START_OK})$ to V .
2. Else, it sends $(\text{sid}, \text{START})$ to \mathcal{S} . Upon receiving the *vote management algorithms* `GenBallot, OpenBallot` from \mathcal{S} , it (a) updates L_{vm} with the algorithms, (b) adds V to L_{ready} , and (c) returns $(\text{sid}, \text{START_OK})$ to V .

Credential generation. This phase is active if $\text{Status}(\text{Cl}, \mathbf{t}, \text{Cred}) = \top$. The predicate `Status` shows us in which phase of the protocol we are based on the current time Cl. It takes as an input the current time Cl, the time vector \mathbf{t} and the phase we want to check if it is active at the moment (all the phases are `Credential, Cast` and `Tally` and can be checked with the acronym `Cred, Cast` and `Tally` respectively). If this phase is active `Status` outputs the special symbol \perp , else it outputs \perp . The predicate is instantiated according to the protocol we want to realize. An example of what this predicate might be is the check if $t_{\text{cast}} < \text{Cl} < t_{\text{open}}$ in the case we want to check if we are in the `Cast` phase.

■ Upon receiving a message $(\text{sid}, \text{GEN_CRED})$ from $V \in \mathbf{V}_{\text{elig}} \setminus \mathbf{V}_{\text{corr}}$ for the first time and after permission from \mathcal{S} via public delayed output, it runs $(\text{cr}, \text{rc}) \leftarrow \text{GenCred}(1^\lambda, \text{elec.par})$ and adds $(V, \text{cr}, \text{rc}, 1)$ to L_{elig} . Here, `cr, rc` play the role of the *private and public part* of the credential, respectively. If there are already tuples $(\cdot, \text{cr}, \cdot, \cdot)$ or $(\cdot, \cdot, \text{rc}, \cdot)$ in L_{elig} or $(\text{cr}, \text{rc}) = \perp$, it sends $(\text{sid}, \text{GEN_CRED}, \perp)$ to V and halts. Else, it adds $(V, \text{cr}, \text{rc}, 1)$ to L_{elig} .

If \mathcal{F}_{STE} receives some credential pair $(\text{sid}, \text{GEN_CRED}, \text{cr}, \text{rc}, V)$ from \mathcal{S} on behalf of a corrupted yet eligible voter V for the first time, if there are no tuples $(\cdot, \text{cr}, \cdot, 1)$ or $(\cdot, \cdot, \text{rc}, 1)$ in L_{elig} , then \mathcal{F}_{STE} adds $(V, \text{cr}, \text{rc}, 0)$ to L_{elig} .

In any case of recording $(V, \text{cr}, \text{rc}, \cdot)$, it sends $(\text{sid}, \text{GEN_CRED}, V, \text{rc})$ to $\langle V_j \rangle_{j \in [n]}$ and \mathcal{S} after permission of \mathcal{S} via public delayed output.

Cast. This phase is active if $\text{Status}(\text{Cl}, \mathbf{t}, \text{Cast}) = \top$ and manages ballot generation, authentication and broadcasting. For ballot authentication, \mathcal{F}_{STE} once computes the *final credential state* St_{fin} by running `UpState` on input St_{gen} and the set of public credentials `rc` included in L_{elig} .

■ Upon receiving a message $(\text{sid}, \text{CAST}, o)$ from an honest and eligible voter V for the first time such that $(V, \text{cr}, \text{rc}, 1) \in L_{\text{elig}}$, it executes the following steps:

1. It generates a *‘dummy’ ballot* as an encryption of 0s by running $v \leftarrow \text{GenBallot}(\text{vote.par}, 0^{|\lambda|}, \lambda)$ and adds the tuple $(V, v, o, 1)$ to the list of the generated ballots L_{gball} . This dummy encryption step captures the semantic security of the `GenBallot` algorithm.

2. It generates an *authentication receipt* as $\sigma \leftarrow \text{AuthBallot}(v, \text{cr}, \text{rc}, St_{\text{fin}})$ for ballot v . If $\text{VrfyBallot}(v, \sigma, St_{\text{fin}}, \text{reg.par}) = 0$, it sends $(\text{sid}, \text{AUTH_BALLOT}, \perp)$ to V and halts.
3. It sends $(\text{sid}, \text{CAST}, v, \sigma)$ to \mathcal{S} . Upon receiving $(\text{sid}, \text{CAST_ALLOWED})$ from \mathcal{S} , it adds $(V, o, v, \text{cr}, \sigma, 1)$ to L_{cast} as the authenticated ballot tuple for V .
4. It sends $(\text{sid}, \text{CAST}, v, \sigma)$ to $\langle V_j \rangle_{j \in [n]}$ and \mathcal{S} . Note that \mathcal{F}_{STE} broadcasts the ballot “anonymously”, i.e. without revealing V ’s identity.

In addition, \mathcal{F}_{STE} allows \mathcal{S} to cast ballots arbitrarily on behalf of a corrupted voter V : when it receives $(\text{sid}, \text{CAST}, o, v, \sigma, V)$ from \mathcal{S} , if there is a tuple $(V, \text{cr}, \text{rc}, 0)$ in L_{elig} , it adds $(V, o, v, \text{cr}, \sigma, 0)$ to L_{cast} , else it adds $(V, o, v, \perp, \sigma, 0)$ to L_{cast} . In any case, it sends $(\text{sid}, \text{CAST}, v, \sigma)$ to $\langle V_j \rangle_{j \in [n]}$ and \mathcal{S} .

Tally. This phase is active if $\text{Status}(\text{Cl}, t, \text{Tally}) = \top$ and manages the opening of all the valid ballots that should be included in the tally (thus, capturing fairness).

■ Upon receiving $(\text{sid}, \text{TALLY})$ from a voter V or \mathcal{S} , \mathcal{F}_{STE} executes the following steps:

1. **For** every tuple $(V, o, v, \text{cr}, \sigma, b)$ in L_{cast} , it runs the ballot verification algorithm $x \leftarrow \text{VrfyBallot}(v, \sigma, St_{\text{fin}})$. If $x = 1$, then \mathcal{F}_{STE} performs the following security checks:
 - (i). If there is no tuple $(V, \text{cr}, \text{rc}, \cdot)$ in L_{elig} , then eligibility has been breached and it returns $(\text{sid}, \text{TALLY}, \perp)$ to V or \mathcal{S} .
 - (ii). If there is a tuple $(V, \text{cr}, \text{rc}, 1)$ in L_{elig} and there is a tuple $(\cdot, o', v', \text{cr}', \sigma', 1)$ in L_{cast} such that $(\text{cr}' = \text{cr}) \wedge (v' \neq v)$, then forgery of some honest ballot has happened, so it returns $(\text{sid}, \text{TALLY}, \perp)$ to V or \mathcal{S} . Note that this check implies that the AuthBallot algorithm should satisfy unforgeability.
 - (iii). Otherwise, it adds (V, o, v, b) to L_{tally} .
As a result, at the end of the **For** loop in this step, L_{tally} contains successfully verified ballots cast only by the intended eligible voters. Note that multiple valid ballots coming from corrupted eligible voters may still exist at this point, something which is handled at the step below.
2. **For** every tuple $(V, \text{cr}, \text{rc}, \cdot)$ in L_{elig} such that there are multiple tuples $(V, o^1, v^1, \cdot), \dots, (V, o^n, v^n, \cdot)$ in L_{tally} , it removes all multiple tuples from L_{tally} except the first one it recorded. This step is important to ensure the one voter-one vote property. Thus, at this point L_{tally} contains valid tuples in one-to-one correspondence with the eligible voters that participated in the election.
3. **For** every (V, o, v, \cdot) in L_{tally} , it first allows \mathcal{S} to open the ballot of any corrupted voter V_j to an alternative opening \tilde{o}_j by updating the corresponding tuple in L_{tally} as $(V_j, \tilde{o}_j, v_j, 0)$. Then, it checks the correctness of the honest voters’ ballot generation (which implies the correctness of the ballot opening algorithm OpenBallot) as follows: if there are two tuples $(V, o, v, 1), (V', o', v', 1)$ in L_{tally} such that $(v' = v) \wedge (o \neq o')$, then it returns $(\text{sid}, \text{TALLY}, \perp)$ to V . Otherwise, if no such conflict exists, it returns $(\text{sid}, \text{TALLY}, \{(o, v) \mid (V, o, v, \cdot) \in L_{\text{tally}}\})$ to V or \mathcal{S} , where $\{(o, v) \mid (V, o, v, \cdot) \in L_{\text{tally}}\}$ is the multiset of eligible voters’ tallied options.

■ Upon receiving $(\text{sid}, \text{LEAKAGE})$ from \mathcal{S} , it reads the time Cl from $\mathcal{G}_{\text{Clock}}$. If $\text{Status}(\text{Cl}, \mathbf{t}, \text{Cred}) = \text{Status}(\text{Cl}, \mathbf{t}, \text{Cast}) = \text{Status}(\text{Cl}, \mathbf{t}, \text{Open}) = \perp$, then it returns to \mathcal{S} all the triples $(v, o, 1)$ such that $(V, v, o, 1) \in L_{\text{gball}} \wedge (V, v, \sigma, 1) \in L_{\text{cast}}$. This leakage illustrates the fact that in “wait” time period, the period where the protocol is between the **Cast** and **Tally** phase, the adversary might be able to open the ballots. On the other hand, this information at that point is not very useful (the adversary can not break fairness) because the adversary can not change its ballot, as the **Cast** phase is over. In protocols where no such periods exist the condition $\text{Status}(\text{Cl}, \mathbf{t}, \text{Cred}) = \text{Status}(\text{Cl}, \mathbf{t}, \text{Cast}) = \text{Status}(\text{Cl}, \mathbf{t}, \text{Open}) = \perp$ is never satisfied.

3.2 Eligibility functionality $\mathcal{F}_{\text{elig}}$

The eligibility functionality $\mathcal{F}_{\text{elig}}$ takes over the following parts of \mathcal{F}_{STE} execution: (i) in **Setup** the eligibility algorithms and the initial credential state, (ii) the entire (private and public) **Credential generation**, (iii) in **Cast** the creation of ballot authentication receipts, (iv) in **Tally** the ballot verification step, so that unforgeability is preserved (cf. Step 1. in \mathcal{F}_{STE} ’s description for $(\text{sid}, \text{TALLY})$ messages), and (v) the linkability of two ballots to the same voter, so that one voter - one vote is preserved. The functionality is presented in Fig. 5.

3.3 Vote management functionality \mathcal{F}_{vm}

The vote management functionality \mathcal{F}_{vm} takes over the following parts of \mathcal{F}_{STE} execution: (i) in **Setup** the configuration of the vote management algorithms, (ii) in **Cast** the secure ballot generation and casting and (iv) in **Tally** the ballot opening step, so that fairness and ballot correctness is preserved (cf. Step 3. in \mathcal{F}_{STE} ’s description for $(\text{sid}, \text{TALLY})$ messages). The functionality is presented in Fig. 6.

3.4 Description of the E-cclesia family

We provide a description of the E-CCLERIA family $\Pi_{\text{E-CCLERIA}}^{\mathcal{F}_{\text{elig}}, \mathcal{F}_{\text{vm}}, \mathcal{G}_{\text{clock}}}$ of STE schemes as a hybrid protocol that makes use of the main modules $\mathcal{F}_{\text{elig}}, \mathcal{F}_{\text{vm}}$. Any pair of real-world protocols $\Pi_{\text{elig}}, \Pi_{\text{vm}}$ that UC-realize $\mathcal{F}_{\text{elig}}, \mathcal{F}_{\text{vm}}$ specifies a member of the family. The description of $\Pi_{\text{E-CCLERIA}}^{\mathcal{F}_{\text{elig}}, \mathcal{F}_{\text{vm}}, \mathcal{G}_{\text{clock}}}$ follows the phases and command interface of \mathcal{F}_{STE} in Subsection 3.1 as described below:

Setup.

- Upon receiving $(\text{sid}, \text{ELECTION_INFO}, \mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}})$ from \mathcal{Z} , if $\mathbf{V}_{\text{elig}} \subseteq \mathbf{V}$ and $t_{\text{cast}} < t_{\text{open}}$, SA sends $(\text{sid}, \text{SETUP_INFO}, \mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}})$ to \mathcal{F}_{vm} .
- Upon receiving $(\text{sid}, \text{ELIGIBLE})$ from the environment \mathcal{Z} , if SA has received $(\mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}})$, it sends $(\text{sid}, \text{ELIGIBLE}, \mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}})$ to $\mathcal{F}_{\text{elig}}$ which sends $\text{reg.par} := (\mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}}, \mathbf{t}, St_{\text{gen}})$ to $\langle V_i \rangle_{i \in [n]}$. Upon receiving reg.par from $\mathcal{F}_{\text{elig}}$, each voter $V \in \mathbf{V}$ stores it as the election parameters elec.par .

The eligibility functionality $\mathcal{F}_{\text{elig}}(\mathbf{V}, \text{define_time}, \text{Status})$.

The functionality initializes the lists of eligible voters $L_{\text{elig}} \leftarrow \emptyset$, of authenticated ballots of eligible voters $L_{\text{auth}} \leftarrow \emptyset$, the value $St_{\text{fin}} = 0$, and its status to ‘init’.

■ Upon receiving $(\text{sid}, \text{CORRUPT}, \mathbf{V}_{\text{corr}})$ from \mathcal{S} , if $\mathbf{V}_{\text{corr}} \subseteq \mathbf{V}$ and status=init, it fixes \mathbf{V}_{corr} as the set of corrupted voters.

■ Upon receiving $(\text{sid}, \text{ELIGIBLE}, \mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}})$ from SA, if $\mathbf{V}_{\text{elig}} \subseteq \mathbf{V}$ and status=init, it sends $(\text{sid}, \text{SETUP_ELIG})$ to \mathcal{S} . Upon receiving $(\text{sid}, \text{SETUP_ELIG}, \text{GenCred}, \text{AuthBallot}, \text{VrfyBallot}, \text{UpState}, St_{\text{gen}})$ from \mathcal{S} , if status=init, then:

1. It computes $\mathbf{t} \leftarrow \text{define_time}(t_{\text{cast}}, t_{\text{open}})$. If $\mathbf{t} \neq \perp$, it sets $\text{reg.par} := (\mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}}, \mathbf{t}, St_{\text{gen}})$ as registration parameters.
2. It sets its status to ‘credential’ and sends $(\text{sid}, \text{ELIG_PAR}, \text{reg.par})$ to $\langle V_j \rangle_{j \in [n]}$ and \mathcal{S} .

