Distributed, end-to-end verifiable, and privacy-preserving internet voting systems

Nikos Chondros [a,], Bingsheng Zhang [b], Thomas Zacharias [c],
Panos Diamantopoulos [a], Stathis Maneas [d], Christos Patsonakis [a],
Alex Delis [a], Aggelos Kiayias [c], Mema Roussopoulos [a]*

[a] *Department of Informatics and Telecommunications, University of Athens, Panepistiomiopolis, Ilisia, 157 84
Athens, Greece*
[b] *School of Computing and Communications, Lancaster University, InfoLab21, Bailrigg, Lancaster LA1 4WA, UK*
[c] *School of Informatics, University of Edinburgh, Office 5.16, 10 Crichton St., Edinburgh EH8 9AB, UK*
[d] *Department of Computer Science, University of Toronto, Bahen Centre for Information Technology, Room 5214, 40
St. George Street, Toronto, ON M5S2E4, Canada*

ABSTRACT

We present the D-DEMOS suite of distributed, privacy-preserving, and end-to-end verifiable e-voting systems; one completely asynchronous and one with minimal timing assumptions but better performance. Their distributed voting operation is human verifiable; a voter can vote over the web, using an unsafe web client stack, without sacrificing her privacy, and get recorded-as-cast assurance. Additionally, a voter can outsource election auditing to third parties, still without sacrificing privacy. We provide a model and security analysis of the systems, implement prototypes of the complete systems, measure their performance experimentally, demonstrate their ability to handle large-scale elections, and demonstrate the performance trade-offs between the two versions.

1. Introduction

Internet voting systems (e.g., Adida, 2008; Chaum, 2001; Chaum et al., 2008; Clarkson et al., 2008; Cramer et al., 1997; Gjøsteen, 2013; Joaquim et al., 2013; Kiayias et al., 2015; Kutylowski and Zagórski, 2010; Zagórski et al., 2013) are a powerful technology to improve the election process and thus provide a fundamental service of e-Government. They have the potential to enhance the democratic process by reducing election costs and by increasing voter participation for social groups that face considerable physical barriers and overseas voters. In addition, several internet voting systems (Adida, 2008; Kiayias et al., 2015; Kutylowski and Zagórski, 2010; Zagórski et al., 2013) allow voters and auditors to directly verify the integrity of the entire election process, providing *end-to-end verifiability*. This is a highly desired property that has emerged in the last decade, where voters can be assured that no entities, even the election authorities, have manipulated the election result. Despite their potential, existing internet voting systems suffer

from single points of failure, which may result in the compromise of voter secrecy, service availability, or integrity of the result (Adida, 2008; Benaloh et al., 2013; Chaum, 2001; Chaum et al., 2008; 2005; Clarkson et al., 2008; Cramer et al., 1997; Fisher et al., 2006; Gjøsteen, 2013; Joaquim et al., 2013; Kiayias et al., 2015; Kutylowski and Zagórski, 2010; Zagórski et al., 2013).

In this paper, we present the design and prototype implementation of the *D-DEMOS* suite of distributed, end-to-end verifiable internet voting systems, with no single point of failure during the election process (that is, besides setup). We set out to overcome two major limitations in existing internet voting systems. The first, is their dependency on centralized components, which is inherent in non-fault-tolerant systems. The second is the requirement of distributed voting systems, for the voter to run special software which processes cryptographic operations on their devices. Overcoming the latter allows votes to be cast with a greater variety of client devices, such as feature phones using SMS, or untrusted public web terminals. Our design is inspired by the novel approach proposed in Kiayias et al. (2015), where the voters are used as a source of randomness to challenge the zero-knowledge proof protocols (Feige et al., 1988). We use the latter to enable end-to-end verifiability.

We design a distributed *Vote Collection* (VC) subsystem that is able to collect votes from voters and assure them their vote was recorded as cast, without requiring any cryptographic operation from the client device. This allows voters to vote via SMS, a simple console client over a telnet session, or a public web terminal, while preserving their privacy. At election end time, *VC* nodes agree on a single set of votes. We introduce two versions of D-DEMOS that differ in how they achieve agreement on the set of cast votes. The D-DEMOS/Async version is completely asynchronous, while D-DEMOS/IC makes minimal synchrony assumptions but is more efficient than the alternative. Once agreement has been achieved, *VC* nodes upload the set of cast votes to a second distributed component, the *Bulletin Board* (BB). This is a replicated service that publishes its data immediately and makes it available to the public forever. Finally, our *trustees* subsystem, comprises a set of persons entrusted with secret keys which can unlock information stored in the *BB*. We share these secret keys among the *trustees*, making sure only an honest majority can uncover information from the *BB*. *Trustees* interact with the *BB* once the votes are uploaded to the latter, to produce and publish the final election tally.

The resulting voting systems are end-to-end verifiable, by the voters themselves and third-party auditors, while preserving voter privacy. To delegate auditing, a voter provides an auditor specific information from her ballot. The auditor, in turn, reads from the distributed *BB* and verifies the complete election process, including the correctness of the election setup by election authorities. Additionally, as the number of auditors increases, the probability of election fraud going undetected diminishes exponentially.

Finally, we implement prototypes of both D-DEMOS voting system versions. We measure their performance experimentally, under a variety of election settings, demonstrating their ability to handle thousands of concurrent connections, and thus manage large-scale elections. We also compare the two systems and emphasize the trade-offs between them, regarding security and performance.

To summarize, we make the following contributions:

- We present a suite of state-of-the-art, end-to-end verifiable, distributed voting systems with no single point of failure besides setup.
- Both systems allow voters to verify their vote was tallied-as-intended without the assistance of special software or trusted devices, and allow external auditors to verify the correctness of the election process. Additionally, both systems allow voters to delegate auditing to a third party auditor, without sacrificing their privacy.
- We provide a model and a security analysis of D-DEMOS/IC.
- We implement prototypes of the systems, measure their performance and demonstrate their ability to handle large-scale elections. Finally, we demonstrate the performance trade-offs between the two versions of the system.

Note that, a preliminary version of one of our systems was used to conduct exit-polls at three voting sites for two national-level elections and is being adopted for use by the largest civil union of workers in Greece, consisting of over a half million members.

The remainder of this paper is organized as follows. Section 2 introduces required background knowledge we reference throughout the paper, while Section 3 presents related work. Section 4 gives an overview of the system components, defines the system and threat model, and describes each system component in detail. Section 5.2 goes over some interesting attack vectors, which help to clarify our design choices. Section 6 describes our prototype implementations and their evaluation, and Section 7 concludes the main body of the paper. Finally, Appendix A provides, for the interested reader, the full proofs of liveness, safety, privacy and end-to-end verifiability of both our systems.

2.   Background

In this section we provide basic background knowledge required to comprehend the system description in the next section. This includes some voting systems terminology, a quick overview of Interactive Consistency, and a series of cryptographic tools we use to design our systems. These tools include additively homomorphic commitment schemes and zero-knowledge proofs, which are used in the System Description (Section 4), and are needed to understand the system design. Additionally, we provide details about collision resistant hash functions, IND-CPA symmetric encryption schemes, and digital signatures, which we use as building blocks for our security proofs in Appendix A.

2.1.   *Voting systems requirements*

An ideal electronic voting system would address a specific list of requirements (see Gritzalis, 2002; Internet Policy Institue, 2001; Neumann, 1993 for an extensive description). Our system addresses the following requirements:

- *End-to-end verifiability*: the voters can verify that their votes were counted as they intended and any party can verify that the election procedure was executed correctly.
- *Privacy*: a party that does not monitor voters during the voting phase of the election, cannot extract information about the voters' ballots. In addition, a voter cannot prove how she voted to any party that did not monitor her during the voting phase of the election[1].
- *Fault tolerance*: the voting system should be resilient to the faulty behavior of up to a number of components or parts, and be both live and safe.

2.2. Interactive consistency

Interactive consistency (IC), first introduced and studied by Pease et al. (1980), is the problem in which n nodes, where up to t may be byzantine, each with its own private value, run an algorithm that allows all non-faulty nodes to infer the values of each other. In our D-DEMOS/IC system, we use the *IC,BC-RBB* algorithm from Diamantopoulos et al. (2015), which achieves IC using a single synchronous round. This algorithm uses two phases to complete. The synchronous *Value Dissemination Phase* comes first, aiming to disperse the values across nodes. Afterwards, an asynchronous *Result Consensus Phase* starts, which results in each honest node holding a vector with every honest node's slot filled with the corresponding value.

2.3. Cryptographic tools

In this text, we use λ as the cryptographic security parameter and we write negl(λ) to denote that a function is negligible in λ, i.e., it is asymptotically smaller than the inverse of any polynomial in λ.

2.3.1. Additively homomorphic commitments
To achieve integrity against a malicious election authority, D-DEMOS utilizes lifted El Gamal (1985) over elliptic curves as a *non-interactive commitment scheme* that achieves the following properties:

1. *Perfectly binding*: no adversary can open a commitment Com(m) of m to a value other than m.
2. *Hiding*: there exists a constant $c < 1$ s.t. the probability that a commitment Com(m) to m leaks information about m to an adversary running in $O(2^{\lambda^c})$ steps is no more than negl(λ).
3. *Additively homomorphic*: $\forall m_1$, m_2, we have that Com(m_1) · Com(m_2) = Com($m_1 + m_2$).

2.3.2. Zero-knowledge proofs
D-DEMOS's security requires the election authority to show the correctness of the election setup to the public without compromising privacy. We enable this kind of verification with the use of zero-knowledge proofs. In a zero-knowledge proof, the prover is trying to convince the verifier that a statement is true, without revealing any information about the statement apart from the fact that it is true

(Quisquater et al., 1990). More specifically, we say an interactive proof system has the *honest-verifier zero-knowledge (HVZK)* property if there exists a probabilistic polynomial time simulator $\mathcal{S}$ that, for any given challenge, can output an accepting proof transcript that is distributed indistinguishable to the real transcript between an honest prover and an honest verifier. Here, we adopt Chaum–Pedersen zero-knowledge proofs (Chaum and Pedersen, 1993), which belong in the special class of Σ protocols (i.e., 3-move public-coin special HVZK proofs), allowing the Election Authority to show that the content inside each commitment is a valid option encoding.

2.3.3. Collision resistant hash functions
Given the security parameter $\lambda \in \mathbb{N}$, we say that a hash function $h: \{0, 1\}^* \mapsto \{0, 1\}^{\ell(\lambda)}$, where $\ell(\lambda)$ is polynomial in λ, is (t, ϵ)-*collision resistant* if for every adversary $\mathcal{A}$ running in time at most t, the probability of $\mathcal{A}$ finding two distinct preimages $m_1 \neq m_2$ such that $h(m_1) = h(m_2)$ is less than ϵ. By the birthday attack, for h to be (t, ϵ)-collision resistant, we necessitate that $t^2/2^{\ell(\lambda)} < \epsilon$. In this work, we use SHA-256 as the instantiation of a $(t, t^2 \cdot 2^{-256})$-collision resistant hash function.