■ Upon receiving $(\text{sid}, \text{GEN_CRED})$ from $V \in \mathbf{V}_{\text{elig}} \setminus \mathbf{V}_{\text{corr}}$, if status=credential, then it reads the time Cl from $\mathcal{G}_{\text{clock}}$. If $\text{Status}(\text{Cl}, \mathbf{t}, \text{Cast}) = \top$, it sets its status to ‘cast’. Otherwise, if $\text{Status}(\text{Cl}, \mathbf{t}, \text{Cred}) = \top$, it executes the following steps:

1. If there is no tuple $(V, \text{cr}', \text{rc}', 1)$ in L_{elig} , it runs $(\text{cr}, \text{rc}) \leftarrow \text{GenCred}(1^\lambda, \text{reg.par})$. If there are tuples $(\cdot, \text{cr}, \cdot, \cdot)$ or $(\cdot, \cdot, \text{rc}, \cdot)$ in L_{elig} or $(\text{cr}, \text{rc}) = \perp$, it sends $(\text{sid}, \text{GEN_CRED}, \perp)$ to V and halts. Else, it adds $(V, \text{cr}, \text{rc}, 1)$ to L_{elig} after permission of \mathcal{S} via public delayed output.
2. It sends $(\text{sid}, \text{GEN_CRED}, V, \text{rc})$ to $\langle V_j \rangle_{j \in [n]}$ and \mathcal{S} after permission of \mathcal{S} via public delayed output.

■ Upon receiving $(\text{sid}, \text{GEN_CRED})$ from $V \in \mathbf{V}_{\text{corr}}$, it forwards the message $(\text{sid}, \text{GEN_CRED}, V)$ to \mathcal{S} .

■ Upon receiving $(\text{sid}, \text{GEN_CRED}, V, \text{cr}, \text{rc})$ from \mathcal{S} , if $V \in \mathbf{V}_{\text{corr}}$, then:

1. If \mathbf{V}_{elig} and there are no tuples $(V, \text{cr}', \text{rc}', 0)$, $(\cdot, \text{cr}, \cdot, 1)$ or $(\cdot, \cdot, \text{rc}, 1)$ in L_{elig} , then it adds $(V, \text{cr}, \text{rc}, 0)$ to L_{elig} .
2. It sends $(\text{sid}, \text{GEN_CRED}, V, \text{rc})$ to $\langle V_j \rangle_{j \in [n]}$ and \mathcal{S} .

■ Upon receiving $(\text{sid}, \text{AUTH_BALLOT}, v)$ from $V \in \mathbf{V}_{\text{elig}} \setminus \mathbf{V}_{\text{corr}}$, if status=cast, then it reads the time Cl from $\mathcal{G}_{\text{clock}}$. If $\text{Status}(\text{Cl}, \mathbf{t}, \text{Open}) = \top$, it sets its status to ‘open’. Otherwise, if $\text{Status}(\text{Cl}, \mathbf{t}, \text{Cast}) = \top$, it executes the following steps:

1. If $St_{\text{fin}} = 0$, then it runs $St_{\text{fin}} \leftarrow \text{UpState}(St_{\text{gen}}, \{\text{rc}(\cdot, \cdot, \text{rc}, \cdot) \in L_{\text{elig}}\})$.
2. If there is some $(V, \text{cr}, \text{rc}, 1) \in L_{\text{elig}}$ but no $(V, v', \text{cr}, \sigma', 1) \in L_{\text{auth}}$, then it runs $\sigma \leftarrow \text{AuthBallot}(v, \text{cr}, \text{rc}, St_{\text{fin}}, \text{reg.par})$. If $\text{VrfyBallot}(v, \sigma, St_{\text{fin}}, \text{reg.par}) = 0$, it sends $(\text{sid}, \text{AUTH_BALLOT}, \perp)$ to V and halts. Else, it (a) adds $(V, v, \text{cr}, \sigma, 1)$ to L_{auth} , and (b) returns $(\text{sid}, \text{AUTH_BALLOT}, v, \sigma)$ to V .

■ Upon receiving $(\text{sid}, \text{AUTH_BALLOT})$ from $V \in \mathbf{V}_{\text{corr}}$, it forwards the message $(\text{sid}, \text{AUTH_BALLOT}, V)$ to \mathcal{S} .

■ Upon receiving $(\text{sid}, \text{AUTH_BALLOT}, V, v, \sigma)$ from \mathcal{S} , if there is a tuple $(V, \text{cr}, \text{rc}, 0)$ in L_{elig} , then it adds $(V, v, \text{cr}, \sigma, 0)$ to L_{auth} . It returns $(\text{sid}, \text{AUTH_BALLOT}, V, v, \sigma)$ to V .

■ Upon receiving $(\text{sid}, \text{VER_BALLOT}, v, \sigma)$ from $V \in \mathbf{V}$:

1. It computes $x \leftarrow \text{VrfyBallot}(v, \sigma, St_{\text{fin}}, \text{reg.par})$.
2. If $x = 1$ and there is no cr such that there are tuples $(\cdot, \text{cr}, \cdot, \cdot) \in L_{\text{elig}}$ and $(\cdot, v, \text{cr}, \sigma, \cdot) \in L_{\text{auth}}$, it sends $(\text{sid}, \text{VER_BALLOT}, v, \sigma, \perp)$ to V and halts.
3. If $x = 1$ and there are tuples $(\cdot, v, \text{cr}, \sigma, \cdot)$, $(\cdot, v', \text{cr}', \sigma', 1) \in L_{\text{auth}}$ such that $\text{cr} = \text{cr}'$ and $v \neq v'$, it sends $(\text{sid}, \text{VER_BALLOT}, v, \sigma, \perp)$ to V and halts.
4. Else, it sends $(\text{sid}, \text{VER_BALLOT}, v, \sigma, x)$ to V .

■ Upon receiving $(\text{sid}, \text{LINK_BALLOTS}, (v, \sigma), (v', \sigma'))$ from $V \in \mathbf{V}$, if there are tuples $(\cdot, v, \text{cr}, \sigma, \cdot)$, $(\cdot, v', \text{cr}', \sigma', \cdot) \in L_{\text{auth}}$ such that $\text{cr} = \text{cr}'$, then it sets $x = 1$, else $x = 0$. Then, it sends $(\text{sid}, \text{LINK_BALLOTS}, (v, \sigma), (v', \sigma'), x)$ to V .

Fig. 5. The eligibility functionality $\mathcal{F}_{\text{elig}}(\mathbf{V})$ interacting with voters \mathbf{V} , SA, and the simulator \mathcal{S} .

The vote management functionality $\mathcal{F}_{\text{vm}}(\mathbf{V}, \text{define_time}, \text{Status})$.

The functionality initializes the lists of generated ballots L_{gball} , cast ballots L_{cast} , opened ballots L_{open} , ballot algorithms L_{vm} , and ready voters L_{ready} as empty. Then it sets its status to ‘exec’.

- Upon receiving $(\text{sid}, \text{CORRUPT}, \mathbf{V}_{\text{corr}})$ from \mathcal{S} , if $\mathbf{V}_{\text{corr}} \subseteq \mathbf{V}$ and $\text{status}=\text{exec}$, it fixes \mathbf{V}_{corr} as the set of corrupted voters.
- Upon receiving $(\text{sid}, \text{SETUP_INFO}, \mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}})$ from SA it computes $\mathbf{t} \leftarrow \text{define_time}(t_{\text{cast}}, t_{\text{open}})$. If $\text{status}=\text{exec}$ and $\mathbf{t} \neq \perp$, it sets $\text{vote.par} := (\mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}}, \mathbf{t})$ as voting parameters and sends $(\text{sid}, \text{SETUP_OK}, \text{vote.par})$ to SA, after permission from \mathcal{S} via delayed output.
- Upon receiving $(\text{sid}, \text{START})$ from voter $V \in \mathbf{V}_{\text{elig}} \setminus \mathbf{V}_{\text{corr}}$, if $V \notin L_{\text{ready}}$:
 1. If $L_{\text{vm}} \neq \emptyset$, it (a) adds V to L_{ready} and (b) returns $(\text{sid}, \text{START_OK})$ to V .
 2. Else, it sends $(\text{sid}, \text{START})$ to \mathcal{S} . Upon receiving $(\text{sid}, \text{SETUP_INFO}, \text{GenBallot}, \text{OpenBallot})$ from \mathcal{S} , it (a) updates L_{vm} with the algorithms, (b) adds V to L_{ready} , and (c) returns $(\text{sid}, \text{START_OK})$ to V .
- Upon receiving $(\text{sid}, \text{GEN_BALLOT}, o)$ from $V \in L_{\text{ready}}$:
 1. If there is no $(V, v', o', 1) \notin L_{\text{gball}}$, it (a) runs $v \leftarrow \text{GenBallot}(\text{vote.par}, 0^{|\text{ol}|}, \lambda)$, (b) adds $(V, v, o, 1)$ to L_{gball} . Upon receiving $(\text{sid}, \text{GEN_BALLOT}, v)$ from \mathcal{S} , it returns $(\text{sid}, \text{GEN_BALLOT}, o, v)$ to V .
 2. Else it returns $(\text{sid}, \text{GEN_BALLOT}, o, \perp)$ to V .
- Upon receiving $(\text{sid}, \text{GEN_BALLOT})$ from $V \in \mathbf{V}_{\text{corr}}$, it sends the message $(\text{sid}, \text{GEN_BALLOT}, V)$ to \mathcal{S} .
- Upon receiving $(\text{sid}, \text{GEN_BALLOT}, o, v, V)$ from \mathcal{S} , if $V \in \mathbf{V}_{\text{corr}}$ it sends $(\text{sid}, \text{GEN_BALLOT}, o, v)$ to V .
- Upon receiving $(\text{sid}, \text{CAST}, v, \sigma)$ from V , if $V \notin \mathbf{V}_{\text{corr}}$ and $\text{status}=\text{exec}$, it reads the time Cl from $\mathcal{G}_{\text{clock}}$. If $\text{Status}(\text{Cl}, \mathbf{t}, \text{Open}) = \top$, it sets the status to ‘open’, otherwise:
 1. If $(V, v, o', 1) \notin L_{\text{gball}}$, it returns $(\text{sid}, \text{CAST}, v, \sigma, \perp)$ to V .
 2. If there is no $(V, v', \sigma', 1) \notin L_{\text{cast}}$, it sends $(\text{sid}, \text{ALLOW_CAST}, v, \sigma)$ to \mathcal{S} . Upon receiving $(\text{sid}, \text{CAST_ALLOWED})$ from \mathcal{S} , it adds $(V, v, \sigma, 1)$ to L_{cast} and sends $(\text{sid}, \text{CAST}, v, \sigma)$ to $\langle V_j \rangle_{j \in [m]}$ and \mathcal{S} .
 3. If there is a tuple $(V, v', \sigma', 1)$ in L_{cast} , it returns $(\text{sid}, \text{CAST}, v, \sigma, \perp)$ to V .
- Upon receiving $(\text{sid}, \text{CAST})$ from $V \in \mathbf{V}_{\text{corr}}$, it sends $(\text{sid}, \text{CAST}, V)$ to \mathcal{S} .
- Upon receiving $(\text{sid}, \text{CAST}, v, \sigma, V)$ from \mathcal{S} , if $V \in \mathbf{V}_{\text{corr}}$ and $\text{status}=\text{exec}$, it reads the time Cl from $\mathcal{G}_{\text{clock}}$. If $\text{Status}(\text{Cl}, \mathbf{t}, \text{Open}) = \top$ it sets status to ‘open’, otherwise it adds $(V, v, \sigma, 0)$ to L_{cast} and sends $(\text{sid}, \text{CAST}, v, \sigma)$ to $\langle V_j \rangle_{j \in [m]}$ and \mathcal{S} .
- Upon receiving $(\text{sid}, \text{OPEN}, v)$ from any party $P \in \mathbf{V} \cup \{\mathcal{S}\}$, it reads the time Cl from $\mathcal{G}_{\text{clock}}$. If $\text{Status}(\text{Cl}, \mathbf{t}, \text{Open}) = \top$, it sets the status to ‘open’ and executes the following steps:
 1. If there is a tuple $(V, v, \sigma, \cdot) \in L_{\text{cast}}$, and a *unique* $(V, v, o, 1) \in L_{\text{gball}}$, it sends $(\text{sid}, \text{OPEN}, v, o)$ to P .
 2. Else, if there is a tuple $(V, v, \sigma, \cdot) \in L_{\text{cast}}$ and at least two tuples $(V, v, o, 1), (V', v', o', 1) \in L_{\text{gball}}$ such that $(v = v') \wedge (o \neq o')$, it sends $(\text{sid}, \text{OPEN}, v, \perp)$ to P .
 3. Else, if there is a tuple $(V, v, \sigma, \cdot) \in L_{\text{cast}}$ but there is no tuple $(V, v, o, 1) \in L_{\text{gball}}$, it sends $(\text{sid}, \text{OPEN}, v)$ to \mathcal{S} . Then it sends the reply it gets from \mathcal{S} to P .
- Upon receiving $(\text{sid}, \text{LEAKAGE})$ from \mathcal{S} , it reads the time Cl from $\mathcal{G}_{\text{Clock}}$. If $\text{Status}(\text{Cl}, \mathbf{t}, \text{Cred}) = \text{Status}(\text{Cl}, \mathbf{t}, \text{Cast}) = \text{Status}(\text{Cl}, \mathbf{t}, \text{Open}) = \perp$ then it returns to \mathcal{S} all the triples $(v, o, 1)$ such that $(V, v, o, 1) \in L_{\text{gball}} \wedge (V, v, \sigma, 1) \in L_{\text{cast}}$.

Fig. 6. The vote management functionality $\mathcal{F}_{\text{vm}}(\mathbf{V})$ interacting with voters \mathbf{V} , SA and the simulator \mathcal{S} .

- Upon receiving $(\text{sid}, \text{START})$ from \mathcal{Z} for the first time, if voter V is in \mathbf{V}_{elig} , then she sends $(\text{sid}, \text{START})$ to \mathcal{F}_{vm} which replies with $(\text{sid}, \text{START_OK})$.

Credential generation. This phase is completely managed by $\mathcal{F}_{\text{elig}}$.

- Upon receiving $(\text{sid}, \text{GEN_CRED})$ from \mathcal{Z} V sends $(\text{sid}, \text{GEN_CRED})$ to $\mathcal{F}_{\text{elig}}$, which in turn sends $(\text{sid}, \text{GEN_CRED}, V, \text{rc})$ to $\langle V_j \rangle_{j \in [n]}$ (or sends $(\text{sid}, \text{GEN_CRED}, \perp)$ to V and halts).

Cast. Here, \mathcal{F}_{vm} and $\mathcal{F}_{\text{elig}}$ combined carry out the ballot generation, authentication and broadcasting tasks.

- Upon receiving $(\text{sid}, \text{CAST}, o)$ from \mathcal{Z} V executes the following steps:
 1. She sends $(\text{sid}, \text{GEN_BALLOT}, o)$ to \mathcal{F}_{vm} which replies with the generated ballot as $(\text{sid}, \text{GEN_BALLOT}, o, v)$ (or sends $(\text{sid}, \text{GEN_BALLOT}, o, \perp)$ to V and halts).
 2. Then, she sends $(\text{sid}, \text{AUTH_BALLOT}, v)$ to $\mathcal{F}_{\text{elig}}$ which replies with the authentication receipt for v as $(\text{sid}, \text{AUTH_BALLOT}, v, \sigma)$ (or sends $(\text{sid}, \text{AUTH_BALLOT}, \perp)$ to V and halts).
 3. Finally, she sends $(\text{sid}, \text{CAST}, v, \sigma)$ to \mathcal{F}_{vm} which broadcasts the message to $\langle V_j \rangle_{j \in [n]}$. In turn, the voters store the received pair (v, σ) .

Tally. In order for the voter to perform self-tallying, she accesses $\mathcal{F}_{\text{elig}}$ for ballot verification and linkability and \mathcal{F}_{vm} for ballot opening.

- Upon receiving a message $(\text{sid}, \text{TALLY})$ from \mathcal{Z} , V executes the following steps:
 1. **For** every tuple $(\text{sid}, \text{CAST}, v, \sigma)$ she has obtained from \mathcal{F}_{vm} , V sends $(\text{sid}, \text{VER_BALLOT}, v, \sigma)$ to $\mathcal{F}_{\text{elig}}$ which replies with $(\text{sid}, \text{VER_BALLOT}, v, \sigma, x)$, where $x \in \{0, 1, \perp\}$.
If there is any ballot verification request such that $\mathcal{F}_{\text{elig}}$ replied with $x = \perp$, then V sets tally to \perp . Otherwise, she includes in her tally set all pairs (v, σ) such that $\mathcal{F}_{\text{elig}}$ replied with $x = 1$.
 2. V discards multiple ballots as follows: for every pair $(v, \sigma), (v', \sigma')$ in her tally set, she sends $(\text{sid}, \text{LINK_BALLOTS}, (v, \sigma), (v', \sigma'))$ to $\mathcal{F}_{\text{elig}}$. If she gets $(\text{sid}, \text{LINK_BALLOTS}, (v, \sigma), (v', \sigma'), 1)$ as a response, then she discards the ballot she received the last out of those two. Clearly, after this pairwise check is completed, all except one of ballots that are linked will be removed from the tally set, so that one voter-one vote is guaranteed.
 3. **For** every pair (v, σ) in the tally set, V sends $(\text{sid}, \text{OPEN}, v)$ to \mathcal{F}_{vm} , which replies with the opening $(\text{sid}, \text{OPEN}, v, o)$. Then, V adds o to the multiset of all opened options (initialized as empty). If at any time \mathcal{F}_{vm} replies with $(\text{sid}, \text{OPEN}, v, \perp)$, then V sets tally to \perp .
 4. Finally, she sets the tally result as the multiset of all opened options.

Security. As already mentioned in the introduction of Section 3, proving that $\Pi_{\text{E-CCLESIA}}^{\mathcal{F}_{\text{elig}}, \mathcal{F}_{\text{vm}}, \mathcal{G}_{\text{clock}}}$ UC-realizes \mathcal{F}_{STE} in the $\{\mathcal{F}_{\text{elig}}, \mathcal{F}_{\text{vm}}, \mathcal{G}_{\text{clock}}\}$ -hybrid model is straightforward given the description of \mathcal{F}_{STE} , $\mathcal{F}_{\text{elig}}$ and \mathcal{F}_{vm} . Below, we provide the theorem statement and proof.

Theorem 1. *The protocol $\Pi_{\text{E-CCLESIA}}^{\mathcal{F}_{\text{elig}}, \mathcal{F}_{\text{vm}}, \mathcal{G}_{\text{clock}}}$ described in Subsection 3.4 UC-realizes \mathcal{F}_{STE} in the $\{\mathcal{F}_{\text{elig}}, \mathcal{F}_{\text{vm}}, \mathcal{G}_{\text{clock}}\}$ -hybrid model.*

Proof. We define a simulator \mathcal{S} as follows: In the cases of public delayed outputs, we assume that \mathcal{S} forwards the message to the adversary \mathcal{A} as if it was from either $\mathcal{F}_{\text{elig}}$ or \mathcal{F}_{vm} and responds to \mathcal{F}_{STE} in the same way as \mathcal{A} . Moreover, whatever command \mathcal{F}_{STE} receives on behalf of a corrupted party, we assume that \mathcal{F}_{STE} forwards that message to \mathcal{S} . Then \mathcal{S} forwards that message to \mathcal{A} as if it was from either $\mathcal{F}_{\text{elig}}$ or \mathcal{F}_{vm} and it returns to \mathcal{F}_{STE} whatever it receives from \mathcal{A} . We describe the details below.

During the **Setup** phase, when \mathcal{S} receives the corruption set \mathbf{V}_{corr} from \mathcal{Z} , it forwards it to \mathcal{A} as if it was from \mathcal{Z} . Then \mathcal{S} plays the role of $\mathcal{F}_{\text{elig}}$ and \mathcal{F}_{vm} and receives back the corruption set from \mathcal{A} . Next, \mathcal{S} forwards the corruption set to \mathcal{F}_{STE} . Upon receiving $(\text{sid}, \text{ELECTION_INFO}, \mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}})$ from \mathcal{F}_{STE} , \mathcal{S} stores the parameters and it forwards the message to \mathcal{A} as if it was from \mathcal{F}_{vm} . Then, \mathcal{S} returns to \mathcal{F}_{STE} whatever it receives from \mathcal{A} . After \mathcal{S} receives $(\text{sid}, \text{ELIGIBLE})$ from \mathcal{F}_{STE} , it plays the role of $\mathcal{F}_{\text{elig}}$ and it sends $(\text{sid}, \text{SETUP_ELIG})$ to \mathcal{A} . Then, upon receiving the eligibility algorithms $(\text{sid}, \text{SETUP_ELIG}, \text{GenCred}, \text{AuthBallot}, \text{VrfyBallot}, \text{UpState}, St_{\text{gen}})$ from \mathcal{A} , \mathcal{S} forwards the message to \mathcal{F}_{STE} . When \mathcal{S} receives $(\text{sid}, \text{START})$ from \mathcal{F}_{STE} , it forwards the message to \mathcal{A} as if it was from \mathcal{F}_{vm} . Upon receiving the vote management algorithms $(\text{sid}, \text{SETUP_INFO}, \text{GenBallot}, \text{OpenBallot})$ from \mathcal{A} , \mathcal{S} forwards the message to \mathcal{F}_{STE} .

During the **Credential generation** phase, when \mathcal{S} receives a credential generation request from \mathcal{F}_{STE} on behalf of a corrupted party V , \mathcal{S} forwards the request to \mathcal{A} as if it was from $\mathcal{F}_{\text{elig}}$. Then, upon receiving $(\text{sid}, \text{GEN_CRED}, V, \text{cr}, \text{rc})$ from \mathcal{A} , \mathcal{S} forwards the message to \mathcal{F}_{STE} .

During the **Cast** phase, when \mathcal{S} receives $(\text{sid}, \text{CAST}, v, \sigma)$ from \mathcal{F}_{STE} , it forwards the message to \mathcal{A} as if it was from \mathcal{F}_{vm} . Then \mathcal{S} returns whatever it receives from \mathcal{A} . Upon receiving a cast ballot request from \mathcal{F}_{STE} on behalf of a corrupted party V , \mathcal{S} forwards the message to \mathcal{A} as if it was from \mathcal{F}_{vm} . After \mathcal{S} receives $(\text{sid}, \text{CAST}, \tilde{v}, \tilde{\sigma}, V)$ from \mathcal{A} , it returns the message $(\text{sid}, \text{CAST}, o, \tilde{v}, \tilde{\sigma}, V)$ to \mathcal{F}_{STE} for some o . The choice of o is irrelevant because during the **Tally** phase \mathcal{S} will be asked by \mathcal{F}_{STE} for a new opening of \tilde{v} . This happens because every time \mathcal{F}_{STE} (\mathcal{F}_{vm} as well) receives a request for opening a malicious ballot, it gives the token to the simulator and it answers according to what \mathcal{S} provides.

During the **Tally** phase, when \mathcal{S} is asked by \mathcal{F}_{STE} for an alternative ballot opening of the tuple $(V, o, v, \cdot) \in L_{\text{tally}}$ where $V \in \mathbf{V}_{\text{corr}}$, \mathcal{S} sends the message $(\text{sid}, \text{OPEN}, v)$ to \mathcal{A} as if it was from \mathcal{F}_{vm} and it returns to \mathcal{F}_{STE} whatever it receives from \mathcal{A} . Upon receiving $(\text{sid}, \text{OPEN}, v)$ from \mathcal{Z} for the first time, \mathcal{S} sends $(\text{sid}, \text{TALLY})$ to \mathcal{F}_{STE} . Upon receiving $(\text{sid}, \text{TALLY}, \{(o, v) | (V, o, v, \cdot) \in L_{\text{tally}}\})$ from \mathcal{F}_{STE} , \mathcal{S} records the tally. Next, \mathcal{S} forwards the message $(\text{sid}, \text{OPEN}, v)$ to \mathcal{A} as if it was from \mathcal{Z} . Upon receiving $(\text{sid}, \text{OPEN}, v)$, \mathcal{S} playing the role of \mathcal{F}_{vm} returns the plaintext for that ballot to \mathcal{A} . If the message doesn't exist in the list previously provided by \mathcal{F}_{STE} , then \mathcal{S} asks \mathcal{A} for the opening of that message as before.

The distribution of messages is exactly the same in both settings, since the algorithms that are used are the same. As a result the simulation is perfect. \square

4 Realizing $\mathcal{F}_{\text{elig}}$ via accumulators

This section describes the primitives needed to build a real protocol that realizes the eligibility functionality. Since cryptographic accumulators do not have a suitable UC treatment in the literature that would fit our purpose, Subsections 4.1 and 4.2 provide an ideal accumulator functionality and link it to a standard definition. Subsection 4.3 presents the Π_{elig} protocol.

4.1 Secure accumulators and \mathcal{F}_{acc}

The purpose of secure accumulators is to provide an object representing a set and create witnesses for specific items being in the set. We use a game-based definition by [7]. In this model, it is assumed that a trusted accumulator manager runs Gen to generate both the public parameters and the secret trapdoor.

Definition 2 (Accumulator). $\text{AC} = (\mathcal{X}_\lambda, \text{Gen})$ is an accumulator scheme for a family of inputs $\{\mathcal{X}_\lambda\}$ if it has the following properties:

1. **Efficient generation**

Gen is an efficient probabilistic algorithm such that $\text{Gen}(1^\lambda) \rightarrow (f, \text{aux}_f)$ for random $f \in F_\lambda$, where F_λ is a family of functions.

2. **Efficient evaluation**

$f \in F_\lambda$ is a polynomial-size circuit such that $f(w, x) \rightarrow a$ for $(w, x) \in \mathcal{U}_f \times \mathcal{X}_\lambda$ and $a \in \mathcal{U}_f$ where \mathcal{U}_f is an efficiently samplable domain for f .

3. **Correctness**

$w \in \mathcal{U}_f$ is a witness for $x \in \mathcal{X}_\lambda$ in the accumulator $a \in \mathcal{U}_f$ under f if $\text{Verify}_f(w, x, a) \rightarrow \top$ where $\text{Verify}_f(w, x, a) := f(w, x) \stackrel{?}{=} a$.

4. **Quasi-commutativity**

For all $f \in F_\lambda$, $u \in \mathcal{U}_f$, $x_1, x_2 \in \mathcal{X}_\lambda$: $f(f(u, x_1), x_2) = f(f(u, x_2), x_1)$. So for $X = \{x_1, \dots, x_m\} \subset \mathcal{X}_\lambda$, $f(u, X) := f(f(\dots(f(u, x_1), \dots), x_m))$.

5. **Witness unforgeability**

Let $\mathcal{U}'_f \times \mathcal{X}'_\lambda$ denote the domains for which the computational procedure for function $f \in F_\lambda$ is defined (so $\mathcal{U}_f \subseteq \mathcal{U}'_f, \mathcal{X}_\lambda \subseteq \mathcal{X}'_\lambda$).

AC is secure if for all PPT adversaries \mathcal{A}_λ :

$$\Pr[f \leftarrow \text{Gen}(1^\lambda); u \leftarrow \mathcal{U}_f; (x, w, X) \leftarrow \mathcal{A}_\lambda(f, \mathcal{U}_f, u) :$$

$$X \subset \mathcal{X}_\lambda; w, a \in \mathcal{U}'_f; x \in \mathcal{X}'_\lambda; x \notin X;$$

$$\text{Verify}_f(u, X, a) = \text{Verify}_f(w, x, a) = \top] = \text{negl}(\lambda).$$

6. **Efficient deletion**

AC is dynamic if there exist efficient algorithms Delete , Update such that if $x, x' \in X$ and $\text{Verify}_f(u, X, a) = \text{Verify}_f(w, x, a) = \top$, then:

$\text{Delete}(\text{aux}_f, a, x') \rightarrow a'$ such that $\text{Verify}_f(u, X \setminus \{x'\}, a') = \top$,

and $\text{Update}(f, a, a', x, x') \rightarrow w'$ such that $\text{Verify}_f(w', x, a') = \top$.

Remark 1. Definition 2 is well known, but it is not the only definition of accumulators in the standard model [15]. The reason for matching our functionality to this particular definition is that it suits the public setting that we wanted

to model the best. As opposed to the more common definitions which involve a central authority responding to requests for accumulation, here the accumulator is a public function which any party can use.

This is not without caveats, as we still rely on trusted setup, and efficient deletions require the manager to use his secret trapdoor. Coming up with a better definition for our usecase (and constructions that satisfy it) is an interesting problem, but out of scope of this work. We also note that quasi-commutativity is often handled as an optional property, but it is implicitly embedded in the way [7] defines witnesses. Extending to non-commutative constructions would hence also require a new definition.

Recently, [3] provided a UC functionality for accumulators, however it is not suitable for our protocol since it requires the accumulator to be managed centrally by some authority.