2.3.4. IND-CPA symmetric encryption schemes
We say that a symmetric encryption scheme $\mathcal{SE}$ achieves (t, q, ϵ)-*indistinguishability against chosen plaintext attacks (IND-CPA)*, if for every adversary $\mathcal{A}$ that (i) runs in time at most t, (ii) makes at most q encryption queries that are pairs of messages $(m_{0,1}, m_{1,1}), \ldots, (m_{0,q}, m_{1,q})$ and (iii) for every encryption query $(m_{0,i}, m_{1,i})$, it receives the encryption of $m_{b,i}$, where b is the outcome of a coin-flip, it holds that:

$$\mathbf{Adv}_{\mathcal{SE}}^{\mathrm{IND-CPA}}(\mathcal{A})$$
$$:= \left| \Pr[\mathcal{A} \text{ outputs } 1 \mid b = 1] - \Pr[\mathcal{A} \text{ outputs } 1 \mid b = 0] \right| < \epsilon,$$

where by $\mathbf{Adv}_{\mathcal{SE}}^{\mathrm{IND-CPA}}(\mathcal{A})$ we denote the *advantage* of $\mathcal{A}$. D-DEMOS applies AES-128-CBC$ encryption, for which a known upper bound is

$$\mathbf{Adv}_{128-\mathrm{AES-CBC\$}}^{\mathrm{IND-CPA}}(\mathcal{A}) \leq 2 \cdot \mathbf{Adv}_{\mathrm{AES-128}}^{\mathrm{PRF}}(\mathcal{B}) + q^2 \cdot 2^{-128},$$

where $\mathbf{Adv}_{\mathrm{AES-128}}^{\mathrm{PRF}}(\mathcal{B})$ is the advantage of an algorithm $\mathcal{B}$ that runs in time at most $t + 129 \cdot q$ and makes at most q queries to break the pseudorandomness of the AES-128 block cipher. A safe conjecture is that $\mathbf{Adv}_{\mathrm{AES-128}}^{\mathrm{PRF}}(\mathcal{B}) \leq (t + 129 \cdot q + q^2) \cdot 2^{-128}$, so in our proofs we assume that AES-128-CBC$ is $(t, q, (2t + 258 \cdot q + 3q^2) \cdot 2^{-128})$-IND-CPA secure. For further details, we refer the reader to (Bellare and Rogaway, 2005, Chapters 3–5).

2.3.5. Digital signature schemes
A digital signature system is said to be secure if it is *existentially unforgeable under a chosen-message attack (EUF-CMA)*. Roughly speaking, this means that an adversary running in polynomial time and adaptively querying signatures for (polynomially many) messages has no more than negl(λ) probability to forge a valid signature for a new message. D-DEMOS/Async utilizes the standard RSA signature scheme, which is EUF-CMA secure under the factoring assumption.

---

[1] In Kiayias et al. (2015), this property is referred as *receipt-freeness*.

3. Related work

3.1. Voting systems

Several end-to-end verifiable e-voting systems have been introduced, e.g., the kiosk-based systems (Benaloh et al., 2013; Chaum et al., 2008; 2005; Fisher et al., 2006; Moran and Naor, 2010) and the internet voting systems (Adida, 2008; Joaquim et al., 2013; Kiayias et al., 2015; Kulyk et al., 2017; Kutylowski and Zagórski, 2010; Zagórski et al., 2013). In all these works, the Bulletin Board (BB) is a single point of failure and has to be trusted.

Dini presents a distributed e-voting system, which however is not end-to-end verifiable (Dini, 2003). In Culnane and Schneider (2014), there is a distributed BB implementation, also handling vote collection, according to the design of the vVote end-to-end verifiable e-voting system (Culnane et al., 2015), which in turn is an adaptation of the Prêt à Voter e-voting system (Chaum et al., 2005). In Culnane and Schneider (2014), the proper operation of the BB during ballot casting requires a trusted device for signature verification. In contrast, our vote collection subsystem is designed so that correct execution of ballot casting can be "human verifiable", i.e., by simply checking the validity of the obtained receipt. Additionally, our vote collection subsystem in D-DEMOS/Async is fully asynchronous, always deciding with exactly $n - f$ inputs, while in Culnane and Schneider (2014), the system uses a synchronous approach based on the FloodSet algorithm from Lynch (1996) to agree on a single version of the state. In this work, we consider secure ballot distribution as out of scope. However, this problem can be circumvented with specialized hardware, such as in Joaquim et al. (2013).

DEMOS (Kiayias et al., 2015) is an end-to-end verifiable e-voting system, which introduces the novel idea of extracting the challenge of the zero-knowledge proof protocols from the voters' random choices; we leverage this idea in our system too. However, DEMOS uses a centralized Election Authority (EA), which maintains all secrets throughout the entire election procedure, collects votes, produces the result and commits to verification data in the BB. Hence, the EA is a single point of failure, and because it knows the voters' votes, it is also a critical privacy vulnerability. In this work, we address these issues by introducing distributed components for vote collection and result tabulation, and we do not assume any trusted component during election. Additionally, DEMOS does not provide any recorded-as-cast feedback to the voter, whereas our system includes such a mechanism.

In Chondros et al. (2016), which is a preliminary version of this work, we present D-DEMOS/Async only. In this work, we also present the design of a new system, called D-DEMOS/IC, its implementation, evaluation and comparison to D-DEMOS/Async. We highlight the performance gains from the new IC approach, and also demonstrate its limitations. Thus, we present a suite of voting systems that provides implementers more options regarding strict safety or higher performance. Additionally, in Section 5.2, we present potential attacks and the ways our voting systems thwart them, to give the reader the intuition behind our design choices. As part of this work, we have completed and written full rigorous security proofs for both systems, which we include in Appendix A. In the conference version of this paper, we provide only a proof sketch for one system (D-DEMOS/Async).

3.2. State machine rplication

Castro and Liskov (2002) introduce a practical Byzantine Fault Tolerant replicated state machine protocol. In the last several years, a number of protocols for Byzantine Fault Tolerant state machine replication have been introduced to improve performance (Cowling et al., 2006; Kotla et al., 2007), robustness (Aublin et al., 2013; Clement et al., 2009b), or both (Aublin et al., 2015; Clement et al., 2009a). Our system does not use the state machine replication approach to handle vote collection, as it would be inevitably more costly. Each of our vote collection nodes can validate a voter's requests on its own. In addition, we are able to process multiple different voters' requests concurrently, without enforcing the total ordering inherent in replicated state machines. Finally, we do not wish voters to use special client-side software to access our system.

4. System description

4.1. Problem definition and goals

We consider an *election* with a single *question* and *m options*, for *n* voters, where voting takes place between a certain *begin* and *end* time (the *voting hours*), and each voter may select a single *option*.

Our major goals in designing our voting system are three:

(1) It has to be end-to-end verifiable, so that anyone can verify the complete election process. Additionally, voters should be able to outsource auditing to third parties, without revealing their voting choice.
(2) It has to be fault-tolerant, so that an attack on system availability and correctness is hard.
(3) Voters should not have to trust the terminals they use to vote, as such devices may be malicious, while still being assured their vote was recorded.

4.2. System overview

We employ an election setup component in our system, which we call the Election Authority (EA), to alleviate the voter from employing any cryptographic operations. The EA initializes all other system components, and then gets immediately destroyed to preserve privacy. The *Vote Collection* (VC) subsystem collects the votes from the voters during election hours, and assures them their vote was *recorded-as-cast*. Our *Bulletin Board* (BB) subsystem, which is a public repository of all election-related information, is used to hold all ballots, votes, and the result, either in encrypted or plain form, allowing any party to read from the BB and verify the complete election process. The VC subsystem uploads all votes to the BB at election end time. Finally, our design includes *trustees*, who are persons entrusted with managing all actions needed until result tabulation and publication, including all actions

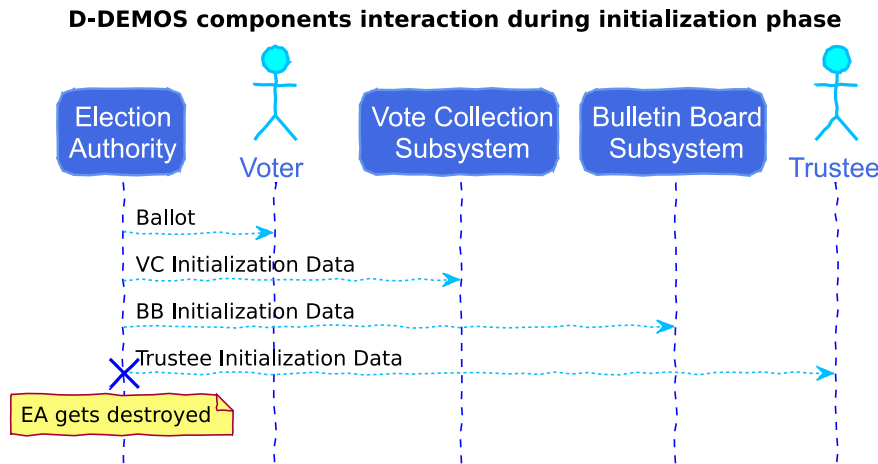**D-DEMOS components interaction during initialization phase**

Fig. 1 – High-level diagram of component interactions during system initialization. Each subsystem is a distributed system of its own, but is depicted as a unified entity in this diagram for brevity.

supporting end-to-end verifiability. *Trustees* hold the keys to uncover any information hidden in the *BB*, and we use threshold cryptography to make sure a malicious minority cannot uncover any secrets or corrupt the process.

Our system starts with the *EA* generating initialization data for every component of our system. The *EA* encodes each election option, and *commits* to it using a commitment scheme, as described below. It encodes the ith option as $\vec{e}_i$, a unit vector where the ith element is 1 and the remaining elements are 0. The commitment of an option encoding is a vector of (lifted) ElGamal ciphertexts (ElGamal, 1985) over elliptic curve that element-wise encrypts a unit vector. Note that this commitment scheme is also additively homomorphic, i.e., the commitment of $e_a + e_b$ can be computed by component-wise multiplying the corresponding commitments of e_a and e_b. The *EA* then creates a votecode and a receipt for each option. Subsequently, the *EA* prepares one ballot for each voter, with two functionally equivalent parts. Each part contains a list of options, along with their corresponding vote codes and receipts. We consider ballot distribution to be outside the scope of this paper, but we do assume ballots, after being produced by the *EA*, are distributed in a secure manner to each voter; thus only each voter knows the vote codes listed in her ballot. We make sure vote codes are not stored in clear form anywhere besides the voter's ballot. We depict this interaction in Fig. 1.

Our *VC* subsystem collects the votes from the voters during election hours, by accepting up to one vote code from each voter (see Fig. 2). The *EA* initializes each *VC* node with the vote codes and the receipts of the voters' ballots. However, it hides the vote codes, using a simple commitment scheme based on symmetric encryption of the plaintext along with a random salt value. This way, each *VC* node can verify if a vote code is indeed part of a specific ballot, but cannot recover any vote code until the voter actually chooses to disclose it. Additionally, we secret-share each receipt across all *VC* nodes using an $(N-f, N)$-VSS (verifiable secret-sharing) scheme with trusted dealer (Schneier, 1996), making sure that a receipt can be recovered and posted back to the voter only when a strong majority of *VC* nodes participates successfully in our voting protocol.

The voter selects one part of her ballot at random, and posts her selected vote code to one of the *VC* nodes. When she receives a receipt, she compares it with the one on her ballot corresponding to the selected vote code. If it matches, she is assured her vote was correctly recorded and will be included in the election tally. The other part of her ballot, the one not used for voting, will be used for auditing purposes. This design is essential for verifiability, in the sense that the *EA* cannot predict which part a voter may use, and the unused part will betray a malicious *EA* with $\frac{1}{2}$ probability per audited ballot.