The definition of the ideal functionality \mathcal{F}_{acc} is shown in Figures 7 and 8. \mathcal{F}_{acc} , parametrized by the input set \mathcal{X}_λ , stores a list L_{acc} of (a, X, w) entries where a is the accumulator value, X is the set of accumulated items and w is either another accumulator or the basis u , a fixed value representing the empty set. Since \mathcal{F}_{acc} is not producing the witnesses itself, to ensure correctness it needs to keep some auxiliary information. L_{sets} is a list of $(a, [S_1, S_2, \dots])$ entries for each a in L_{acc} , where $[S_1, S_2, \dots]$ is a list of all known representations of the contents of the set accumulated in a . L_{ref} is a list of (a, A_r) entries which represents references to accumulators that have not been expanded yet. We call an accumulator *expanded* if it has an entry in L_{sets} , i.e. there exists some representation of its contents. If it is not expanded, there is no such entry, but A_r may appear (syntactically) as part of some list S_i for another entry. We denote all such unexpanded references in capitals and with the $_r$ suffix. We call a set *fully-expanded* if it does not contain references to any unexpanded accumulators.

To illustrate the use of these lists, we give an example execution for the accumulation of items x_1, x_2, x_3 :

1. *Accumulate* $\{x_2, x_3\}$ *with witness* a : \mathcal{F}_{acc} computes $f(a, \{x_2, x_3\}) = b$, adds $(b, \{x_2, x_3\}, a)$ to L_{acc} , $(b, [A_r \cup \{x_2, x_3\}])$ to L_{sets} and (a, A_r) to L_{ref} .
2. *Accumulate* x_1 *with witness* u : \mathcal{F}_{acc} computes $f(u, x_1) = a$ and adds $(a, \{x_1\}, u)$ to L_{acc} and $(a, [\{x_1\}])$ to L_{sets} . Since (a, A_r) is in L_{ref} , the first entry in L_{sets} expands to $(b, [\{x_1, x_2, x_3\}])$ and (a, A_r) is removed from L_{ref} .
3. *Accumulate* $\{x_1, x_2, x_3\}$ *with witness* u : \mathcal{F}_{acc} computes $f(u, \{x_1, x_2, x_3\}) = c$. If $c = b$, it adds $(c, \{x_1, x_2, x_3\}, u)$ to L_{acc} , otherwise it aborts.

We can now describe the individual commands.

Setup. Upon initialization, \mathcal{F}_{acc} asks the adversary for the parameters and for the accumulation and verification algorithms. It distributes these to any party that asks. The party M that first calls the functionality is the designated accumulator manager. The **Delete** algorithm implicitly includes the trapdoor (*aux_f* in Definition 2).

Setup.

The lists L_{acc} , L_{sets} and L_{ref} are initialized as empty.

■ Upon receiving (sid, SETUP) from the manager M for the first time, it sends (sid, SETUP, \mathcal{X}_λ) to \mathcal{S} . Upon receiving (sid, ALGS, (\mathcal{U}_f , u), f , Verify, Delete, Update) from \mathcal{S} , it sets $\text{params} = (\mathcal{X}_\lambda, \mathcal{U}_f, u)$ and sends (sid, ALGS, params , f , Verify, Delete, Update) to M .

■ Upon receiving (sid, PARAMS) from any party P , it returns (sid, PARAMS, params , f , Verify, Update) to P .

Accumulation.

■ Upon receiving (sid, ACCUM, w , X) from party P :

1. If $X \not\subseteq \mathcal{X}'_\lambda$ or $w \notin \mathcal{U}'_f$, it aborts.
2. If there is $(a, X, w) \in L_{acc}$ for some a , it returns (sid, ACCUM, a) to P .
3. It computes $a \leftarrow f(w, X)$.
4. If $a \notin \mathcal{U}_f$ or $\text{Verify}(w, X, a) \neq \top$, it aborts.
5. If $w = u$, it creates $L_a \leftarrow [X]$.
Else if there is $(w, L_w) \in L_{sets}$ for some L_w , it creates $L_a \leftarrow [W \cup X : W \in L_w]$.
Else it adds (w, W_r) to L_{ref} , and creates $L_a \leftarrow [W_r \cup X]$.
6. If there is $(a, L'_a) \in L_{sets}$ for some L'_a , it replaces the pair with $(a, L'_a \parallel L_a)$ (i.e. it appends L_a to L'_a) unless $L'_a = L_a$.
Else if for any $A \in L_a$, there is $(a', L'_a) \in L_{sets}$ for some $a' \neq a$ and $A \in L'_a$, it aborts.
Else if $(a, A_r) \in L_{ref}$ for some A_r :
 - (a) For every $(a', L'_a) \in L_{sets}$, it replaces each occurrence of A_r in L'_a with items from L_a (so each $A' = A_r \cup \dots$ can expand into multiple $E = A \cup \dots$).
 - (b) For any E , if there is $(a'', L''_a) \in L_{sets}$ for some $a'' \neq a'$ such that $E \in L''_a$, it aborts.
 - (c) It removes (a, A_r) from L_{ref} and it adds (a, L_a) to L_{sets} .
Else, it adds (a, L_a) to L_{sets} .
7. It adds (a, X, w) to L_{acc} and returns (sid, ACCUM, a) to P .

Deletion.

■ Upon receiving (sid, DELETE, a , x , A) from manager M :

1. If $A \not\subseteq \mathcal{X}'_\lambda$, $a \notin \mathcal{U}_f$ or $x \notin A$, it aborts.
2. If there is no $(a, \{x\}, w) \in L_{acc}$ for some w , and no $(a, L_a) \in L_{sets}$ such that $A \in L_a$, it aborts.
3. It computes $a' \leftarrow \text{Delete}(a, x)$.
4. If $a' = u$ and $A = \{x\}$, it skips to Step 7.
5. If $\text{Verify}(u, A \setminus \{x\}, a') \neq \top$, it aborts.
6. It runs Step 6 of ACCUM for a' and $L_a = [A \setminus \{x\}]$.
7. It returns (sid, DELETED, a') to M .

■ Upon receiving (sid, UPDATE, a , a' , x , x' , A) from P :

1. If $A \not\subseteq \mathcal{X}'_\lambda$, $a, a' \notin \mathcal{U}_f$ or $x, x' \notin A$, it aborts.
2. It runs Step 2 of DELETE.
3. It computes $w' \leftarrow \text{Update}(a, a', x, x')$. If $\text{Verify}(w', \{x\}, a') \neq \top$, it aborts.
4. It runs Steps 5 to 7 of ACCUM for w' and a' .
5. It returns (sid, UPDATED, w') to P .

Fig. 7. Setup, accumulation and deletion commands of \mathcal{F}_{acc} .

Verification.

- Upon receiving $(\text{sid}, \text{VERIFY}, w, X, a)$ from party P :
 1. If $X \not\subseteq \mathcal{X}'_\lambda$, $a \notin \mathcal{U}_f$ or $w \notin \mathcal{U}'_f$, it aborts.
 2. If $(a, X, w) \in L_{\text{acc}}$, it returns $(\text{sid}, \text{VERIFIED}, \top)$ to P .
 3. If $\text{Verify}(w, X, a) = \top$ and (a) or (b) holds, it aborts:
 - (a) There is $(a, L_a) \in L_{\text{sets}}$ such that for some fully-expanded $A \in L_a$ we have $X \not\subseteq A$.
 - (b) There are $(w, L_w), (a, L_a) \in L_{\text{sets}}$ such that for some $W \in L_w$ and $A \in L_a$ we have $A = W \cup V$ where V is a fully-expanded set with $V \not\subseteq X$.
 4. It returns $(\text{sid}, \text{VERIFIED}, \text{Verify}(w, X, a))$ to P .

Fig. 8. Verification commands of \mathcal{F}_{acc} .

Accumulation. In an **ACCUM** call, \mathcal{F}_{acc} first checks if the given input has an associated accumulator value in L_{acc} and returns it if that's the case. If not, it computes the value and checks that it verifies. Step 5 ensures the propagation of known sets throughout the L_{sets} list.

For example, suppose \mathcal{F}_{acc} is given witness w and input $\{z\}$. If $(w, [X_r \cup \{y\}, \{x, y\}]) \in L_{\text{sets}}$ and $(x, X_r) \in L_{\text{ref}}$, then the L_a corresponding to the newly calculated accumulator a will have to be of the form $[X_r \cup \{y, z\}, \{x, y, z\}]$. Step 6 ensures that the output is consistent and does not invalidate previous values, and that references to incomplete sets are updated correctly. Continuing the example, if $(a, [Y_r \cup \{x, z\}]) \in L_{\text{sets}}$, \mathcal{F}_{acc} appends the newly calculated sets to the existing list, replacing the old entry with $(a, [Y_r \cup \{x, z\}, X_r \cup \{y, z\}, \{x, y, z\}])$. If there is no such entry, but there is $a' \neq a$ such that $(a', [\{x, y, z\}]) \in L_{\text{sets}}$, then correctness is violated and \mathcal{F}_{acc} aborts. The third condition deals with the case when the computed a has a reference $(a, A_r) \in L_{\text{ref}}$. The inputs w and $\{z\}$ reveal something about the contents of A_r , which lets \mathcal{F}_{acc} update all the entries in L_{sets} which contain A_r as long as there are no conflicts. A conflict could arise if there were $(a', [A_r \cup \{s\}]), (a'', [\{x, y, z, s\}]) \in L_{\text{sets}}$ with $a' \neq a''$, because $A_r \cup \{s\}$ expands to $\{x, y, z, s\}$.

Deletion. To match the model of [7], only the manager M can delete, but any party P can update their witness after a deletion. Definition 2 only permits deletion of items we know are accumulated so both commands require the calling party to specify the set A corresponding to a .

Verification. There are two ways a forgery could occur. The adversary could try to verify an accumulator either with an item that has not been accumulated in it (e.g. the tuple $(u, \{x\}, a)$ when there is $(a, [\{y, z\}]) \in L_{\text{sets}}$ such that $x \notin \{y, z\}$), or with an invalid witness (e.g. $(w, \{x\}, a)$ when there are $(w, [\{x, y\}]), (a, [\{x, y, z\}]) \in L_{\text{sets}}$ such that $z \neq x$).

4.2 A protocol that realizes \mathcal{F}_{acc}

In Fig. 9, we define a real protocol Π_{AC} that uses an accumulator scheme in the CRS model. Theorem 2 captures the equivalence between our functionality and secure accumulator schemes, and in particular implies that \mathcal{F}_{acc} can be realized by a strong RSA accumulator construction [7]. The proof can be found in Supplementary Material A.1.

The accumulation protocol Π_{AC} for $\text{AC} = (\mathcal{X}_\lambda, \text{Gen})$.

- Upon receiving (sid, SETUP) for the first time, M sends (sid, CRS) to $\mathcal{F}_{\text{CRS}}^{\text{gen}}$ to receive (sid, CRS, $(f, \mathcal{U}_f, u, \text{aux}_f)$). It sets $\text{params} = (\mathcal{X}_\lambda, \mathcal{U}_f, u)$ and returns (sid, ALGS, params , f , Verify , $\text{Delete}(\text{aux}_f, \cdot, \cdot)$, $\text{Update}(f, \cdot, \cdot, \cdot, \cdot)$).
- Upon receiving (sid, PARAMS), P sends (sid, CRS) to $\mathcal{F}_{\text{CRS}}^{\text{gen}}$ to receive (CRS, sid, (f, \mathcal{U}_f, u)). It sets $\text{params} = (\mathcal{X}_\lambda, \mathcal{U}_f, u)$ and returns (sid, PARAMS, params , f , Verify , Update).
- Upon receiving (sid, ACCUM, w, X), P aborts if $X \not\subseteq \mathcal{X}'_\lambda$ or $w \notin \mathcal{U}'_f$. Otherwise it returns (sid, ACCUM, $f(w, X)$).
- Upon receiving (sid, DELETE, a, x', A), M aborts if $A \not\subseteq \mathcal{X}'_\lambda$, $a \notin \mathcal{U}'_f$, $x' \notin A$ or if $\text{Verify}(u, A, a) \neq \top$. Otherwise it returns (sid, DELETED, $\text{Delete}(\text{aux}_f, a, x')$).
- Upon receiving (sid, UPDATE, a, a', x, x', A), P aborts if $A \not\subseteq \mathcal{X}'_\lambda$, $a, a' \notin \mathcal{U}'_f$, $x, x' \notin A$ or if $\text{Verify}(u, A, a) \neq \top$. Otherwise it returns (sid, UPDATED, $\text{Update}(a, a', x, x')$).
- Upon receiving (sid, VERIFY, w, X, a), P aborts if $X \not\subseteq \mathcal{X}'_\lambda$, $w \notin \mathcal{U}'_f$. Else it returns (sid, VERIFIED, $\text{Verify}(w, X, a)$).

Fig. 9. Definition of Π_{AC} in the $\mathcal{F}_{\text{CRS}}^{\text{gen}}$ -hybrid model.

Theorem 2. Π_{AC} UC-realizes \mathcal{F}_{acc} in the $\mathcal{F}_{\text{CRS}}^{\text{gen}}$ -hybrid model if and only if $\text{AC} = (\mathcal{X}_\lambda, \text{Gen})$ is a secure dynamic accumulator scheme.

4.3 A protocol that realizes $\mathcal{F}_{\text{elig}}$

To realize $\mathcal{F}_{\text{elig}}$, we need two other primitives besides accumulators, both of which have UC formulations in the literature: *non-interactive commitments* [5] and (non-interactive zero-knowledge) *signatures of knowledge* [11].

We note the command interfaces that they provide in Fig. 10. \mathcal{F}_{NIC} allows a party to commit to a message m via a commitment c that can be verified using the opening o , and \mathcal{F}_{SoK} allows a party to produce a signature σ on message m if they know the opening r to some commitment c (which commits to some value S) which has been accumulated in a with witness w . \mathcal{F}_{SoK} can internally use any functionality that represents a language of statements on which signatures of knowledge can be made. In our case, it uses both \mathcal{F}_{acc} and \mathcal{F}_{NIC} . The model of [11] is easily extended to the case of two “internal” functionalities, since we can

define the language that our \mathcal{F}_{SoK} accepts as $\mathcal{L} = \{(a, S, c, w, r) : (w, \{c\}, a) \in \mathcal{L}_{\text{acc}} \wedge (c, S, r) \in \mathcal{L}_{\text{NIC}}\}$, where \mathcal{L}_{acc} , \mathcal{L}_{NIC} are the languages accepted by \mathcal{F}_{acc} and \mathcal{F}_{NIC} , respectively.

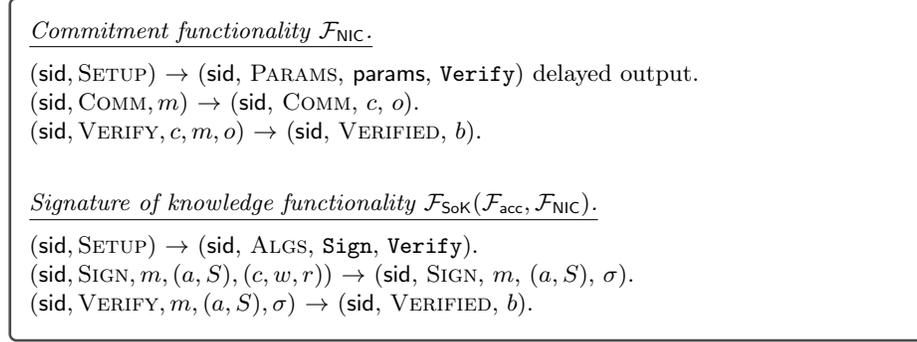


Fig. 10. Interfaces provided by \mathcal{F}_{NIC} and $\mathcal{F}_{\text{SoK}}(\mathcal{F}_{\text{acc}}, \mathcal{F}_{\text{NIC}})$.

Fig. 11 defines an eligibility protocol that realizes $\mathcal{F}_{\text{elig}}$. For ease of notation, when we say P calls \mathcal{F}_{acc} or \mathcal{F}_{NIC} , we mean that the communication goes through \mathcal{F}_{SoK} since the functionalities are embedded within. [11] describes formally how this is achieved. Π_{elig} also uses the functionalities for anonymous broadcast, the global clock and the certification functionality instantiated for SA and all V which we denote by $\{\mathcal{F}_{\text{cert}}^P\}$, where $P \in \{V, \text{SA}\}$. Note that when we write $\mathcal{F}_{\text{cert}}^P$ we mean $\mathcal{F}_{\text{cert}}(P, \mathbf{V} \cup \text{SA})$.

Remark 2. Note that in the definitions of $\mathcal{F}_{\text{elig}}$ and Π_{elig} , we implicitly assume that the commands with which the environment can activate the corrupted parties are the same as for the honest parties. However, for corrupted parties, $\mathcal{F}_{\text{elig}}$ will first reach out to the simulator to request their internal state, which the simulator will provide (this part is captured explicitly), and then $\mathcal{F}_{\text{elig}}$ will output as usual based on what \mathcal{S} provides.

Theorem 3. Π_{elig} UC-realizes $\mathcal{F}_{\text{elig}}$ in the $\{\mathcal{F}_{\text{SoK}}(\mathcal{F}_{\text{acc}}, \mathcal{F}_{\text{NIC}}), \mathcal{F}_{\text{an.BC}}, \{\mathcal{F}_{\text{cert}}^P\}, \mathcal{G}_{\text{clock}}\}$ -hybrid model.

Proof. In a direct proof of UC-realizability, given a real adversary \mathcal{A} we have to construct a simulator \mathcal{S} that can make the ideal execution indistinguishable from the real one from the point of view of the environment \mathcal{Z} . Since we are in a hybrid model, \mathcal{S} will internally simulate the real execution for \mathcal{A} by playing the part of voters and authorities as well as the part of the hybrid functionalities. In case of corrupted voters, it will pass any requests from \mathcal{A} onto $\mathcal{F}_{\text{elig}}$.

The eligibility protocol $\Pi_{\text{elig}}^{\mathcal{F}_{\text{SoK}}(\mathcal{F}_{\text{acc}}, \mathcal{F}_{\text{NIC}}), \mathcal{F}_{\text{an.BC}}, \{\mathcal{F}_{\text{cert}}^P\}, \mathcal{G}_{\text{clock}}}$.

All parties initialize the set $C \leftarrow \emptyset$. Moreover, all parties have hard coded the predicates (`define_time`, `Status`) If at any point a hybrid functionality aborts, the party returns \perp to \mathcal{Z} .

■ Upon receiving $(\text{sid}, \text{ELIGIBLE}, \mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}})$ from \mathcal{Z} , if $\mathbf{V}_{\text{elig}} \subseteq \mathbf{V}$ and SA's `status=init`, SA runs:

1. It sets its status to 'credential'.
2. It computes $\mathbf{t} \leftarrow \text{define_time}(t_{\text{cast}}, t_{\text{open}})$.
3. It sends $(\text{sid}, \text{SETUP})$ to \mathcal{F}_{acc} to get $(\text{sid}, \text{ALGS}, \text{params}_{\text{acc}}, f, \text{Verify}_{\text{acc}}, \text{Delete}, \text{Update})$. It parses $\text{params}_{\text{acc}}$ as $(\mathcal{X}_\lambda, \mathcal{U}_f, u)$. It sends $(\text{sid}, \text{SETUP})$ to \mathcal{F}_{NIC} to get $(\text{sid}, \text{PARAMS}, \text{params}_{\text{NIC}}, \text{Verify}_{\text{NIC}})$. It parses $\text{params}_{\text{NIC}}$ as $(\mathcal{M}, \mathcal{R})$. It sends $(\text{sid}, \text{SETUP})$ to \mathcal{F}_{SoK} to get $(\text{sid}, \text{ALGS}, \text{Sign}, \text{Verify}_{\text{SoK}})$. It saves all algorithms and parameters except `Delete` in St_{gen} .
4. It sets $\text{reg.par} := (\mathbf{V}_{\text{elig}}, t_{\text{cast}}, t_{\text{open}}, \mathbf{t}, St_{\text{gen}})$.
5. It sends $((\text{SA}, \text{sid}), \text{SETUP})$ to $\mathcal{F}_{\text{cert}}^{\text{SA}}$, waits for confirmation and then sends $((\text{SA}, \text{sid}), \text{SIGN}, (\text{SA}, \text{reg.par}))$ back to receive $((\text{SA}, \text{sid}), \text{SIGNATURE}, (\text{SA}, \text{reg.par}), s)$. Then it sends $(\text{sid}, \text{BROADCAST}, (\text{SA}, \text{reg.par}, s))$ to $\mathcal{F}_{\text{an.BC}}$.
6. It returns $(\text{sid}, \text{ELIG.PAR}, \text{reg.par})$ to \mathcal{Z} .

■ Upon receiving $(\text{sid}, \text{BROADCAST}, (\text{SA}, \text{reg.par}, s))$ from $\mathcal{F}_{\text{an.BC}}$, V sends $((\text{SA}, \text{sid}), \text{VERIFY}, (\text{SA}, \text{reg.par}), s)$ to $\mathcal{F}_{\text{cert}}^{\text{SA}}$. If it returns 1, she stores `reg.par` and sets her status to 'credential'.

■ Upon receiving $(\text{sid}, \text{GEN.CRED})$ from \mathcal{Z} , if V 's `status=credential`, she reads `Cl` from $\mathcal{G}_{\text{clock}}$. If $\text{Status}(\text{Cl}, \mathbf{t}, \text{Cast}) = \top$, she sets her status to 'cast'. Otherwise, if $\text{Status}(\text{Cl}, \mathbf{t}, \text{Cred}) = \top$, then V acts as follows:

1. She computes $S \stackrel{\$}{\leftarrow} \mathcal{M}$ and sends $(\text{sid}, \text{COMM}, S)$ to \mathcal{F}_{NIC} to get $(\text{sid}, \text{COMM}, c, r)$. If $c \notin \mathcal{X}_\lambda$, it repeats this step until it does.
2. She sends $((V, \text{sid}), \text{SETUP})$ to $\mathcal{F}_{\text{cert}}^V$, waits for confirmation and then sends $((V, \text{sid}), \text{SIGN}, (V, c))$ back to receive $((V, \text{sid}), \text{SIGNATURE}, (V, c), s)$. Then it sends $(\text{sid}, \text{BROADCAST}, (V, (c, s)))$ to $\mathcal{F}_{\text{an.BC}}$.

■ Upon receiving $(\text{sid}, \text{BROADCAST}, (V', (c', s')))$ from $\mathcal{F}_{\text{an.BC}}$, V sends $((V', \text{sid}), \text{VERIFY}, (V', c'), s')$ to $\mathcal{F}_{\text{cert}}^{V'}$. If it returns 1, it adds c' to C and returns $(\text{sid}, \text{GEN.CRED}, V', c')$ to \mathcal{Z} .

■ Upon receiving $(\text{sid}, \text{AUTH.BALLOT}, v)$ from \mathcal{Z} , if V 's `status=cast`, it reads `Cl` from $\mathcal{G}_{\text{clock}}$. If $\text{Status}(\text{Cl}, \mathbf{t}, \text{Open}) = \top$, it sets its status to 'open'. Otherwise, if $\text{Status}(\text{Cl}, \mathbf{t}, \text{Cast}) = \top$, then V :

1. She sends $(\text{sid}, \text{ACCUM}, u, C \setminus \{c\})$ to \mathcal{F}_{acc} to get $(\text{sid}, \text{ACCUM}, w)$, and $(\text{sid}, \text{ACCUM}, w, \{c\})$ to get $(\text{sid}, \text{ACCUM}, a)$.
2. She sends $(\text{sid}, \text{SIGN}, v, (a, S), (c, w, r))$ to \mathcal{F}_{SoK} to get $(\text{sid}, \text{SIGN}, v, (a, S), \phi)$. She sets $\sigma := (\phi, S)$.
3. She returns $(\text{sid}, \text{AUTH.BALLOT}, v, \sigma)$ to \mathcal{Z} .

■ Upon receiving $(\text{sid}, \text{VER.BALLOT}, v, \sigma)$ from \mathcal{Z} , V :

1. She parses σ as (ϕ, S) and sends $(\text{sid}, \text{ACCUM}, u, C)$ to \mathcal{F}_{acc} to get $(\text{sid}, \text{ACCUM}, a)$ and $(\text{sid}, \text{VERIFY}, v, (a, S), \phi)$ to \mathcal{F}_{SoK} to get $(\text{sid}, \text{VERIFIED}, x)$.
2. She returns $(\text{sid}, \text{VER.BALLOT}, v, \sigma, x)$ to \mathcal{Z} .

■ Upon receiving $(\text{sid}, \text{LINK.BALLOTS}, (v, \sigma), (v', \sigma'))$ from \mathcal{Z} , V :

1. She parses σ as (ϕ, S) and σ' as (ϕ', S') . If $S = S'$, she sets $x = 1$, otherwise $x = 0$.
2. She returns $(\text{sid}, \text{LINK.BALLOTS}, (v, \sigma), (v', \sigma'), x)$ to \mathcal{Z} .

Fig. 11. Definition of Π_{elig} in the $\{\mathcal{F}_{\text{SoK}}(\mathcal{F}_{\text{acc}}, \mathcal{F}_{\text{NIC}}), \mathcal{F}_{\text{an.BC}}, \{\mathcal{F}_{\text{cert}}^P\}, \mathcal{G}_{\text{clock}}\}$ -hybrid model in the presence of \mathbf{V} and SA.

Since both $\mathcal{F}_{\text{elig}}$ and all the functionalities we are working with expect the adversary to provide algorithms and parameters during setup, \mathcal{S} will have to ask \mathcal{A} for the algorithms for \mathcal{F}_{acc} , \mathcal{F}_{NIC} and \mathcal{F}_{SoK} and use them to define the algorithms that $\mathcal{F}_{\text{elig}}$ expects, as shown in Fig. 12.

We do not describe the reasoning behind the hybrid functionalities' algorithms here, as we are referring to the definitions established in [5] and [11] respectively for \mathcal{F}_{NIC} and \mathcal{F}_{SoK} . TrapCom , TrapOpen and $\text{Verify}_{\text{NIC}}$ belong to \mathcal{F}_{NIC} and SimSign , $\text{Verify}_{\text{SoK}}$ and Extract are from \mathcal{F}_{SoK} , and $\text{Verify}_{\text{cert}}$ and $\text{Sign}_{\text{cert}}$ are given to all $\mathcal{F}_{\text{cert}}^P$ (which maybe different for each party). The rest have been explained in the context of \mathcal{F}_{acc} in Subsection 4.1.

GenCred($1^\lambda, \text{reg.par}$):

1. $(\text{comm}, \text{info}) \leftarrow \text{TrapCom}(\text{sid}, \text{params}_{\text{NIC}}, \text{trapdoor})$.
2. Let $(\mathcal{M}, \mathcal{R}) \leftarrow \text{params}_{\text{NIC}}$ and compute $S \xleftarrow{\$} \mathcal{M}$.
3. $\text{open} \leftarrow \text{TrapOpen}(\text{sid}, S, \text{info})$.
4. If $\text{Verify}_{\text{NIC}}(\text{params}_{\text{NIC}}, \text{comm}, S, \text{open}) \neq \top$, return \perp .
5. Return $(\text{cr}, \text{rc}) := ((S, \text{open}), (\text{comm}))$.

UpState(St_{gen}, C):

1. Return $St_{\text{fin}} := C$.

AuthBallot($v, \text{cr}, \text{rc}, St_{\text{fin}}, \text{reg.par}$):

1. Let $(S, r) := \text{cr}$, $(c, s) := \text{rc}$ and $C := St_{\text{fin}}$.
2. Compute $w \leftarrow f(u, C \setminus \{c\})$ and $a \leftarrow f(w, \{c\})$.
3. Compute $a' \leftarrow f(u, C)$. If $a' \neq a$, return \perp .
4. If $\text{Verify}_{\text{acc}}(u, C \setminus \{c\}, w) \neq \top$, $\text{Verify}_{\text{acc}}(w, \{c\}, a) \neq \top$ or $\text{Verify}_{\text{acc}}(u, C, a) \neq \top$, return \perp .
5. Else compute $\phi \leftarrow \text{SimSign}(\mathcal{L}, v, (a, S))$.
6. If $\text{Verify}_{\text{SoK}}(\mathcal{L}, v, (a, S), \phi) \neq \top$, return \perp .
7. Else return $\sigma := (\phi, S)$.

VrfyBallot($v, \sigma, St_{\text{fin}}, \text{reg.par}$):

1. Let $C := St_{\text{fin}}$ and $(\phi, S) := \sigma$. Compute $a \leftarrow f(u, C)$.
2. If $\text{Verify}_{\text{acc}}(u, C, a) \neq \top$, return \perp .
3. Else compute $(c, w, r) \leftarrow \text{Extract}(\mathcal{L}, v, (a, S), \phi)$.
4. If $\text{Verify}_{\text{acc}}(w, \{c\}, a) \neq \top$ or $\text{Verify}_{\text{NIC}}(\text{params}_{\text{NIC}}, c, S, r) \neq \top$, return \perp .
5. Else return $x \leftarrow \text{Verify}_{\text{SoK}}(\mathcal{L}, v, (a, S), \phi)$.

Fig. 12. Algorithms supplied by \mathcal{S} to $\mathcal{F}_{\text{elig}}$ during Setup, using algorithms supplied by \mathcal{A} to \mathcal{F}_{acc} , \mathcal{F}_{NIC} and \mathcal{F}_{SoK} .