Our second distributed subsystem is the *BB*, which is a replicated service of isolated nodes. Each *BB* node is initialized from the *EA* with vote codes and associated option encodings in committed form (again, for vote code secrecy), and each *BB* node provides public access to its stored information. At election end time, *VC* nodes run our Vote Set Consensus protocol, which guarantees all *VC* nodes agree on a single set of voted vote codes. After agreement, each *VC* node uploads this set to every *BB* node, which in turn publishes this set once it receives the same copy from enough *VC* nodes (see Fig. 3).

Our third distributed subsystem is a set of *trustees*, who are persons entrusted with managing all actions needed after vote collection, until result tabulation and publication; this includes all actions supporting end-to-end verifiability. Secrets that may uncover information in the *BB* are shared across *trustees*, making sure malicious *trustees* under a certain threshold cannot uncover and disclose sensitive information. We use Pedersen's Verifiable linear Secret Sharing (VSS) (Pedersen, 1991) to split the election data among the *trustees*. In a (k, n)-VSS, at least k shares are required to reconstruct the original data, and any collection of less than k shares leaks no information about the original data. Moreover, Pedersen's VSS is additively homomorphic, i.e., one can compute the share of $a + b$ by adding the share of a and the share of b respectively. This approach allows *trustees* to perform homomorphic "addition" on the option-encodings of cast vote codes, and contribute back a share of the opening of the homomorphic "total". Once enough *trustees* upload their shares of the "total", the election tally is uncovered and published at each *BB* node (see Fig. 4).

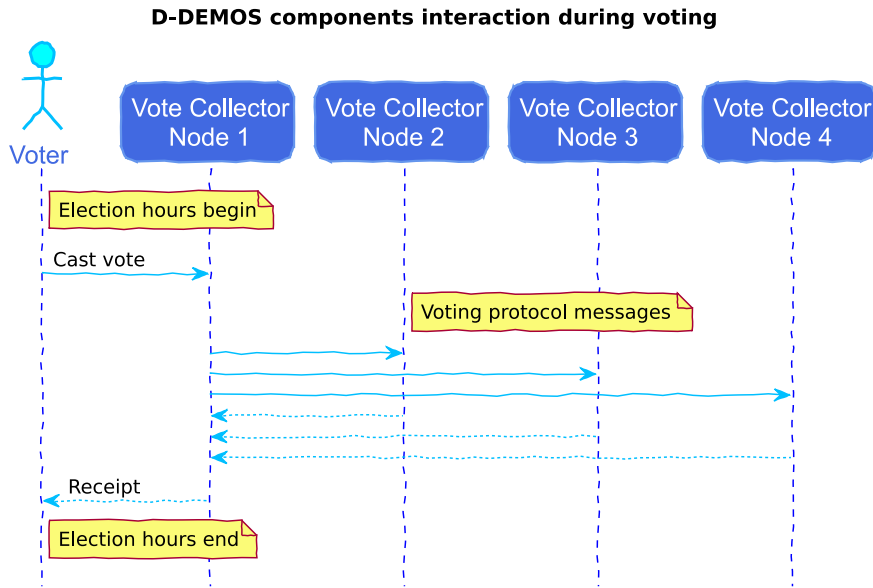**D-DEMOS components interaction during voting**

Fig. 2 – High-level diagram of component interactions during the voting phase. Message exchanges between VC nodes are simplified for this diagram. In this diagram, there are 4 VC nodes, tolerating up to 1 fault.

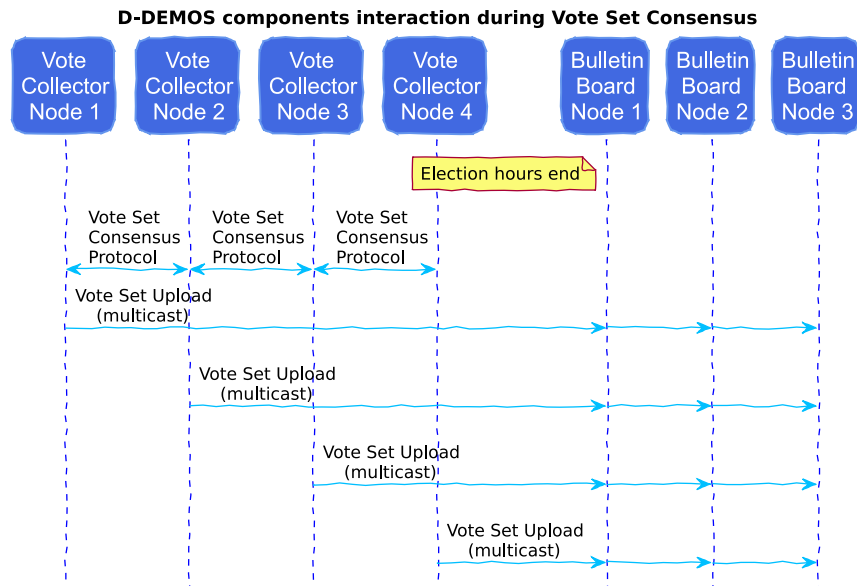**D-DEMOS components interaction during Vote Set Consensus**

Fig. 3 – High-level diagram of component interactions during the vote set consensus phase. 4 VC nodes and 3 BB nodes are shown, where each subsystem tolerates 1 fault. "VSC" stands for "Vote Set Consensus". After agreeing on a single Vote Set S, each VC node uploads S to every BB node. Messages are simplified for this diagram.

By the design of the VC, BB and trustee subsystems, and given that all fault tolerance thresholds are met, our system adheres to the following contract with the voters: *Any honest voter who receives a valid receipt from a VC node, is assured her vote will be published on the BB, and thus it will be included in the election tally.*

To ensure voter privacy, the system cannot reveal the content inside an option encoding commitment at any point. However, a malicious EA might put an arbitrary value (say 9000 votes for option 1) inside such a commitment, causing an incorrect tally result. To prevent this, we utilize the Chaum–

Pedersen zero-knowledge proof (Chaum and Pedersen, 1993), allowing the EA to show that the content inside each commitment is a valid option encoding, without revealing its actual content. Namely, the prover uses Sigma OR proof to show that each ElGamal ciphertext encrypts either 0 or 1, and the sum of all elements in a vector is 1. Our zero knowledge proof is organized as follows. First, the EA posts the initial part of the proofs on the BB. Second, during the election, each voter's A/B part choice is viewed as a source of randomness, 0/1, and all the voters' choices are collected and used as the challenge of our zero knowledge proof. Finally, the *trustees* will jointly

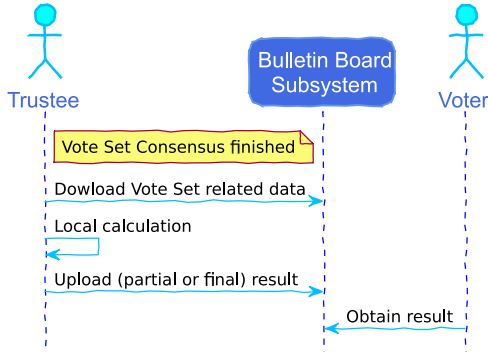**D-DEMOS components interaction towards result publication**

Fig. 4 – High-level diagram of *trustee* interactions with the *BB*, towards result tabulation and publication. *Trustees* are more than one, and interact with the *BB* in any order. The *BB* is a distributed system of its own, but is depicted as a unified entity in this diagram for brevity.

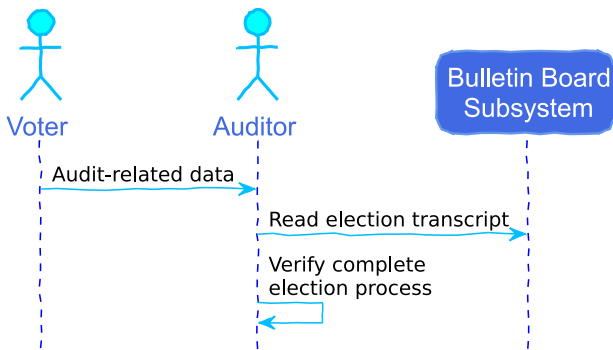**D-DEMOS components interaction for auditing**

Fig. 5 – High-level diagram of the system auditing. Voters send Auditors audit-related data that does not violate the voter's privacy. Auditors in turn read from the *BB* and verify the complete election process. The *BB* is a distributed system of its own, but is depicted as a unified entity in this diagram for brevity.

produce the final part of the proofs and post it on the *BB* before the opening of the tally. Hence, everyone can verify those proofs on the *BB*. We omit the zero-knowledge proof components in this paper and refer the interested reader to Chaum and Pedersen (1993) for details.

Our design allows any voter to read information from the *BB*, combine it with her private ballot, and verify her ballot was included in the tally. Additionally, any third-party auditor can read the *BB* and verify the complete election process (see Fig. 5). As the number of auditors increases, the probability of election fraud going undetected diminishes exponentially. For example, even if only 10 people audit, with each one having $\frac{1}{2}$ probability of detecting ballot fraud, the probability of ballot fraud going undetected is only $\frac{1}{2}^{10} = 0.00097$. Thus, even if the *EA* is malicious and, e.g., tries to point all vote codes to a specific option, this faulty setup will be detected because of the end-to-end verifiability of the complete system.

In this paper, we present two different versions of our voting system, with different performance and security trade-offs. In the first version, called *D-DEMOS/IC*, Vote Set Consensus is realized by an algorithm achieving Interactive Consistency, and thus requiring synchronization. The second version, *D-DEMOS/Async*, uses an asynchronous binary consensus algorithm for Vote Set Consensus, and thus is completely asynchronous. The performance trade-offs between the two are analyzed in Section 6.2.

4.3. System and threat model

We assume a fully connected network, where each node can reach any other node with which it needs to communicate. The network can drop, delay, duplicate, or deliver messages out of order. However, we assume messages are eventually delivered, provided the sender keeps retransmitting them. For all nodes, we make no assumptions regarding processor speeds.

We assume the EA sets up the election and is destroyed upon completion of the setup, as it does not directly interact with the remaining components of the system, thus reducing the attack surface of the privacy of the voting system as a whole. We also assume initialization data for every system component is relayed to it via untappable channels. We assume the adversary does not have the computational power to violate the security of any underlying cryptographic primitives. We place no bound on the number of faulty nodes the adversary can coordinate, as long as the number of malicious nodes of each subsystem is below its corresponding fault threshold. Let N_v, N_b, and N_t be the number of VC nodes, BB nodes, and trustees respectively. The voters are denoted by $V_\ell, \ell = 1, \ldots, n$.

For both versions of our system, we assume the clocks of VC nodes are synchronized with real world time; this is needed to prohibit voters from casting votes outside election hours. For the safety of *D-DEMOS/Async* version, we make no further timing assumptions. To ensure liveness, we assume the adversary cannot delay communication between honest nodes above a certain threshold.

For the *D-DEMOS/IC* version, we use the *IC,BC-RBB* algorithm achieving Interactive Consistency (IC) from Diamantopoulos et al. (2015), which requires a single synchronization point after the beginning of the algorithm. To accommodate this, we use the election-end time as the starting point of IC, and additionally assume the adversary cannot cause clock drifts between VC nodes also for safety, besides liveness. This is because lost messages in the first round of *IC,BC-RBB* are considered failures of the sending node.

Formally, we assume there exists a *global clock* variable Clock $\in \mathbb{N}$, and that every VC node, BB node and voter X is equipped with an *internal clock* variable Clock$[X] \in \mathbb{N}$. We define the following two events on the clocks:

(i) The event Init(X) : Clock$[X]$ ← Clock, that initializes a node X by synchronizing its internal clock with the global clock.

(ii) The event Inc(i) : $i \leftarrow i + 1$, that causes some clock i to advance by one time unit.

The adversarial setting.

(1) The EA initializes every VC node, BB node, trustee of the D-DEMOS system by running $\mathsf{Init}(\cdot)$ in all clocks for synchronization. Then, EA prepares the voters' ballots and all the VC nodes', BB nodes', and trustees' initialization data. Finally, it forwards the ballots for ballot distribution to the voters V_ℓ, $\ell = 1, \ldots, n$.

(2) $\mathcal{A}$ corrupts a fixed subset of VC nodes, a fixed subset of BB nodes, and a fixed subset of trustees. In addition, it defines a fixed subset of corrupt voters $\mathcal{V}_{\mathrm{corr}}$.

(3) When an honest node X wants to transmit a message $\mathbf{M}$ to an honest node Y, then it just sends $(X, \mathbf{M}, Y)$ to $\mathcal{A}$.

(4) $\mathcal{A}$ may invoke the events $\mathsf{Inc}(\mathsf{Clock})$ or $\mathsf{Inc}(\mathsf{Clock}[X])$, for any node X. Moreover, $\mathcal{A}$ may write on the incoming network tape of any honest component node of D-DEMOS.

(5) For every voter V_ℓ:

   (a) If $V_\ell \in \mathcal{V}_{\mathrm{corr}}$, then $\mathcal{A}$ fully controls V_ℓ.

   (b) If $V_\ell \notin \mathcal{V}_{\mathrm{corr}}$, then $\mathcal{A}$ may initialize V_ℓ by running $\mathsf{Init}(V_\ell)$ only once. If this happens, then the only control of $\mathcal{A}$ over V_ℓ is $\mathsf{Inc}(\mathsf{Clock}[V_\ell])$ invocations. Upon initialization, V_ℓ engages in the voting protocol.

Fig. 6 – The adversarial setting for the adversary $\mathcal{A}$ acting upon the distributed bulletin board system.

The adversarial setting for $\mathcal{A}$ upon D-DEMOS is defined in Fig. 6.

The description in Fig. 6 poses no restrictions on the control the adversary has over all internal clocks, or the number of nodes that it may corrupt (arbitrary denial of service attacks or full corruption of D-DEMOS nodes are possible). Therefore, it is necessary to strengthen the model so that we can perform a meaningful security analysis and prove the properties (liveness, safety, end-to-end verifiability, and voter privacy) that D-DEMOS achieves. Namely, we require the following:

A. Fault tolerance. We consider arbitrary (Byzantine) failures, because we expect our system to be deployed across separate administrative domains. For each of the subsystems, we have the following fault tolerance thresholds:
   - The number of faulty VC nodes, f_v, is strictly less than 1/3 of N_v i.e., for fixed f_v:

   $$N_v \geq 3f_v + 1.$$

   - The number of faulty BB nodes, f_b, is strictly less than 1/2 of N_b i.e., for fixed f_b:

   $$N_b \geq 2f_b + 1.$$

   - For the trustees' subsystem, we apply h_t out-of N_t threshold secret sharing, where h_t is the number of honest trustees, thus we tolerate $f_t = N_t - h_t$ malicious trustees.

B. Bounded synchronization loss. For the liveness of D-DEMOS (both versions), all system participants are aware of a value T_{end} such that for each node X, if $\mathsf{Clock}[X] \geq T_{\mathrm{end}}$, then X considers that the election has ended. In addition, the safety of D-DEMOS/IC version, assumes two timing points, a starting point (that we set as T_{end}) and a *barrier*,

denoted by T_{barrier}, that determine the beginning of the *Value Dissemination* phase and the transition to the *Result Consensus* phase of the underlying Interactive Consistency protocol (see Section 2.2), respectively.

For the above reasons, we bound the drift on the nodes' internal clocks, assuming an upper bound Δ of the drift of all honest nodes' internal clocks with respect to the global clock. Formally, we have that $\mathcal{A}$ may invoke the events $\mathsf{Inc}(\mathsf{Clock})$ or $\mathsf{Inc}(\mathsf{Clock}[X])$ for every node X, under the restriction that $|\mathsf{Clock}[X] - \mathsf{Clock}| \leq \Delta$, where $|\cdot|$ denotes the absolute value.

C. Bounded communication delay. For the liveness of D-DEMOS (both versions) and the safety of D-DEMOS/IC, we need to ensure eventual message delivery in bounded time. Therefore, we assume that there exists an upper bound δ on the time that $\mathcal{A}$ can delay the delivery of the messages between honest nodes. Formally, when the honest node X sends $(X, \mathbf{M}, Y)$ to $\mathcal{A}$, if the value of the global clock is T, then $\mathcal{A}$ must write $\mathbf{M}$ on the incoming network tape of Y by the time that $\mathsf{Clock} = T + \delta$. We note that δ should be a reasonably small value for liveness, while for safety of D-DEMOS/IC it suffices to be dominated by the predetermined timeouts of the VC nodes.

For clarity, we recap the aforementioned requirements in Fig. 7.

4.4. Election authority

EA produces the initialization data for each election entity in the setup phase. To enhance the system robustness, we let the EA generate all the public/private key pairs for all the system components (except voters) without relying on external PKI support. We use zero knowledge proofs to ensure the correctness of all the initialization data produced by the EA.

Requirement	D-DEMOS/IC		D-DEMOS/Async	
	Liveness	Safety	Liveness	Safety
Fault tolerance of the VC subsystem	✓	✓	✓	✓
Fault tolerance of the BB subsystem		✓		✓
Fault tolerance of the trustees' subsystem		✓		✓
Bounded synchronization loss	✓	✓	✓	
Bounded communication delay	✓	✓	✓	

Fig. 7 – Requirements for the liveness and safety of D-DEMOS/IC and D-DEMOS/Async.

4.4.1. Voter ballots

The EA generates one ballot ballot_ℓ for each voter ℓ, and assigns a unique 64-bit $\mathsf{serial-no}_\ell$ to it. As shown below, each ballot consists of two parts: Part A and Part B. Each part contains a list of m $\langle \mathsf{vote-code}, \mathsf{option}, \mathsf{receipt} \rangle$ tuples, one tuple for each election option. The EA generates the vote-code as a 128-bit random number, unique within the ballot, and the receipt as 64-bit random number.

$\mathsf{serial-no}_\ell$

Part A

$\mathsf{vote-code}_{\ell,1}$	$\mathsf{option}_{\ell,1}$	$\mathsf{receipt}_{\ell,1}$
...	...	...
$\mathsf{vote-code}_{\ell,m}$	$\mathsf{option}_{\ell,m}$	$\mathsf{receipt}_{\ell,m}$

Part B

$\mathsf{vote-code}_{\ell,1}$	$\mathsf{option}_{\ell,1}$	$\mathsf{receipt}_{\ell,1}$
...	...	...
$\mathsf{vote-code}_{\ell,m}$	$\mathsf{option}_{\ell,m}$	$\mathsf{receipt}_{\ell,m}$

4.4.2. BB initialization data

The initialization data for all BB nodes is identical, and each BB node publishes its initialization data immediately. The BB's data is used to show the correspondence between the vote codes and their associated cryptographic payload. This payload comprises the committed option encodings, and their respective zero knowledge proofs of valid encoding (first move of the prover), as described in Section 4.2. However, the vote codes must be kept secret during the election, to prevent the adversary from "stealing" the voters' ballots and using the stolen vote codes to vote. To achieve this, the EA first randomly picks a 128-bit key, msk, and encrypts each $\mathsf{vote-code}$ using AES-128-CBC with random initialization vector (AES-128-CBC\$) encryption, denoted as $[\mathsf{vote-code}]_{\mathsf{msk}}$. Each BB node is given $H_{\mathsf{msk}} \leftarrow SHA256(\mathsf{msk}, \mathsf{salt}_{\mathsf{msk}})$ and $\mathsf{salt}_{\mathsf{msk}}$, where $\mathsf{salt}_{\mathsf{msk}}$ is a fresh 64-bit random salt. Hence, each BB node can be assured the key it reconstructs from VC key-shares (see below) is indeed the key that was used to encrypt these vote-codes.

The rest of the BB initialization data is as follows: for each $\mathsf{serial-no}_\ell$, and for each ballot part, there is a *shuffled* list of $\langle [\mathsf{vote-code}_{\ell,\pi_\ell^X(j)}]_{\mathsf{msk}}, \mathsf{payload}_{\ell,\pi_\ell^X(j)} \rangle$ tuples, where $\pi_\ell^X \in S_m$ is a random permutation (X is A or B).

$(H_{\mathsf{msk}}, \mathsf{salt}_{\mathsf{msk}})$

$\mathsf{serial-no}_\ell$

Part A

$[\mathsf{vote-code}_{\ell,\pi_\ell^A(1)}]_{\mathsf{msk}}$	$\mathsf{payload}_{\ell,\pi_\ell^A(1)}$
⋮	⋮
$[\mathsf{vote-code}_{\ell,\pi_\ell^A(m)}]_{\mathsf{msk}}$	$\mathsf{payload}_{\ell,\pi_\ell^A(m)}$

Part B

$[\mathsf{vote-code}_{\ell,\pi_\ell^B(1)}]_{\mathsf{msk}}$	$\mathsf{payload}_{\ell,\pi_\ell^B(1)}$
⋮	⋮
$[\mathsf{vote-code}_{\ell,\pi_\ell^B(m)}]_{\mathsf{msk}}$	$\mathsf{payload}_{\ell,\pi_\ell^B(m)}$

We shuffle the list of tuples of each part to ensure voter's privacy. This way, nobody can guess the voter's choice from the position of the cast vote-code in this list.