Then \mathcal{S} continues observing the ideal execution as allowed by $\mathcal{F}_{\text{elig}}$ and its task is to use this to play the part of real voters in the simulated execution with \mathcal{A} . Using the commitment algorithms obtained earlier by \mathcal{S} , $\mathcal{F}_{\text{elig}}$ can generate valid credentials for all eligible voters who have been activated in the ideal execution. On the other hand, in the ideal execution when \mathcal{S} is asked by

$\mathcal{F}_{\text{elig}}$ if it allows the credential to be generated via public delayed output, \mathcal{S} signs the credential and checks the signature with the algorithms provided by \mathcal{A} . Note that in the ideal execution there are not any signatures involved with each voter’s credential because we assume that $\mathcal{F}_{\text{elig}}$ broadcasts and authenticates the public part of the credential at the same time. If the verification fails, \mathcal{S} does not allow the generation of the credential. Similarly, when \mathcal{S} is asked by $\mathcal{F}_{\text{elig}}$ if it allows to send the credential via public delay output to the other parties, \mathcal{S} asks \mathcal{A} if it allows the broadcast as if it was from $\mathcal{F}_{\text{an,BC}}$ and responds appropriately to $\mathcal{F}_{\text{elig}}$. Note that \mathcal{S} has no visibility into the interaction between voters and $\mathcal{F}_{\text{elig}}$ with respect to AUTH_BALLOT and VER_BALLOT commands, just like \mathcal{A} would not in the real protocol. That is not the case for \mathcal{Z} , as it can activate parties at will and supplies their inputs, so indistinguishability of the outputs of those commands is ensured by the matching of the algorithms described earlier. In this way, properties guaranteed by $\mathcal{F}_{\text{elig}}$ directly translate to (one or several) properties provided by the hybrid functionalities, e.g. eligibility relies on binding commitments and ballot unforgeability depends on unforgeability of signatures of knowledge.

A detailed construction of \mathcal{S} is given in Supplementary Material A.2. \square

5 Realizing \mathcal{F}_{vm} via TLE

We define one of our main building blocks, the \mathcal{F}_{TLE} functionality that captures the security properties of a time-lock encryption (TLE) scheme. Intuitively speaking, a TLE scheme is a pair of algorithms $(e, d) := (\text{TL.Enc}, \text{TL.Dec})$ that a party P_i can use in order to encrypt a message m with time label τ_{dec} such that everyone can decrypt that message after the current time exceeds τ_{dec} . The decryption is possible because after the time τ_{dec} , a witness w_{dec} is available. According to [25], the production of that witness is the responsibility of a *computational reference clock (CRC)* which at time τ_{dec} produces the witness and provides it to all the parties upon request.

In Subsection 5.1, we provide a definition of the *ideal TLE functionality* \mathcal{F}_{TLE} . In Subsection 5.2, we construct a real-world protocol that UC-realizes the vote management functionality \mathcal{F}_{vm} (cf. Subsection 3.3) via \mathcal{F}_{TLE} .

5.1 Definition of \mathcal{F}_{TLE}

We provide the first UC treatment of TLE by defining the functionality \mathcal{F}_{TLE} , following the approach of [6,8]. The functionality is described in Fig. 13, and at a high level operates as follows.

Initially, the simulator \mathcal{S} provides \mathcal{F}_{TLE} with the corrupted voter set and the encryption/decryption algorithms. Each time an encryption query issued by an uncorrupted party is given to \mathcal{F}_{TLE} , the functionality records the message and the encryption *not of the message*, but of the all zero string and returns the ciphertext. This illustrates the fact that the ciphertext itself does not contain any information about the message. Next, it handles the decryption queries

The time-lock encryption functionality $\mathcal{F}_{\text{TLE},1\lambda}^{\text{leak}}$.

The functionality initializes the lists of algorithms L_{algo} and ready voters L_{ready} as empty.

- Upon receiving $(\text{sid}, \text{CORRUPT}, \mathbf{P}_{\text{corr}})$ from \mathcal{S} , it records the corrupted set \mathbf{P}_{corr} .
- Upon receiving $(\text{sid}, \text{START})$ from a party $P \notin L_{\text{ready}}$, it adds $P \rightarrow L_{\text{ready}}$ and:
 1. If $L_{\text{algo}} \equiv \emptyset$, it forwards the message to \mathcal{S} . Upon receiving $(\text{sid}, \text{ENC}, \text{DEC}, e_{\mathcal{S}}, d_{\mathcal{S}})$ from \mathcal{S} , it registers the tuple $(e_{\mathcal{S}}, d_{\mathcal{S}}) \rightarrow L_{\text{algo}}$. It returns $(\text{sid}, \text{START_OK})$ to P .
 2. Else it returns $(\text{sid}, \text{START_OK})$ to P .
- Upon receiving $(\text{sid}, \text{ENC}, m, \tau_{\text{dec}})$ from $P \notin \mathbf{P}_{\text{corr}}$:
 1. If $\tau_{\text{dec}} < 0$ or $P \notin L_{\text{ready}}$, it returns $(\text{sid}, \text{ENC}, m, \tau_{\text{dec}}, \perp)$ to P .
 2. Else, it adds $(m, c \leftarrow e_{\mathcal{S}}(0^{|m|}, \tau_{\text{dec}}; r), \tau_{\text{dec}}, r) \rightarrow L_{\text{rec}}$, where $r \xleftarrow{\$} \{0, 1\}^{\text{poly}(\lambda)}$ is the randomness of encryption, and returns $(\text{sid}, \text{ENC}, m, \tau_{\text{dec}}, c)$ to P .
- Upon receiving $(\text{sid}, \text{DEC}, c, \tau)$ from $P \notin \mathbf{P}_{\text{corr}}$:
 1. If $\tau < 0$ or $P_i \notin L_{\text{ready}}$, it returns $(\text{sid}, \text{DEC}, c, \tau, \perp)$ to P . Else, it reads the time Cl from $\mathcal{G}_{\text{clock}}$ and:
 - (a) If $\text{Cl} < \tau$, it sends $(\text{sid}, \text{DEC}, c, \tau, \text{MORE_TIME})$ to P .
 - (b) If $\text{Cl} \geq \tau$, then
 - If there are two tuples $(m_1, c, \tau_1, r_1), (m_2, c, \tau_2, r_2)$ in L_{rec} such that $m_1 \neq m_2$ where $\tau \geq \max\{\tau_1, \tau_2\}$, it returns to P $(\text{sid}, \text{DEC}, c, \tau, \perp)$.
 - If no tuple (\cdot, c, \cdot, \cdot) is recorded in L_{rec} , it sends $(\text{sid}, \text{DEC}, c, \tau)$ to \mathcal{S} and returns to P whatever it receives from \mathcal{S} .
 - If there is a unique tuple $(m, c, \tau_{\text{dec}}, r)$ in L_{rec} , then if $\tau \geq \tau_{\text{dec}}$, it returns $(\text{sid}, \text{DEC}, c, \tau, m)$ to P . Else, if $\text{Cl} < \tau_{\text{dec}}$, it returns $(\text{sid}, \text{DEC}, c, \tau, \text{MORE_TIME})$ to P . Else, if $\text{Cl} \geq \tau_{\text{dec}} > \tau$, it returns $(\text{sid}, \text{DEC}, c, \tau, \text{INVALID_TIME})$ to P .
- Whatever message it receives from $P \in \mathbf{P}_{\text{corr}}$, it forwards it to \mathcal{S} and vice versa.
- Upon receiving $(\text{sid}, \text{LEAKAGE})$ from \mathcal{S} , it reads the time Cl from $\mathcal{G}_{\text{clock}}$ and returns $(\text{sid}, \text{LEAKAGE}, \{(m, c, \tau_{\text{dec}})\}_{\tau_{\text{dec}}: \tau_{\text{dec}} \leq \text{leak}(\text{Cl})})$ to \mathcal{S} .

Fig. 13. Functionality \mathcal{F}_{TLE} parameterized by λ , a leakage function leak , interacting with simulator \mathcal{S} , parties in \mathbf{P} , and global clock $\mathcal{G}_{\text{clock}}$.

appropriately, unless it finds two messages recorded with the same ciphertext, in which case it outputs \perp . This implies that the encryption/decryption algorithms should satisfy correctness. In the cases of encryption/decryption queries issued by corrupted parties, \mathcal{F}_{TLE} responds according to the instructions of \mathcal{S} . Moreover, on demand, \mathcal{F}_{TLE} gives the record of all messages with encryption up to $\text{leak}(\text{Cl})$ to \mathcal{S} , where Cl is the current time and leak a leakage function that takes as an input a time and gives as an output a time of at least higher value. This function leak captures the fact that in some cases the adversary can decrypt messages before the time comes. Ideally the “best” leak function with respect to security is the identity one, the one that gives no real advantage to the adversary.

Unfortunately, there are some time lock encryption schemes where the adversary can decrypt a little bit earlier than the honest voters. For example time lock encryption that is based on bitcoin. This is happening because the adversary can locally compute some witnesses (selfish mining) without announce them to the rest of the parties. In the context of e-voting, it means that the adversary for these time lock instantiations can decrypt the ballots before the tally phase, learn the election outcome and change his vote (violation of fairness). Fortunately, even these time lock encryption schemes are useful in voting if we introduce a “time wait” period, that’s it a period where none can cast a ballot and ballots can not be opened either. In that period, the adversary can open the ballot but he can’t change his vote because the casting period is over. That’s the reason why we defined the predicates `define_time`, `Status`. As we see next, these predicates are instantiated with the leakage function in order our \mathcal{F}_{TLE} to be able to realise \mathcal{F}_{vm} .

Candidate for time lock encryption scheame. A candidate UC realization of our ideal functionality \mathcal{F}_{TLE} could be achieved by using the constructions of Liu’s *et al.* [25] which are based on witness encryption [18] or Shamir’s *et al.* time-lock puzzles [30]. Of course, many challenges arise when trying to transfer these constructions to the UC setting.

One of these challenges is to capture the limited resources of an adversary. Specifically, the concept of time-lock encryption assumes an adversary that executes a specific number of steps. Then, the main argument is that for these adversaries, it is hard to “break” the properties of a time-lock encryption scheme. On the other hand, in the UC framework, the adversary and the environment can execute any number of polynomial steps with respect to a security parameter. We can tackle this difficulty by modeling the desired computation the adversary does in order to “break” the time-lock encryption scheme as an ideal functionality. For example, if the desired computation is a hash evaluation of a value x , we can model this computation as the random oracle functionality, namely \mathcal{F}_{RO} . Then, we can restrict the access to that functionality by introducing a *functionality wrapper* in the spirit of [2]. The functionality that is responsible for these computations is local, which means that the environment doesn’t have direct access to it. So any computation request must be done via the adversary. On the other hand, the access of the adversary is restricted by the functionality wrapper.

As it can been seen in our \mathcal{F}_{TLE} definition, the encryption of honestly generated ciphertexts is not based on the actual messages but instead on zero string messages. This is done in order to capture semantic security of the encryption algorithm. Hence, no information about the actual plaintext can be revealed from the ciphertext. On the other hand, when a certain amount of time has passed, these ciphertexts must be opened to the actual messages. As a result, any construction of time-lock encryption that realizes our functionality must allow a simulator to “equivocate” the messages. For this reason, we can extend every construction of time-lock encryption scheme to a time-lock non commit-

ting encryption scheme by using techniques similar to [28] in the random oracle model.

Lastly, a computational reference clock [25] must be introduced in order to realize our \mathcal{F}_{TLE} . This computational reference clock can be modeled as an ideal functionality, namely \mathcal{F}_{CRC} , and has the following properties: i) it produces witnesses after a constant period of time by reading $\mathcal{G}_{\text{clock}}$, ii) an adversary cannot outperform it. That means that an adversary cannot produce witnesses faster than \mathcal{F}_{CRC} in a specific amount of time which can be defined as a round. In turn, \mathcal{F}_{CRC} could be realized with a construction based on Bitcoin [2] by using a collection of blocks as a witness. We leave the concrete and formal UC-realization of \mathcal{F}_{TLE} as future work.

5.2 A protocol $\Pi_{\text{vm}}^{\mathcal{F}_{\text{TLE}}, \{\mathcal{F}_{\text{cert}}^P, \mathcal{F}_{\text{an.BC}}, \mathcal{G}_{\text{clock}}\}}$ that realizes \mathcal{F}_{vm}

We present the description of the protocol $\Pi_{\text{vm}}^{\mathcal{F}_{\text{TLE}}, \{\mathcal{F}_{\text{cert}}^P, \mathcal{F}_{\text{an.BC}}, \mathcal{G}_{\text{clock}}\}}$ and we provide a proof that it UC-realizes \mathcal{F}_{vm} in the $\{\mathcal{F}_{\text{TLE}}, \{\mathcal{F}_{\text{cert}}^P\}, \mathcal{F}_{\text{an.BC}}, \mathcal{G}_{\text{clock}}\}$ -hybrid model. The description of $\Pi_{\text{vm}}^{\mathcal{F}_{\text{TLE}}, \{\mathcal{F}_{\text{cert}}^P, \mathcal{F}_{\text{an.BC}}, \mathcal{G}_{\text{clock}}\}}$ shown in Fig. 14 follows the phases and the command interface of \mathcal{F}_{vm} in Subsection 3.3.

Definition 3. Let leak be a leakage function and $t_{\text{cast}}, t_{\text{open}}$ be time points. We say that the pair of functions $\text{Status}_{\text{leak}}, \text{define_time}_{\text{leak}}$ with respect to a leakage function leak is *phase preserving* if the following property holds:

$$\forall \text{Cl such that } \text{Status}_{\text{leak}}(\text{Cl}, t, \text{Cast}) = \top \Rightarrow \text{Status}_{\text{leak}}(\text{leak}(\text{Cl}), t, \text{Open}) = \perp,$$

where $t \leftarrow \text{define_time}_{\text{leak}}(t_{\text{cast}}, t_{\text{open}})$

Theorem 4. Π_{vm} UC-realizes \mathcal{F}_{vm} in the $\{\mathcal{F}_{\text{TLE}}^{\text{leak}}, \mathcal{F}_{\text{an.BC}}, \{\mathcal{F}_{\text{cert}}^P\}, \mathcal{G}_{\text{clock}}\}$ -hybrid model given that the pair of functions $\text{Status}_{\text{leak}}, \text{define_time}_{\text{leak}}$ is *phase preserving*.

Proof. In cases where a corrupted party receives an input and we do not describe her behaviour, we assume that the message is sent to \mathcal{S} from \mathcal{F}_{vm} and \mathcal{S} forwards that message to \mathcal{A} as if it was from that party. Then \mathcal{S} returns to \mathcal{F}_{vm} whatever receives from \mathcal{A} .