4.4.3. VC initialization data

The EA uses an $(N_v - f_v, N_v)$-VSS (Verifiable Secret-Sharing) scheme to split msk and every $\mathsf{receipt}_{\ell,j}$ into N_v shares, denoted as $(\|\mathsf{msk}\|_1, \ldots, \|\mathsf{msk}\|_{N_v})$ and $(\|\mathsf{receipt}_{\ell,j}\|_1, \ldots, \|\mathsf{receipt}_{\ell,j}\|_{N_v})$ respectively. For each $\mathsf{vote-code}_{\ell,j}$ in each ballot, the EA also computes $H_{\ell,j} \leftarrow SHA256(\mathsf{vote-code}_{\ell,j}, \mathsf{salt}_{\ell,j})$, where $\mathsf{salt}_{\ell,j}$ is a 64-bit random number. $H_{\ell j}$ allows each VC node to validate a $\mathsf{vote-code}_{\ell,j}$ individually (without network communication), while still keeping the $\mathsf{vote-code}_{\ell,j}$ secret. To preserve voter privacy, these tuples are also shuffled using π_ℓ^X. The initialization data for VC_i is structured as below:

$\|\mathsf{msk}\|_i$

$\mathsf{serial-no}_\ell$

Part A

$(H_{\ell,\pi_\ell^A(1)}, \mathsf{salt}_{\ell,\pi_\ell^A(1)})$	$\|\mathsf{receipt}_{\ell,\pi_\ell^A(1)}\|_i$
...	...
$(H_{\ell,\pi_\ell^A(m)}, \mathsf{salt}_{\ell,\pi_\ell^A(m)})$	$\|\mathsf{receipt}_{\ell,\pi_\ell^A(m)}\|_i$

Part B

$(H_{\ell,\pi_\ell^B(1)}, \mathsf{salt}_{\ell,\pi_\ell^B(1)})$	$\|\mathsf{receipt}_{\ell,\pi_\ell^B(1)}\|_i$
...	...
$(H_{\ell,\pi_\ell^B(m)}, \mathsf{salt}_{\ell,\pi_\ell^B(m)})$	$\|\mathsf{receipt}_{\ell,\pi_\ell^B(m)}\|_i$

4.4.4. Trustee initialization data

The EA uses (h_t, N_t)-VSS to split the opening of encoded option commitments $Com(\bar{e}_i)$ into N_t shares, denoted as

$(\|\vec{e}_i^i\|_1, \ldots, \|\vec{e}_i^i\|_{N_t})$. The initialization data for Trustee$_i$ is structured as below:

serial−no$_\ell$	
Part A	
$Com(\vec{e}_{\pi_\ell^A(i)})$	$\|\vec{e}_{\pi_\ell^A(i)}\|_\ell$
...	...
Part B	
$Com(\vec{e}_{\pi_\ell^B(i)})$	$\|\vec{e}_{\pi_\ell^B(i)}\|_\ell$
...	...

Similarly, the state of zero knowledge proofs for ballot correctness is shared among the *trustees* using (h_t, N_t)-VSS. For further details, we refer the interested reader to Chaum and Pedersen (1993).

4.5. Vote Collectors

The Vote Collection subsystem comprises N_v nodes that collect the votes from the voters and, at election end time, agree on a single set of cast vote codes and upload it to the Bulletin Board. In the following subsections, we present two different versions of the *VC* subsystem, one with a timing assumption (*D-DEMOS/IC*) and one fully asynchronous (*D-DEMOS/Async*).

4.5.1. Vote Collectors for D-DEMOS/IC

VC is a distributed system of N_v nodes, running our *voting* and *vote-set consensus* protocols. *VC* nodes have private and authenticated channels to each other, and a public (unsecured) channel for voters. The algorithms implementing our *D-DEMOS/IC voting* protocol are presented in Algorithm 1. For simplicity, we present our algorithms operating for a single election.

The *voting* protocol starts when a voter submits a VOTE⟨serial−no, vote−code⟩ message to a *VC* node. We call this node the *responder*, as it is responsible for delivering the receipt to the voter. The *VC* node confirms the current system time is within the defined election hours, and locates the ballot with the specified serial−no. It verifies this ballot has not been used for this election, either with the same or a different vote code. Then, it compares the vote−code against every hashed vote code in each ballot line, until it locates the correct entry and obtains the corresponding receipt−share. Next, it marks the ballot as pending for the specific vote−code. Finally, it multicasts a VOTE_P⟨serial−no, vote−code, receipt−share⟩ message to all *VC* nodes, disclosing its share of the receipt. If the located ballot is marked as voted for the specific vote−code, the *VC* node sends the stored receipt to the voter without any further interaction with other *VC* nodes.

Each *VC* node that receives a VOTE_P message, first validates the received receipt−share according to the verifiable secret sharing scheme used. Then, it performs the same validations as the responder, and multicasts another VOTE_P message (only once), disclosing its share of the receipt. When a node collects $h_v = N_v − f_v$ valid shares, it uses the verifiable secret sharing reconstruction algorithm to reconstruct the receipt (the secret) and marks the ballot as voted for the specific vote−code. Additionally, the *responder* node sends this receipt back to the voter. A message flow diagram of our *voting* protocol is depicted in Fig. 8. As is evident from the diagram, the time from the multicast of the first VOTE_P message until

---

Algorithm 1 Vote Collector algorithms for D-DEMOS/IC.

```
1:  procedure ON VOTE(serial−no, vote−code) from source:
2:      if SysTime() between start and end
3:          b :=locateBallot(serial−no)
4:          if b.status == NotVoted
5:              l := ballot.VerifyVoteCode(vote−code)
6:              if l ≠ null
7:                  b.status := Pending
8:                  b.used−vc := vote−code
9:                  b.lrs := {}              ▷ list of receipt shares
10:                 sendAll(VOTE_P⟨serial−no, vote−code, l.share⟩)
11:                 wait for (N_v − f_v) VOTE_P messages, fill b.lrs
12:                 b.receipt := Rec(b.lrs)
13:                 b.status := Voted
14:                 send(source, b.receipt)
15:          else if b.status == Voted AND b.used−vc == vote−code
16:              send (source, ballot.receipt)

17: procedure ON VOTE_P(serial−no, vote−code, share) from source:
18:     if SysTime() between start and end
19:         b :=locateBallot(serial−no)
20:         if b.status == NotVoted
21:             l := ballot.VerifyVoteCode(vote−code)
22:             if l ≠ null
23:                 b.status := Pending
24:                 b.used−vc := vote−code
25:                 b.lrs.Append(share)
26:                 sendAll(VOTE_P⟨serial−no, vote−code, l.share⟩)
27:         else if b.status == Voted AND b.used−vc == vote−code
28:             b.lrs.Append(share)
29:             if size(b.lrs) >= N_v − f_v
30:                 b.receipt := Rec(b.lrs)
31:                 b.status := Voted

32: function BALLOT::VERIFYVOTECODE(vote−code)
33:     for l = 1 to ballot_lines do
34:         if lines[l].hash == h(vote−code||lines[l].salt) return l
        return null
```
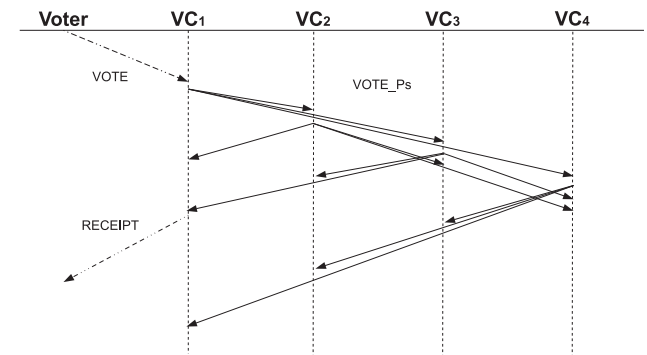

Fig. 8 – Diagram of message exchanges for a single vote during the D-DEMOS/IC vote collection phase.

---

collecting all receipt shares, is only slightly longer than a single round-trip between two *VC* nodes.

At election end time, each *VC* node stops processing VOTE and VOTE_P messages, and initiates the *vote-set consensus* protocol. It creates a set VS_i of ⟨serial−no, vote−code⟩ tuples, including all *voted* and *pending* ballots. Then, it participates in the Interactive Consistency (IC) protocol of Diamantopoulos et al. (2015), with this set. At the end of IC, each node contains

Cross-tabulate $\langle VS_1, \ldots, VS_n \rangle$ per ballot, creating a list of vote codes for each ballot. Perform the following actions for each ballot:

1) If the list contains two or more distinct vote codes, mark the ballot as NotVoted and exit.

2) If a vote code vc_a appears at least $N_v - 2f_v$ times, mark the ballot as Voted for vc_a and exit.

3) Otherwise, mark the ballot as NotVoted and exit.

Fig. 9 – High level description of algorithm after IC.

a vector $\langle VS_1, \ldots, VS_n \rangle$ with the Vote Set of each node, and follows the algorithm of Fig. 9. Step 1 makes sure any ballot with multiple submitted vote codes is discarded. Since vote codes are private, and cannot be guessed by malicious vote collectors, the only way for multiple vote codes to appear is if malicious voters are involved, against whom our system is not obliged to respect our *contract*.

With a single vote code remaining, step 2 considers the threshold above which to consider a ballot as voted for a specific vote code. We select the $N_v - 2f_v$ threshold for which we are certain that even the following extreme scenario is handled. If the *responder* is malicious, submits a receipt to an honest voter, but denies it during *vote-set consensus*, the remaining $N_v - 2f_v$ honest *VC* nodes that revealed their receipt shares for the generation of the receipt, are enough for the system to accept the vote code (receipt generation requires $N_v - f_v$ nodes, of which f_v may be malicious, thus $N_v - 2f_v$ are necessarily honest).

Finally, step 3 makes sure vote codes that occur less than $N_v - 2f_v$ times are discarded. Under this threshold, there is no way a receipt was ever generated.

At the end of this algorithm, each node submits the resulting set of *voted* $\langle$serial−no, vote−code$\rangle$ tuples to each *BB* node, which concludes its operation for the specific election.

4.5.2. Vote Collectors for D-DEMOS/Async

We make the following enhancements to the Vote Collection subsystem, to achieve the completely asynchronous version *D-DEMOS/Async*. During voting we introduce another step, which guarantees only a single vote code can be accepted (towards producing a receipt) for a given ballot, using a *uniqueness certificate* (see below). We also employ an asynchronous binary consensus primitive to achieve Vote Set Consensus.

More specifically, during voting, the *responder VC* node validates the submitted vote code, but before disclosing its receipt share, it multicasts an ENDORSE$\langle$serial−no, vote−code$\rangle$ message to all *VC* nodes. Each *VC* node, after making sure it has not endorsed another vote code for this ballot, responds with an ENDORSEMENT$\langle$serial−no, vote−code, sig$_{VC_i}\rangle$ message, where sig$_{VC_i}$ is a digital signature of the specific serial-no and vote-code, with VC_i's private key. The responder collects $N_v - f_v$ valid signatures, and places them in a *uniqueness certificate* UCERT for this ballot. It then discloses its receipt share via the VOTE_P message, but also attaches the formed UCERT in the message.

Each VC node that receives a VOTE_P message, first verifies the validity of UCERT and discards the message on error. On

success, it proceeds as per the *D-DEMOS/IC* protocol (validating the receipt share it receives and then disclosing its own receipt share).

The algorithms implementing our *D-DEMOS/Async voting* protocol are presented in Algorithm 2.

Algorithm 2 Vote Collector algorithms for D-DEMOS/Async.

```
 1: procedure ON VOTE(serial−no, vote−code) from source:
 2:     if SysTime() between start and end
 3:         b :=locateBallot(serial−no)
 4:         if b.status == NotVoted
 5:             l := ballot.VerifyVoteCode(vote−code)
 6:             if l ≠ null
 7:                 b.UCERT := {}              ▷ Uniqueness certificate
 8:                 sendAll(ENDORSE⟨serial−no, vote−code⟩)
 9:                 wait for (N_v − f_v) valid replies, fill b.UCERT
10:                 b.status := Pending
11:                 b.used−vc := vote−code
12:                 b.lrs := {}                ▷ list of receipt shares
13:                 sendAll(VOTE_P⟨serial−no, vote−code, l.share⟩)
14:                 wait for (N_v − f_v) VOTE_P messages, fill b.lrs
15:                 b.receipt := Rec(b.lrs)
16:                 b.status := Voted
17:                 send(source, b.receipt)
18:         else if b.status == Voted AND b.used−vc == vote−code
19:             send (source, ballot.receipt)

20: procedure ON VOTE_P(serial−no, vote−code, share, UCERT) from
        source:
21:     if UCERT is not valid
22:         return
23:     if SysTime() between start and end
24:         b :=locateBallot(serial−no)
25:         if b.status == NotVoted
26:             l := ballot.VerifyVoteCode(vote−code)
27:             if l ≠ null
28:                 b.status := Pending
29:                 b.used−vc := vote−code
30:                 b.lrs.Append(share)
31:                 sendAll(VOTE_P⟨serial−no, vote−code, l.share⟩)
32:         else if b.status == Voted AND b.used−vc == vote−code
33:             b.lrs.Append(share)
34:             if size(b.lrs) >= N_v − f_v
35:                 b.receipt := Rec(b.lrs)
36:                 b.status := Voted

37: function BALLOT::VERIFYVOTECODE(vote−code)
38:     for l = 1 to ballot_lines do
39:         if lines[l].hash == h(vote−code||lines[l].salt) return l
        return null
```
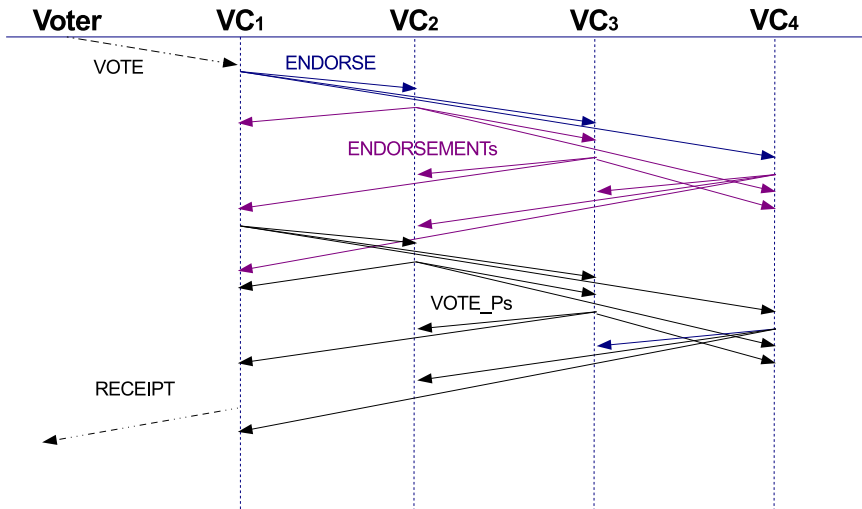
Fig. 10 – Diagram of message exchanges for a single vote during the D-DEMOS/Async vote collection phase.

1) Send ANNOUNCE⟨serial-no, vote-code, UCERT⟩ to all nodes. The vote-code will be *null* if the node knows of no vote code for this ballot.

2) Wait for $N_v - f_v$ such messages. If any such message contains a valid vote code vc_a, accompanied by a valid UCERT, change the local state immediately, by setting vc_a as the vote code used for this ballot.

3) Participate in a Binary Consensus protocol, with the subject "Is there a valid vote code for this ballot?". Enter with an opinion of 1, if a valid vote code is locally known, or a 0 otherwise.

4) If the result of Binary Consensus is 0, consider the ballot not voted.

5) Else, if the result of Binary Consensus is 1, consider the ballot voted. There are two sub-cases here:

    a) If vote code vc_a, accompanied by a valid UCERT is locally known, mark the ballot voted for vc_a.

    b) If, however, vc_a is not known, send a RECOVER-REQUEST⟨serial-no⟩ message to all *VC* nodes, wait for the first valid RECOVER-RESPONSE⟨serial-no, vc_a, UCERT⟩ response, and update the local state accordingly.

Fig. 11 – High level description of algorithm for asynchronous vote set consensus.

The voting process is outlined in the diagram of Fig. 10, where we now see two round-trips are needed before the receipt is reconstructed and posted to the voter. The formation of a valid UCERT gives our algorithms the following guarantees:

(a) No matter how many responders and vote codes are active at the same time for the same ballot, if a UCERT is formed for vote code vc_a, no other uniqueness certificate for any vote code different than vc_a can be formed.

(b) By verifying the UCERT before disclosing a VC node's receipt share, we guarantee the voter's receipt cannot be reconstructed unless a valid UCERT is present.

At election end time, each VC node stops processing ENDORSE, ENDORSEMENT, VOTE and VOTE_P messages, and

follows the *vote-set consensus* algorithm in Fig. 11, for each registered ballot.

Steps 1 and 2 ensure used vote codes are dispersed across nodes. Recall our receipt generation requires $N_v - f_v$ shares to be revealed by distinct *VC* nodes, of which at least $N_v - 2f_v$ are honest. Note that any two $N_v - f_v$ subsets of N_v have at least one honest node in common. Because of this, if a receipt was generated, at least one honest node's ANNOUNCE will be processed by every honest node, and all honest *VC* nodes will obtain the corresponding vote code in these two steps. Consequently, all honest nodes enter step 3 with an opinion of 1 and binary consensus is guaranteed to deliver 1 as the resulting value, thus safeguarding our contract against the voters. In any case, step 3 guarantees all *VC* nodes arrive at the same conclusion, on whether this ballot is voted or not.

In the algorithm outlined above, the result from binary consensus is translated from 0/1 to a status of "not-voted" or

a unique valid vote code, in steps 4 and 5. The 5b case of this translation, in particular, requires additional explanation. Assume, for example, that a voter submitted a valid vote code vc_a, but a receipt was not generated before election end time. In this case, an honest vote collector node VC_i may not be aware of vc_a at step 3, as steps 1 and 2 do not make any guarantees in this case. Thus, VC_i may rightfully enter consensus with a value of 0. However, when honest nodes' opinions are mixed, the consensus algorithm may produce any result. In case the result is 1, VC_i will not possess the correct vote code vc_a, and thus will not be able to properly translate the result. This is what our recovery sub-protocol is designed for. VC_i will issue a RECOVER-REQUEST multicast, and we claim that another honest node, VC_h exists that *possesses* vc_a and *replies* with it. The reason for the existence of an honest VC_h is straightforward and stems from the properties of the binary consensus problem definition. If all honest nodes enter binary consensus with the same opinion a, the result of any consensus algorithm is guaranteed to be a. Since we have an honest node VC_i, that entered consensus with a value of 0, but a result of 1 was produced, there has to exist another honest node VC_h that entered consensus with an opinion of 1. Since VC_h is honest, it must *possess* vc_a, along with the corresponding UCERT (as no other vote code vc_b can be active at the same time for this ballot). Again, because VC_h is honest, it will follow the protocol and *reply* with a well formed RECOVER-REPLY. Additionally, the existence of UCERT guarantees that any malicious replies can be safely identified and discarded.

As per *D-DEMOS/IC*, at the end of this algorithm, each node submits the resulting set of *voted* ⟨serial−no, vote−code⟩ tuples to each BB node, which concludes its operation for the specific election.

4.6. Voter

We expect the voter, who has received a ballot from *EA*, to know the URLs of at least $f_v + 1$ VC nodes. To vote, she picks one part of the ballot at random, selects the vote code representing her chosen option, and loops, selecting a VC node at random and posting the vote code, until she receives a valid receipt. After the election, the voter can verify two things from the updated BB. First, she can verify her cast vote code is included in the tally set. Second, she can verify that the unused part of her ballot, as "opened" at the BB, matches the copy she received before the election started. This step verifies that the vote codes are associated with the expected options as printed in the ballot. Finally, the voter can delegate both of these checks to an *auditor*, without sacrificing her privacy. This is because the cast vote code does not reveal her choice, and because the unused part of the ballot is completely unrelated to the used one.

4.7. Bulletin Board

A BB node functions as a public repository of election-specific information. It can be read via a public and anonymous channel, while writes happen over an authenticated channel, implemented with PKI originating from the voting system. BB nodes are independent from each other, as a BB node never directly contacts another BB node. Readers are expected to issue a read request to all BB nodes, and trust the reply that comes from the majority. Writers are also expected to write to all BB nodes; their submissions are always verified, and explained in more detail below.

After the setup phase, each BB node publishes its initialization data. During election hours, BB nodes remain inert. After the voting phase, each BB node receives from each VC node, the final vote-code set and the shares of msk. Once it receives $f_v + 1$ identical final vote code sets, it accepts and publishes the final vote code set. Once it receives $N_v − f_v$ valid key shares (again from VC nodes), it reconstructs msk, and decrypts and publishes all the encrypted vote codes in its initialization data.

At this point, the cryptographic payloads corresponding to the cast vote codes are made available to the *trustees*. *Trustees*, in turn, read from the BB subsystem, perform their individual calculations and then write to the BB nodes; these writes are verified by the *trustees'* keys, generated by the *EA*. Once enough *trustees* have posted valid data, the BB node combines them and publishes the final election result.

We intentionally designed our BB nodes to be as simple as possible for the reader, refraining from using a *Replicated State Machine*, which would require readers to run algorithm-specific software. The robustness of BB nodes comes from controlling all write accesses to them. Writes from VC nodes are verified against their honest majority threshold. Further writes are allowed only from *trustees*, verified by their keys.

Finally, a reader of our BB nodes should post her read request to all nodes, and accept what the majority responds with ($f_b + 1$ is enough). We acknowledge there might be temporary state divergence (among BB nodes), from the time a writer updates the first BB node, until the same writer updates the last BB node. However, given our thresholds, this should be only momentary, alleviated with simple retries. Thus, if there is no reply backed by a clear majority, the reader should retry until there is one.

4.8. Trustees

After the end of election hours, each *trustee* fetches all the election data from the BB subsystem and verifies its validity. For each ballot, there are two possible valid outcomes: (i) one of the A/B parts are voted, (ii) none of the A/B parts are voted. If both A/B parts of a ballot are marked as voted, then the ballot is considered as invalid and is discarded. Similarly, *trustees* also discard those ballots where more than one commitments in an A/B part are marked as voted.

In case (i), for each encoded option commitment in the unused part, Trustee$_\ell$ submits its corresponding share of the opening of the commitment to the BB. For each encoded option commitment in the voted part, Trustee$_\ell$ computes and posts the share of the final message of the corresponding zero knowledge proof, showing the validity of those commitments. Meanwhile, those commitments marked as voted are collected to a tally set $\mathbf{E}_{tally}$. In case (ii), for each encoded option commitment in both parts, Trustee$_\ell$ submits its corresponding share of the opening of the commitment to the BB. Finally,

denote $D_{tally}^{(\ell)}$ as Trustee$_\ell$'s set of shares of option encoding commitment openings, corresponding to the commitments in E_{tally}. Trustee$_\ell$ computes the opening share for E_{sum} as $T_\ell = \sum_{D \in D_{tally}^{(\ell)}}$ and then submits T_ℓ to each *BB* node.

4.9.   *Auditors*

Auditors are participants of our system who can verify the election process. The role of the auditor can be assumed by voters or any other party. After election end time, auditors read information from the *BB* and verify the correct execution of the election, by verifying the following:

1. within each opened ballot, no two vote codes are the same;
2. there are no two submitted vote codes associated with any single ballot part;
3. within each ballot, no more than one part has been used;
4. all the openings of the commitments are valid;
5. all the zero-knowledge proofs associated with the used ballot parts are completed and valid.

In case they received audit information (an unused ballot part and a cast vote code) from voters who wish to delegate verification, they can also verify:

6. the submitted vote codes are consistent with the ones received from the voters;
7. the openings of the unused ballot parts are consistent with the ones received from the voters.