We describe the ideal adversary \mathcal{S} . When \mathcal{S} receives the corruption vector from \mathcal{Z} , \mathcal{S} forwards it to \mathcal{A} as if it was from \mathcal{Z} . When \mathcal{S} receives back the corruption vector from \mathcal{A} playing the role of both of $\mathcal{F}_{\text{TLE}}, \mathcal{F}_{\text{cert}}^{\text{SA}}$, \mathcal{S} forwards it to \mathcal{F}_{vm} . When \mathcal{S} receives the setup information $(\text{sid}, \text{SETUP_INFO}, t_{\text{cast}}, t_{\text{open}}, \mathbf{V}_{\text{elig}})$ from \mathcal{F}_{vm} , \mathcal{S} sends $((\text{SA}, \text{sid}), \text{SETUP})$ to \mathcal{A} as if it was from $\mathcal{F}_{\text{cert}}^{\text{SA}}$. Upon receiving $((\text{SA}, \text{sid}), \text{ALGORITHMS}, \text{Verify}, \text{Sign})$ from \mathcal{A} playing the role of $\mathcal{F}_{\text{cert}}^{\text{SA}}$, \mathcal{S} stores the algorithms $(\text{Verify}, \text{Sign})$ and produces the signature $\sigma \leftarrow \text{Sign}(t_{\text{cast}}, t_{\text{open}}, \mathbf{V}_{\text{elig}})$. Then \mathcal{S} asks \mathcal{A} if it allows the cast of the message $(t_{\text{cast}}, t_{\text{open}}, \mathbf{V}_{\text{elig}}, \sigma)$ as if it was from $\mathcal{F}_{\text{an.BC}}$.

Then, \mathcal{S} responds to \mathcal{F}_{vm} according to the answer of \mathcal{A} .

The vote management protocol $\Pi_{\text{vm}}^{\mathcal{F}_{\text{TLE}}^{\text{leak}}, \mathcal{F}_{\text{an.BC}}, \{\mathcal{F}_{\text{cert}}^i\}, \mathcal{G}_{\text{clock}}}$.

Each voter V maintains a list of the cast ballots L_{cast}^V initially as empty. Each voter initializes its status to 'exec'.

■ Upon receiving $(\text{sid}, \text{SETUP_INFO}, t_{\text{cast}}, t_{\text{open}}, \mathbf{V}_{\text{elig}})$ for the first time from \mathcal{Z} , SA sends $((\text{SA}, \text{sid}), \text{SETUP})$ to $\mathcal{F}_{\text{cert}}^{\text{SA}}$, and waits to receive $((\text{SA}, \text{sid}), \text{SETUP})$ from $\mathcal{F}_{\text{cert}}^{\text{SA}}$. Then SA computes $t \leftarrow \text{define_time}_{\text{leak}}(t_{\text{cast}}, t_{\text{open}})$ and if $t \neq \perp$ it sends $((\text{SA}, \text{sid}), \text{SIGN}, (\text{SA}, t_{\text{cast}}, t_{\text{open}}, t, \mathbf{V}_{\text{elig}}))$ to $\mathcal{F}_{\text{cert}}^{\text{SA}}$. Upon receiving $((\text{SA}, \text{sid}), \text{SIGNATURE}, (\text{SA}, t_{\text{cast}}, t_{\text{open}}, t, \mathbf{V}_{\text{elig}}), s)$ from $\mathcal{F}_{\text{cert}}^{\text{SA}}$, it sends $(\text{sid}, \text{BROADCAST}, (\text{SA}, t_{\text{cast}}, t_{\text{open}}, t, \mathbf{V}_{\text{elig}}, s))$ to $\mathcal{F}_{\text{an.BC}}$.

■ Upon receiving $(\text{sid}, \text{BROADCAST}, (\text{SA}, t_{\text{cast}}, t_{\text{open}}, t, \mathbf{V}_{\text{elig}}, s))$ from $\mathcal{F}_{\text{an.BC}}$, V sends $((\text{SA}, \text{sid}), \text{VERIFY}, (\text{SA}, t_{\text{cast}}, t_{\text{open}}, t, \mathbf{V}_{\text{elig}}, s))$ to $\mathcal{F}_{\text{cert}}^{\text{SA}}$. If it returns 1, she stores $t_{\text{cast}}, t_{\text{open}}, t, \mathbf{V}_{\text{elig}}$.

■ Upon receiving $(\text{sid}, \text{START})$ from \mathcal{Z} , V forwards the message to \mathcal{F}_{TLE} . Upon receiving $(\text{sid}, \text{START_OK})$ from \mathcal{F}_{TLE} , V outputs the received message.

■ Upon receiving $(\text{sid}, \text{GEN_BALLOT}, o)$ from \mathcal{Z} , V does:

1. If this is the first time receiving this command-message and $V \in \mathbf{V}_{\text{elig}}$, V sends to \mathcal{F}_{TLE} $(\text{sid}, \text{ENC}, o, t_{\text{open}})$. Upon receiving $(\text{sid}, \text{ENC}, o, t_{\text{open}}, v)$ from \mathcal{F}_{TLE} , V stores the pair $(V, v, o, 1)$ and returns the message $(\text{sid}, \text{GEN_BALLOT}, o, v)$ to \mathcal{Z} .
2. Else, V returns to \mathcal{Z} $(\text{sid}, \text{GEN_BALLOT}, o, \perp)$.

■ Upon receiving $(\text{sid}, \text{CAST}, v, \sigma)$ from \mathcal{Z} , if her status is 'exec', V reads the time Cl from $\mathcal{G}_{\text{clock}}$. If $\text{Status}_{\text{leak}}(\text{Cl}, t, \text{Open}) = \top$, V sets the status to 'open'. Otherwise V does:

1. If there is a tuple of the form $(V, v, \cdot, 1)$ stored and it is the first time receiving this command-message, V sends $(\text{sid}, \text{CAST}, v, \sigma)$ to $\mathcal{F}_{\text{an.BC}}$. Upon receiving $(\text{sid}, \text{CAST}, v, \sigma)$ from $\mathcal{F}_{\text{an.BC}}$, V^* stores the tuple (v, σ) to $L_{\text{cast}}^{V^*}$.
2. Else, V returns $(\text{sid}, \text{CAST}, v, \sigma, \perp)$ to \mathcal{Z} .

■ Upon receiving $(\text{sid}, \text{OPEN}, v^*)$ from \mathcal{Z} , if there is a tuple $(v^*, \sigma^*) \in L_{\text{cast}}^V$, V sends $(\text{sid}, \text{DEC}, v^*, t_{\text{open}})$ to \mathcal{F}_{TLE} .

1. Upon receiving $(\text{sid}, \text{DEC}, v^*, t_{\text{open}}, o^*)$ from \mathcal{F}_{TLE} , V returns the message $(\text{sid}, \text{OPEN}, v^*, o^*)$ to \mathcal{Z} .
2. Upon receiving $(\text{sid}, \text{DEC}, v^*, t_{\text{open}}, \perp)$ from \mathcal{F}_{TLE} , V returns the message $(\text{sid}, \text{OPEN}, v^*, \perp)$ to \mathcal{Z} .

Fig. 14. Definition of Π_{vm} in the $\{\mathcal{F}_{\text{TLE}}^{\text{leak}}, \mathcal{F}_{\text{an.BC}}, \{\mathcal{F}_{\text{cert}}^i\}, \mathcal{G}_{\text{clock}}\}$ -hybrid model in the presence of \mathbf{V} and SA.

Next, if \mathcal{S} receives $(\text{sid}, \text{START})$ from \mathcal{F}_{vm} , \mathcal{S} forwards the message to \mathcal{A} as if it was from \mathcal{F}_{TLE} . Upon receiving the time lock algorithms $(\text{sid}, \text{ENC}, \text{DEC}, e_{\mathcal{A}}, d_{\mathcal{A}})$ from \mathcal{A} , \mathcal{S} returns the message $(\text{sid}, \text{SETUP_INFO}, \text{GenBallot} = e_{\mathcal{A}}, \text{OpenBallot} = d_{\mathcal{A}})$ to \mathcal{F}_{vm} .

Upon receiving a GEN_BALLOT request from \mathcal{F}_{vm} on behalf of a corrupted party V_i , \mathcal{S} forwards the message to \mathcal{A} as if it was from \mathcal{F}_{TLE} and it returns the response of \mathcal{A} to \mathcal{F}_{vm} . Upon receiving $(\text{sid}, \text{ALLOW_CAST}, v, \sigma)$ from \mathcal{F}_{vm} , \mathcal{S} asks

\mathcal{A} if it allows the broadcast of (v, σ) as if it was from $\mathcal{F}_{\text{an.BC}}$. If the broadcast is allowed, \mathcal{S} sends $(\text{sid}, \text{CAST_ALLOWED})$ to \mathcal{F}_{vm} . When \mathcal{S} receives a `CAST` request from \mathcal{F}_{vm} on behalf of a corrupted party V_i , it forwards the message to \mathcal{A} as if it was from $\mathcal{F}_{\text{an.BC}}$ and it returns the message it received from \mathcal{A} back to \mathcal{F}_{vm} .

Upon receiving $(\text{sid}, \text{OPEN}, v)$ from \mathcal{F}_{vm} (where v is a ballot not generated by \mathcal{F}_{vm}), \mathcal{S} sends $(\text{sid}, \text{DEC}, v, t_{\text{open}})$ to \mathcal{A} as if it was from \mathcal{F}_{TLE} . When \mathcal{S} receives $(\text{sid}, \text{DEC}, v, t_{\text{open}}, o)$ from \mathcal{A} , it returns the message $(\text{sid}, \text{OPEN}, v, o)$ to \mathcal{F}_{TLE} .

Upon receiving $(\text{sid}, \text{LEAKAGE})$ from \mathcal{Z} , \mathcal{S} forwards the message to \mathcal{A} as if it was from \mathcal{Z} . Upon receiving $(\text{sid}, \text{LEAKAGE})$ from \mathcal{A} , \mathcal{S} reads the time Cl from $\mathcal{G}_{\text{clock}}$. Then \mathcal{S} playing the role of \mathcal{F}_{TLE} returns to \mathcal{A} all the maliciously generated cipher texts with time labeling until time $\text{leak}(\text{Cl})$. If $\text{Status}_{\text{leak}}(\text{Cl}, \mathbf{t}, \text{Open}) = \text{Status}_{\text{leak}}(\text{Cl}, \mathbf{t}, \text{Cast}) = \text{Status}_{\text{leak}}(\text{Cl}, \mathbf{t}, \text{Cred}) = \perp$, then \mathcal{S} can request and give also the honest generated cipher texts from \mathcal{F}_{vm} and returns them to \mathcal{A} as if it was from \mathcal{F}_{TLE} . Note that all honest parties will use \mathcal{F}_{TLE} with the same time labeling for encryption requests and at most once. The only way simulation fails is when the $\text{Status}_{\text{leak}}(\text{Cl}, \mathbf{t}, \text{Cast}) = \top$ and $\text{Status}_{\text{leak}}(\text{leak}(\text{Cl}), \mathbf{t}, \text{Open}) = \top$ because in that case \mathcal{S} can not retrieve the honestly generated plain texts from \mathcal{F}_{vm} so it can't give a response on behalf of \mathcal{F}_{TLE} to \mathcal{A} on a $(\text{sid}, \text{LEAKAGE})$ command. By the definition of $\text{Status}_{\text{leak}}$ and $\text{define_time}_{\text{leak}}$ this is impossible to happen.

The distribution of messages is exactly the same in both the ideal and the hybrid setting, as the algorithms $\text{GenBallot} \equiv e_{\mathcal{A}}$ and $\text{OpenBallot} \equiv d_{\mathcal{A}}$ that are used are the same. As a result the simulation is perfect. \square

References

1. Ben Adida. Helios: Web-based open-audit voting. In *USENIX security symposium*, pages 335–348, 2008.
2. Christian Badertscher, Ueli Maurer, Daniel Tschudi, and Vassilis Zikas. Bitcoin as a transaction ledger: A composable treatment. In *CRYPTO*, pages 324–356, 2017.
3. Foteini Baldimtsi, Ran Canetti, and Sophia Yakoubov. Universally composable accumulators. *IACR Cryptology ePrint Archive*, 2018:1241, 2018.
4. Jonathan Bannet, David W. Price, Algis Rudys, Justin Singer, and Dan S. Wallach. Hack-a-vote: Security issues with electronic voting systems. *IEEE Security & Privacy*, 2(1):32–37, 2004.
5. Jan Camenisch, Maria Dubovitskaya, and Alfredo Rial. UC commitments for modular protocol design and applications to revocation and attribute tokens. In *CRYPTO*, 2016.
6. Jan Camenisch, Anja Lehmann, Gregory Neven, and Kai Samelin. Uc-secure non-interactive public-key encryption. In *CSF 2017*, pages 217–233. IEEE Computer Society, 2017.
7. Jan Camenisch and Anna Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In *CRYPTO*, pages 61–76, 2002.
8. Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, 2001.
9. Ran Canetti. Universally composable signatures, certification and authentication. *IACR Cryptology ePrint Archive*, 2003:239, 2003.

10. Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. Universally composable security with global setup. In *TCC*, pages 61–85, 2007.
11. Melissa Chase and Anna Lysyanskaya. On signatures of knowledge. In *CRYPTO*, pages 78–96, 2006.
12. David Chaum. Blind signatures for untraceable payments. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *Advances in Cryptology*, pages 199–203, Boston, MA, 1983. Springer US.
13. Véronique Cortier, David Galindo, Ralf Küsters, Johannes Müller, and Tomasz Truderung. Sok: Verifiability notions for e-voting protocols. In *IEEE S&P*, 2016.
14. David Derler, Christian Hanser, and Daniel Slamanig. Revisiting cryptographic accumulators, additional properties and relations to other primitives. In *CT-RSA*, 2015.
15. David Derler, Christian Hanser, and Daniel Slamanig. Revisiting cryptographic accumulators, additional properties and relations to other primitives. In *CT-RSA*, 2015.
16. Saghar Estehghari and Yvo Desmedt. Exploiting the client vulnerabilities in internet e-voting systems: Hacking helios 2.0 as an example. In *EVT/WOTE*, 2010.
17. Nelly Fazio and Antonio Nicolosi. Cryptographic accumulators: Definitions, constructions and applications. *Paper written for course at New York University: www.cs.nyu.edu/nicolosi/papers/accumulators.pdf*, 2002.
18. Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In *STOC*, 2013.
19. Michael T. Goodrich, Roberto Tamassia, and Jasminka Hasic. An efficient dynamic and distributed cryptographic accumulator. In *ISC*, 2002.
20. Jens Groth. Efficient maximal privacy in boardroom voting and anonymous broadcast. In *FC*, 2004.
21. Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-resistant electronic elections. In *Proc. of the ACM workshop on privacy in the electronic society*, pages 61–70, 2005.
22. Aggelos Kiayias and Moti Yung. Self-tallying elections and perfect ballot secrecy. In *PKC*, 2002.
23. Aggelos Kiayias, Thomas Zacharias, and Bingsheng Zhang. End-to-end verifiable elections in the standard model. In *EUROCRYPT*, 2015.
24. Czesław Kościelny, Mirosław Kurkowski, and Marian Srebrny. *Foundations of Symmetric Cryptography*, pages 77–118. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
25. Jia Liu, Tibor Jager, Saqib A. Kakvi, and Bogdan Warinschi. How to build time-lock encryption. *Designs, Codes and Cryptography*, Feb 2018.
26. Ian Miers, Christina Garman, Matthew Green, and Aviel D. Rubin. Zerocoin: Anonymous distributed e-cash from bitcoin. In *IEEE S&P*, 2013.
27. Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, <http://bitcoin.org/bitcoin.pdf>.
28. Jesper Buus Nielsen. Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In *CRYPTO*, 2002.
29. Tatsuaki Okamoto. Receipt-free electronic voting schemes for large scale elections. In *Security Protocols*, 1998.
30. R. L. Rivest, A. Shamir, and D. A. Wagner. Time-lock puzzles and timed-release crypto. Technical report, Cambridge, MA, USA, 1996.
31. Peter Y. A. Ryan and Steve A. Schneider. Prêt-à-voter with re-encryption mixes. In *ESORICS*, 2006.