5.   Discussion

5.1.   *Potential attacks*

In this section, we outline some of the possible attacks against the D-DEMOS systems, and the way our systems thwart them. This is a high level discussion, aiming to help the reader understand *why* our systems work reliably. In Appendix A, we provide the formal proofs of correctness and privacy, which are the foundation of this discussion.

In this high-level description, we intentionally do not focus on Denial-of-Service attacks, as these kind of attacks attempt to stop the system from producing a result, or stop voters from casting their votes. Although these attacks are important (Huang et al., 2007), they cannot be hidden, as voters will notice immediately the system not responding (either because of our receipt mechanism and our liveness property, or because of lack of information in the *BB*). Instead, we focus on attacks on the correctness of the election result, as these have consequences typical voters cannot identify easily. In this discussion, we assume the fault thresholds of Section 4.3 are not violated, and the attacker cannot violate the security of the underlying cryptographic primitives.

Additionally, D-DEMOS is immune to phishing attacks (Kumaraguru et al., 2010; Yue and Wang, 2010), as the voter does not disclose credentials that can be used to vote *any* option, but a vote code that corresponds to a *specific* option.

In this section, we focus on correctness, noting that our systems' privacy is achieved by the security of our cryptographic schemes (see Sections 2.3 and A.4 for details), and the partial initialization data that each node of the distributed subsystems receives at the setup phase.

5.1.1.   *Malicious Election Authority component*
At a high level, the EA produces vote codes and corresponding receipts. Vote codes are pointers to the associated cryptographic payload, which includes *option encodings*. Options encodings are used to produce the tally using homomorphic addition. If the EA miss-encodes any option, it will be identified by the Zero-Knowledge proof validation performed by the Auditors.

The EA may instead try to "point" a vote code to a valid but different option encoding (than the one described in the voter's ballot), in an attempt to manipulate the result. In this case, the EA cannot predict which one of the two parts the voter will use. Recall that the unused part of the ballot will be opened in the *BB* by the *trustees*, and thus the voters can read and verify the correctness of their unused ballot parts.

As explained in detail in Appendix A.3, if none of the above attacks take place, there is perfect consistency between each voter's ballot and its corresponding information on the *BB*. Because of this, as well as the correctness and the perfect hiding property of our commitment scheme, the homomorphic tally will be opened to the actual election result.

5.1.2.   *Malicious voter*
A malicious voter can try to submit multiple vote codes to the VC subsystem, attempting to cause disagreement between its nodes. In this case, a receipt *may* be generated, depending on the order of delivery of network messages. Note that, our safety *contract* allows our system to either accept only one vote code for this ballot, or discard the ballot altogether, as the voter is malicious and our contract holds only against honest voters.

In the D-DEMOS/IC case, this is resolved at the *Vote Set Consensus* phase. During the *voting* phase, each VC node accepts only the first vote code it receives (via either a VOTE or a VOTE_P message), and attempts to follow our *voting* protocol. This results in the generation of at most one receipt, for one of the posted vote codes. However, during *Vote Set Consensus*, honest VC nodes will typically identify the multiple posted vote codes and discard the ballot altogether, even if a receipt was indeed generated. If the ballot is not discarded (e.g., because malicious vote collector nodes hid the extra vote codes and honest nodes knew only of one), our $N_v - 2f_v$ threshold guarantees that no vote codes with generated receipts are discarded.

In the D-DEMOS/Async case, this is resolved completely at the *voting* phase. Each VC node still accepts only the first vote code it receives, but additionally attempts to build a UCERT for it. As the generation of a UCERT is guaranteed to be successful only for a single vote code, the outcome of the *voting* protocol will be either no UCERT being built, resulting in considering the ballot as not-voted, or a single UCERT generated.

Thus, the two systems behave differently in the case of multiple posted vote codes, as D-DEMOS/IC typically discards

such ballots, while D-DEMOS/Async may process some of them, when a UCERT is successfully built.

5.1.3. *Malicious Vote Collector*

A malicious *VC* node cannot easily guess the vote codes in the voters' ballots, as they are randomly generated. Additionally, because vote codes are encrypted in the local state of each *VC* node, the latter cannot decode and use them. Note that, a vote code in a voter's ballot is considered private until the voter decides to use it and transmits it over the network. From this point on, the vote code can be intercepted by the attacker, as the only power it gives him is to cast it.

A malicious *VC* node can obtain vote codes from colluding malicious voters. In this case, the only possible attack on correctness is exactly the same as if it originated from the malicious voter herself, and we already described our counter-measures in Section 5.1.2.

A malicious *VC* node may become a *responder*. In this case, this *VC* node may *selectively* transmit the cast vote code to a subset of the remaining *VC* nodes, potentially including all the other malicious and colluding nodes, and deliver the receipt to an honest voter. Consequently, the attacker controlling the malicious entities, may try to "confuse" the honest *VC* nodes and have them disagree on whether the ballot is voted or not, by having all malicious *VC* nodes lie at *vote set consensus* time, reporting the ballot as not voted.

Recall that, for the receipt to be generated, $N_v - f_v$ *VC* nodes need to cooperate, of which up to f_v may be malicious. This leaves $N_v - 2f_v$ honest nodes always present.

In the case of D-DEMOS/IC, these $N_v - 2f_v$ honest nodes will show up in the per ballot cross-tabulation, and will drive the decision to mark the ballot as voted (note that, in the algorithm of Fig. 9, $N_v - 2f_v$ is the lower threshold for a ballot to be marked as voted). In the case of D-DEMOS/Async, we include the ANNOUNCE-exchanging phase before the consensus algorithm, to guarantee at least one of the $N_v - 2f_v$ honest nodes' ANNOUNCE message will be processed by every honest node. In this case, all honest nodes will agree on entering consensus that the ballot is voted, which guarantees the outcome of consensus to be in accordance.

5.1.4. *Malicious* BB *nodes and* trustees

Malicious entities between both the *BB* nodes and the *trustees* cannot influence the security of both systems. The reason is, a node of each of these two subsystems does not communicate with the remaining nodes of the same subsystem, and thus cannot influence either the correctness, or progress of the system as a whole.

5.2. *System limitations*

We believe the D-DEMOS suite of e-voting systems are an important step towards achieving robust e-voting systems. Still, they have a few limitations which we list in this section.

First, we chose not to address secure ballot distribution in this work. The *EA*, which is common to both versions of the D-DEMOS suite, outputs the ballots and we assume they are distributed out-of-band, in a privacy-preserving manner to all voters. This issue can be circumvented with specialized hardware, such as in Joaquim et al. (2013), or by simply printing the ballots and having the voters pick them at random. In any case, we leave this issue as future work.

Second, in this work, we target *1-out-of-m* elections, in which voters can choose only one out of *m* options from their ballots. Our system could be extended to support a *k-out-of-m* type of voting scheme, where voters can select more than one option, each equally weighted during tallying. The homomorphic tally we employ, has the advantages of lightweight computation and communication complexity, but cannot support more complex voting schemes such as STV. To support more complex voting schemes, we could construct a mix-net based variant of our system. However, this is out of scope of this paper and we leave it as future work.

Additionally, we use 128-bit sized vote-codes, which would produce human readable codes of 22 characters given Base64 encoding. This is obviously a sizable character string for humans to input and verify. However, the 128-bit size we chose is simply a security parameter, as it affects the ability of the attacker to guess vote codes (see Section A.2). Thus, smaller, or even larger sized vote-codes can be used depending on the security requirements of the specific application. For details on the trade-off between vote-code size and security, we refer the interested reader to Section A.2.3.

Finally, our systems share the same administration burdens as all fault-tolerant systems. Deployment requires carefully select the correct number of nodes for each subsystem, ahead of time. This decision affects robustness, performance, and ease of administration.

6. Implementation and evaluation

6.1. *Implementation*

Voting system: We implement the Election Authority component of our system as a standalone C++ application, and all other components in Java. Whenever we store data structures on disk, or transmit them over the network, we use Google Protocol Buffers (Inc., 2015) to encode and decode them efficiently. We use the MIRACL library (MIRACL, 2015) for elliptic-curve cryptographic operations. In all applications requiring a database, we use the PostgreSQL relational database system (community, 2015b).

We build an *asynchronous communications stack* (ACS) on top of Java, using Netty (Community, 2015a) and the asynchronous PostgreSQL driver from Laisi (2015), using TLS based authenticated channels for inter-node communication, and a public HTTP channel for public access. This infrastructure uses connection-oriented sockets, but allows the applications running on the upper layers to operate in a message-oriented fashion. We use this infrastructure to implement *VC* and *BB* nodes. We implement "verifiable secret sharing with honest dealer", by utilizing Shamir's Secret Share library implementation (Tiemens, 2015), and having the *EA* sign each share.

For D-DEMOS/IC, we use the implementation of *IC,BC-RBB* (Interactive Consistency algorithm, using asynchronous binary consensus and reliable broadcast without signatures) from Diamantopoulos et al. (2015). We use the election end time as a synchronization point to start the algorithm, and

configure the timeout of the first phase of the algorithm according to the number of *VC* nodes and the number of ballots in the election. For D-DEMOS/Async, we implement Bracha's Binary Consensus directly on top of the ACS, and we use that to implement our Vote Set Consensus algorithm (depicted in Fig. 11). We introduce a version of Binary Consensus that operates in batches of arbitrary size; this way, we achieve greater network efficiency.

Additionally, we batch most of the asynchronous vote set consensus "announce" phase's messages. If this phase was implemented without optimization, it would result in a message complexity of $n*N_v$ (individual ANNOUNCE messages), imposing a significant network load. This is because each node has to multicast an ANNOUNCE message for each ballot, and wait for $n(N_v - f_v)$ replies to progress. To optimize it, we have each node consult its local database and diagnose cases where another node already knows the correct vote code and UCERT for a specific ballot. This is feasible because when a node VC_b discloses its share using the VOTE_P message, it also includes the UCERT, and this fact is recorded in the recipient's node (VC_a) database along with the sender node's share. For these cases, we produce ANNOUNCE_RANGE messages addressed to individual nodes, having the source node VC_a announce a range of ballot serial numbers as voted, a fact that is already known to the recipient node VC_b (because VC_a located recorded VOTE_P messages from VC_b). We use the same mechanism to announce ranges of not-voted ballots.

Trustee android application: In addition to the web interface for *trustees*, we also implement a specialized *Trustee* Android application. We re-use the MIRACL library on Android and provide a simple user interface for *trustees*, where they use a single button press to perform each of their required tasks: download their initialization data from the *EA*, download election data from the *BB*, calculate their cryptographic contribution to the result opening, and finally upload their share of the opening to the *BB*.

Web browser replicated service reader: Our choice to model the Bulletin Board as a replicated service of non-cooperating nodes puts the burden of response verification on the reader of the service; a human reader is expected to manually issue a read request to all nodes, then compare the responses and pick the one posted by the majority of nodes. To alleviate this burden, we implement a web browser extension which automates this task, as an extension for Mozilla Firefox. The user sets up the list of URLs for the replicated service. The add-on (1) intercepts any HTTP request towards any of these URLs, (2) issues the same request to the rest of the nodes, and (3) captures the responses, compares them in binary form, and routes the response coming from the majority, as a response to the original request posted by the user. Majority is defined by the number of defined URL prefixes; for 3 such URLs, the first 2 equal replies suffice.