32. Alan Szepieniec and Bart Preneel. New techniques for electronic voting. Washington, D.C., 2015. USENIX Association.
33. Edward Tremel. Real-world performance of cryptographic accumulators. *Undergraduate Honors Thesis, Brown University*, 2013.
34. Scott Wolchok, Eric Wustrow, Dawn Isabel, and J. Alex Halderman. Attacking the washington, D.C. internet voting system. In *FC*, 2012.
35. Y. Zhang, J. Katz, and C. Papamanthou. An expressive (zero-knowledge) set accumulator. In *Euro S&E;P*, 2017.

A Supplementary material for Section 4

A.1 Proof of Theorem 2

The presented proofs for realizing \mathcal{F}_{acc} follow the general structure of the realizability proofs for the signature functionality in the 2005 version of [8]. We first prove a lemma that will be used to prove one direction of the equivalence.

Lemma 1. *If AC is an accumulator scheme and \mathcal{F}_{acc} is given parameters and algorithms according to AC, then \mathcal{F}_{acc} never aborts on a valid ACCUM query.*

Proof. Suppose we have a query $(\text{sid}, \text{ACCUM}, w, X)$ for $w \in \mathcal{U}'_f$, $X \subset \mathcal{X}'_\lambda$. f is deterministic and \mathcal{F}_{acc} only adds properly calculated values to L_{acc} , so \mathcal{F}_{acc} either returns $a = f(w, X)$ or aborts. Aborts can only occur in Step 6, so we show that they are never satisfied:

1. *Suppose there is $A \in L_a$ and $(a', L'_a) \in L_{\text{sets}}$ such that $A \in L'_a$ and $a \neq a'$.* After Step 5, A is one of three forms: $A = W \cup X$ where W is empty, or there is $(w, W) \in L_{\text{ref}}$ for some w , or there is $(w, [W_1, W_2, \dots]) \in L_{\text{sets}}$ with $W = W_i$ for some i . Further, we can write $W = W' \cup Y$ for some (possibly empty) fully-expanded set $Y = \{y_1, y_2, \dots\}$. Due to the way that items are added or modified within L_a , there can be at most one unexpanded witness reference in each, and in all cases X is added to each set. Since $A \in L'_a$, there must have been a sequence of queries which accumulated all elements of Y as well as X into a' , so that $f(w', Y \cup X) = a'$. From the current query we know that $f(w, X) = a$ and $f(w', Y) = w$, which implies that $a = f(w, X) = f(f(w', Y), X) = f(w', Y \cup X) = a'$, clearly a contradiction.

2. *Suppose $(a, W) \in L_{\text{ref}}$ and there is $A \in L_a$, $(a', L'_a) \in L_{\text{sets}}$ such that $W \cup B \in L'_a$ for some $B = \{b_1, b_2, \dots\}$ and $(a'', L''_a) \in L_{\text{sets}}$ such that $A \cup B \in L''_a$ and $a' \neq a''$.* $A \cup B \in L''_a$ implies that $f(w', Y \cup X \cup B) = a''$. The current query implies that $f(w, X)$ is the accumulator for W , and $W \cup B \in L'_a$ gives us $f(a, B) = a'$. Combining the last two we get $a' = f(f(w, X), B) = f(f(w', Y), X \cup B) = f(w', Y \cup X \cup B) = a''$, which is a contradiction.

No abort can occur, so (a, X, w) is recorded in L_{acc} and $(\text{sid}, \text{ACCUM}, a)$ is returned. \square

Theorem 5. *If $\text{AC} = (\mathcal{X}_\lambda, \text{Gen})$ is a secure dynamic accumulator scheme, then Π_{AC} UC-realizes \mathcal{F}_{acc} in the $\mathcal{F}_{\text{CRS}}^{\text{Gen}}$ -hybrid model.*

Proof. We prove the contrapositive, i.e. if Π_{AC} does not UC-realize \mathcal{F}_{acc} , then AC is not a secure dynamic accumulator scheme, so at least one property is not satisfied. Let Π stand for Π_{AC} and \mathcal{F} for the ideal protocol for \mathcal{F}_{acc} . By Definition 1 we know that there exist an adversary \mathcal{A} such that for every simulator \mathcal{S} there exists an environment \mathcal{Z} which can distinguish between \mathcal{F} and Π with non-negligible probability. So, for a specific simulator \mathcal{S} , as we define below, we use such environment \mathcal{Z} to construct an adversary \mathcal{G} against witness unforgeability because \mathcal{Z} must trigger a VERIFY query during its run if it is to succeed. We consider \mathcal{S} which supplies the correct algorithms from the AC scheme to \mathcal{F} .

We construct \mathcal{G} by simulating the ideal execution with \mathcal{F} and \mathcal{S} for \mathcal{Z} . Like \mathcal{S} , \mathcal{G} also runs a simulated \mathcal{A} . When \mathcal{Z} activates M or P with SETUP, \mathcal{G} returns algorithms as given by \mathcal{S} . \mathcal{G} answers all $(\text{sid}, \text{ACCUM}, w', X')$ calls with $(\text{sid}, \text{ACCUM}, f(w', X'))$. When \mathcal{Z} sends $(\text{sid}, \text{VERIFY}, v, Y, a)$, \mathcal{G} checks if $\text{Verify}_f(v, Y, a) = \top$ and one of the following holds:

1. There is $(a, L_a) \in L_{\text{sets}}$ such that $Y \not\subseteq A$ for some fully expanded $A \in L_a$, so that $x \in Y$ but $x \notin A$.

Let $w = f(v, Y \setminus \{x\})$ so $f(w, x) = f(v, Y) = a = f(u, A)$.

2. There are (v, L_w) and $(a, L_a) \in L_{\text{sets}}$ such that $A = W \cup V$ and $V \not\subseteq Y$ for some $A \in L_a, W \in L_w$ and some fully-expanded V , so that $x \in V$ but $x \notin Y$. Similarly, $w = f(v, V \setminus \{x\})$ so $f(w, x) = f(v, V) = a = f(v, Y) = f(u, Y \setminus W)$.

If it does, \mathcal{G} returns the forgery (x, w, X) for $X = A$ or $X = Y \setminus W$ and halts. If not, it continues the simulation. Let E be the event that \mathcal{Z} runs SETUP and a $(\text{sid}, \text{VERIFY}, v, Y, a)$ query such that $\text{Verify}(v, Y, a) = \top$ and a forgery condition holds. By Lemma 1, \mathcal{F} never aborts a well-formed ACCUM query and its outputs are consistent with Π . VERIFY is identical between \mathcal{F} and Π unless a forgery condition is triggered. Hence the view of \mathcal{Z} of the real execution is indistinguishable from the ideal one as long as E doesn't occur. We know that \mathcal{Z} can distinguish with non-negligible probability, E must occur with the same probability and so \mathcal{G} obtains a forgery. \square

Theorem 6. *If Π_{AC} UC-realizes \mathcal{F}_{acc} in the $\mathcal{F}_{\text{CRS}}^{\text{Gen}}$ -hybrid model, then $\text{AC} = (\mathcal{X}_\lambda, \text{Gen})$ is a secure dynamic accumulator scheme.*

Proof. We show the contrapositive. Lack of efficient generation or evaluation would fail SETUP, so we only consider commutativity, witness unforgeability and the dynamic procedures. In each case, we construct \mathcal{Z} that distinguishes for all \mathcal{S} , with an inactive \mathcal{A} . \mathcal{Z} always first obtains algorithms on behalf of M and P .

Suppose AC does not satisfy commutativity, so there are $f \in F_\lambda, u \in \mathcal{U}_f$ and $x_1, x_2 \in \mathcal{X}_\lambda$ such that $f(f(u, x_1), x_2) \neq f(f(u, x_2), x_1)$. \mathcal{Z} activates P with ACCUM for the following pairs of inputs and outputs: $(u, x_1) \rightarrow a, (u, x_2) \rightarrow a', (a, x_2) \rightarrow b$ and $(a', x_1) \rightarrow b'$. In the execution with \mathcal{F} , \mathcal{Z} has no control over what algorithms \mathcal{S} supplies, but regardless of that \mathcal{F} can only abort or return values such that $b = b'$. The parties executing Π use $g \leftarrow \text{Gen}(1^\lambda)$ obtained from $\mathcal{F}_{\text{CRS}}^{\text{Gen}}$. Since Gen returns random functions from F_λ and we are only guaranteed the existence of one f that fails commutativity for our u, x_1, x_2 , the probability

of Gen returning f has a lower bound of $\frac{1}{|F_\lambda|}$. The size of F_λ is polynomially constrained by 1^λ , so the probability cannot be negligible. Then Π returns $b = f(a, x_2) = f(f(u, x_1), x_2) \neq f(f(u, x_2), x_1) = f(a', x_1) = b'$. If the protocol returns $b = b'$ or aborts, \mathcal{Z} outputs 0, otherwise it outputs 1, succeeding at distinguishing between \mathcal{F} and Π with non-negligible probability, depending on if Π uses a non-commutative accumulation function.

Suppose AC does not satisfy witness-unforgeability. \mathcal{Z} runs an instance of $\mathcal{G}(f, \mathcal{U}_f, u)$ to get a forgery (x, w, X) with non-negligible probability. \mathcal{Z} activates P with $(\text{sid}, \text{ACCUM}, u, X)$ to get $(\text{sid}, \text{ACCUM}, a)$, and then with $(\text{sid}, \text{VERIFY}, w, \{x\}, a)$. In the execution with \mathcal{F} , there is no entry for $(a, \{x\}, w)$ in L_{acc} . If $\text{Verify}(w, \{x\}, a) = \perp$ for Verify supplied by \mathcal{S} , \mathcal{F} outputs \perp . Otherwise, it finds there is $(a, [X])$ in L_{sets} such that $x \notin X$ and aborts. Π computes $\text{Verify}(w, \{x\}, a) = \top$ since $a = f(u, X)$ and returns \top . Hence if the protocol returns \perp or aborts, \mathcal{Z} outputs 0, otherwise it outputs 1, succeeding in distinguishing with non-negligible probability depending only on whether \mathcal{G} supplies the forgery.

Lastly, if AC has algorithms Delete , Update such that their outputs do not always verify, then as in the previous two cases we use this to construct a distinguishing environment \mathcal{Z} . It activates parties to accumulate the inputs for which Delete or Update produce invalid outputs, then activates M to delete those items and observes if the protocol aborts (as \mathcal{F}_{acc} will if Verify outputs \perp) or returns the expected output. \square

A.2 Simulator from proof of Theorem 3

\mathcal{S} proceeds as follows:

1. Send $(\text{sid}, \text{CORRUPT}, \mathbf{V}_{\text{corr}})$ to $\mathcal{F}_{\text{elig}}$ where \mathbf{V}_{corr} is the set of voters that \mathcal{A} has decided to corrupt at the beginning of the real execution.
2. Wait to receive $(\text{sid}, \text{SETUP_ELIG})$ from $\mathcal{F}_{\text{elig}}$.
3. Simulate the ‘init’ stage of Π_{elig} :
 - Simulate setup of the hybrid functionalities \mathcal{F}_{acc} , \mathcal{F}_{NIC} , $\{\mathcal{F}_{\text{cert}}^{\text{SA}}\}$ and \mathcal{F}_{SoK} by requesting algorithms and parameters from \mathcal{A} for each of them. Store all except the accumulator Delete algorithm in St_{gen} .
 - Use the received algorithms to define the new algorithms as in Fig. 12 and send $(\text{sid}, \text{SETUP_ELIG}, \text{GenCred}, \text{AuthBallot}, \text{VrfyBallot}, \text{UpState}, St_{\text{gen}})$ to $\mathcal{F}_{\text{elig}}$.
 - Wait to receive $(\text{sid}, \text{ELIG_PAR}, \text{reg.par})$ from $\mathcal{F}_{\text{elig}}$.
4. Begin receiving $(\text{sid}, \text{GEN_CRED}, V_i, rc_i)$ from $\mathcal{F}_{\text{elig}}$ for honest voters V_i . Sign and verify the credential with the algorithms provided upon request from \mathcal{A} by playing the role of $\mathcal{F}_{\text{cert}}^{V_i}$. If verification succeeds, then collate their public credentials in the set C .
5. Send $(\text{sid}, \text{GEN_CRED}, V'_i, cr'_i, rc'_i)$ to $\mathcal{F}_{\text{elig}}$ for each corrupted voter V'_i , where rc'_i is the public credential generated by \mathcal{A} and cr'_i is the private credential generated by \mathcal{A} or a random unique value, if \mathcal{A} has generated the public part in another way.

6. \mathcal{S} has no visibility into AUTH_BALLOT queries between $\mathcal{F}_{\text{elig}}$ and the dummy honest voters.
7. Simulate the ‘cast’ stage of Π_{elig} :
 - Respond to \mathcal{A} ’s requests to hybrid functionalities involved in authentication.
 - Verification and linking of ballots is performed by the voters themselves, so \mathcal{S} can again only respond to \mathcal{A} ’s requests to hybrid functionalities. The answers should be consistent with those of $\mathcal{F}_{\text{elig}}$, since they are computed using the same algorithms.
8. \mathcal{S} can provide the authentication receipts on behalf of corrupted voters by sending $(\text{sid}, \text{AUTH_BALLOT}, V'_i, v'_i, \sigma'_i)$ to $\mathcal{F}_{\text{elig}}$, using the outputs given by \mathcal{A} in the simulation.

Hence the ideal and the real distribution will be the same.