With the above approach, the user never sees a wrong reply, as it is filtered out by the extension. Also note this process will be repeated for all dependencies of the initial web page (images, scripts, CSS), as long as they come from the same source (with the same URL prefix), verifying the complete user visual experience in the browser.

Note that, this mechanism is required only when reading data from the Bulletin Board, such as the election result, or audit information. This mechanism is neither needed nor used during voting, where the voter interacts with the Vote Collection subsystem using our voting protocol.

6.2. Evaluation

We experimentally evaluate the performance of our voting system, focusing mostly on our vote collection algorithm, which is the most performance critical part. We conduct our experiments using a cluster of 12 machines, connected over a Gigabit Ethernet switch. The first 4 are equipped with Hexa-core Intel Xeon E5-2420 @ 1.90 GHz, 16 GB RAM, and one 1TB SATA disk, running CentOS 7 Linux, and we use them to run our VC nodes. The remaining 8 comprise dual Intel(R) Xeon(TM) CPUs @ 2.80 GHz, with 4 GB of main memory, and two 50GB disks, running CentOS 6 Linux, and we use them as clients.

We implement a multi-threaded voting client to simulate concurrency. This client starts the requested number of threads, each of which loads its corresponding ballots from disk and waits for a signal to start. From then on, the thread enters a loop where it picks one VC node and vote code at random, requests the voting page from the selected VC (HTTP GET), submits its vote (HTTP POST), and waits for the reply (receipt). This simulates multiple concurrent voters casting their votes in parallel, and gives an understanding of the behavior of the system under the corresponding load. We employ the PostgreSQL RDBMS (community, 2015b) to store all VC initialization data from the *EA*.

We start off by demonstrating our system's capability of handling large-scale elections. To this end, we generate election data for referendums, i.e., $m = 2$, and vary the total number of ballots n from 50 million to 250 million. This causes the database size to increase accordingly and impact queries. We fix the number of concurrent clients to 400 and cast a total of 200,000 ballots, which are enough for our system to reach its steady-state operation (larger experiments result in the same throughput). Fig. 12 shows the throughput of both D-DEMOS/IC and D-DEMOS/Async declines slowly, even with a five-fold increase in the number of eligible voters. The cause of the decline is the increase of the database size. Note that, an "operation" in this experiment is the casting of a vote to a *VC*-node, including obtaining the generated receipt. This holds for all subsequent experiments where we report throughput in operations per second.

In our second experiment, we explore the effect of m, i.e., the number of election options, on system performance. We vary the number of options from $m = 2$ to $m = 10$. Each election has a total of $n = 200,000$ ballots which we spread evenly across 400 concurrent clients. As illustrated in Fig. 13, our vote collection protocol manages to deliver approximately the same throughput regardless of the value of m, for both D-DEMOS/IC and D-DEMOS/Async. Notice that the major extra overhead m induces during vote collection, is the increase in the number of hash verifications during vote code validation, as there are more vote codes per ballot. The increase in number of options has a minor impact on the database size as well (as each ballot has $2m$ options).

Next, we evaluate the scalability of our vote collection protocol by varying the number of vote collectors and concurrent

D-DEMOS/IC throughput versus n, LAN   D-DEMOS/Async throughput versus n, LAN

Fig. 12 – Vote collection throughput graphs for D-DEMOS/IC (a) and D-DEMOS/Async(b), versus the number of total election ballots n.

D-DEMOS/IC throughput versus m, LAN   D-DEMOS/Async throughput versus m, LAN
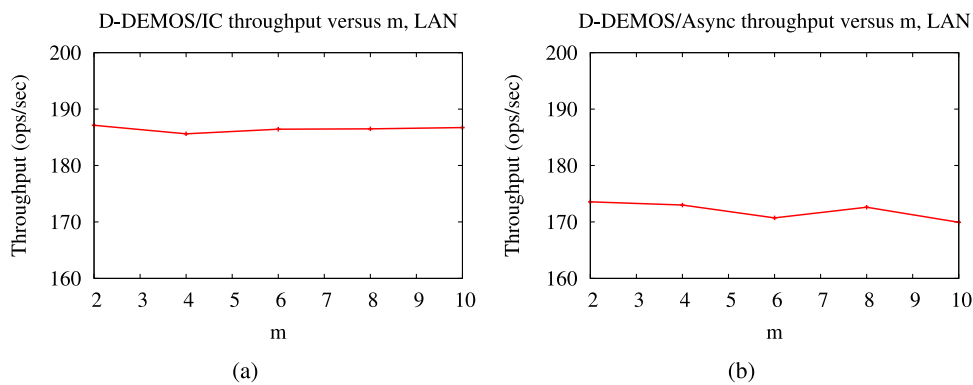
Fig. 13 – Vote collection throughput graphs for D-DEMOS/IC (a) and D-DEMOS/Async(b), versus the number of election options m.

clients. We eliminate the database, by caching the election data in memory and servicing voters from the cache, to measure the net communication and processing costs of our voting protocol. We vary the number of VC nodes from 4 to 16, and distribute them across the 4 physical machines. Note that, co-located nodes are unable to produce vote receipts via local messages only, since the $N_v - f_v$ threshold cannot be satisfied, i.e., cross-machine communication is still the dominant factor in receipt generation. For election data, we use the dataset with $n = 200,000$ ballots and $m = 4$ options, which is enough for our system to reach its steady state.

In Fig. 14, we plot the average response time of both our vote collection protocols, versus the number of vote collectors, under different concurrency levels, ranging from 500 to 2000 concurrent clients. Results for both systems illustrate an almost linear increase in the client-perceived latency, for all concurrency scenarios, up to 13 VC nodes. From this point on, when four logical VC nodes are placed on a single physical machine, we notice a non-linear increase in latency. We attribute this to the overloading of the memory bus, a resource shared among all processors of the system, which services all (in-memory) database operations. D-DEMOS/IC has a smaller response time with its single round intra-VC node communication, while D-DEMOS/Async is slightly slower due to the extra Uniqueness Certificate producing round.

Fig. 15 shows the throughput of both our vote collection protocols, versus the number of vote collectors, under differ-

ent concurrency levels. We observe that, in terms of overall system throughput, the penalty of tolerating extra failures (increasing the number of vote collectors) manifests early on. We notice an almost 50% decline in system throughput from 4 to 7 VC nodes for D-DEMOS/IC, and a bigger one for D-DEMOS/Async. However, further increases in the number of vote collectors lead to a much smoother, linear decrease. Overall, D-DEMOS/IC achieves better throughput than D-DEMOS/Async, due to exchanging fewer messages and lacking signature operations.

In Fig. 16, we plot a different view of both our systems' throughput, this time versus the concurrency level (ranging from 100 to 2000). Plots represent number of VC node settings (4–16), thus different fault tolerance levels. Results show both our systems have the nice property of delivering nearly constant throughput, regardless of the incoming request load, for a given number of VC nodes.

We repeat the same experiment by emulating a WAN environment using *netem* (Hemminger et al., 2005), a network emulator for Linux. We inject a uniform latency of 25 ms (typical for US coast-to-coast communication (Grigorik, 2013)) for each network packet exchanged between vote collector nodes, and present our results in Figs. 17–19. A simple comparison between LAN and WAN plots illustrates our system manages to deliver the same level of throughput and average response time, regardless of the increased intra-VC communication latency.
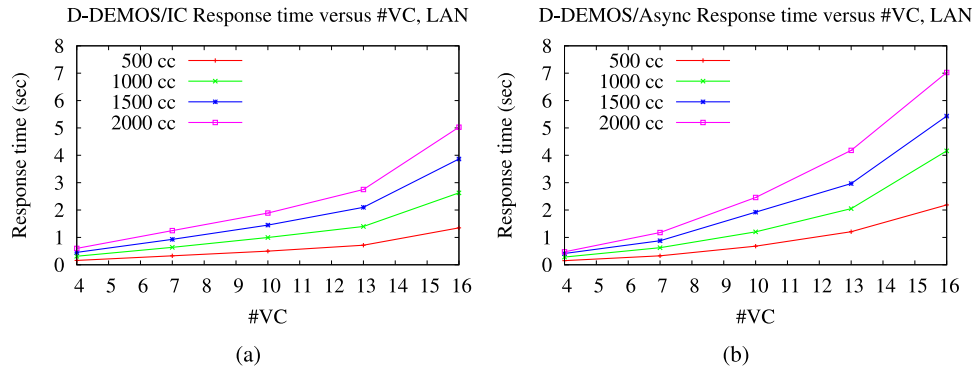
Fig. 14 – Vote Collection response time of D-DEMOS/IC (a) and D-DEMOS/Async (b), versus the number of VC nodes, under a LAN setting. Election parameters are $n = 200{,}000$ and $m = 4$.


Fig. 15 – Vote Collection throughput of D-DEMOS/IC (a) and D-DEMOS/Async (b), versus the number of VC nodes, under a LAN setting. Election parameters are $n = 200{,}000$ and $m = 4$.
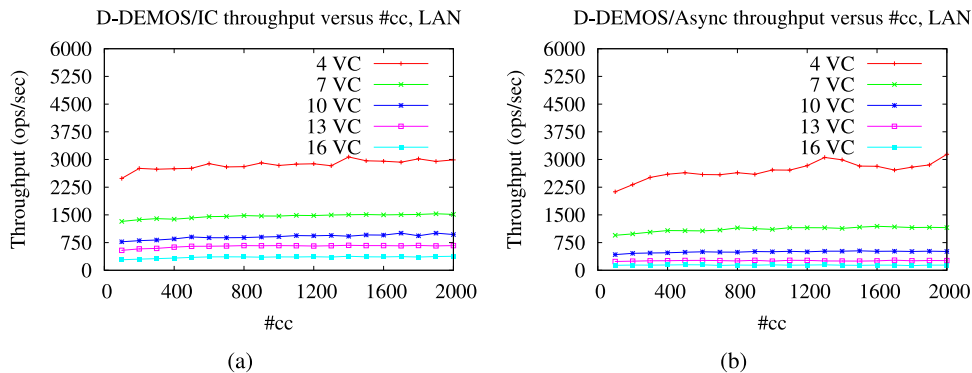

Fig. 16 – Vote Collection throughput of D-DEMOS/IC (a) and D-DEMOS/Async (b), versus the number of concurrent clients, under a LAN setting. Plots illustrate performance for different cardinalities of VC nodes, thus different fault tolerance settings. Election parameters are $n = 200{,}000$ and $m = 4$.

The benefits of the in memory approach, expressed both in terms of sub-second client (voter) response time and increased system throughput, make it an attractive alternative to the more standard database setup. For instance, in cases where high-end server machines are available, it would be possible to service mid to large scale elections completely from memory. We estimate the size of the in-memory representation of a $n = 200\,K$ ballot election, with $m = 4$ options, at approximately 322 MB (see Maneas, 2015 for derivation

details). In this size, we include 64-bit Java pointers overhead, as we are using simple hash-maps of plain Java classes. This size can be decreased considerably in a more elaborate implementation, where data is serialized by Google Protocol Buffers, for example.

Finally, in Fig. 20, we illustrate a breakdown of the duration of each phase of the complete voting system (D-DEMOS/IC and D-DEMOS/Async), versus the total number of ballots cast. We assume immediate phase succession, i.e., the vote collec-

Fig. 17 – Vote Collection response time of D-DEMOS/IC (a) and D-DEMOS/Async (b), versus the number of VC nodes, under a WAN setting. Election parameters are $n = 200,000$ and $m = 4$.
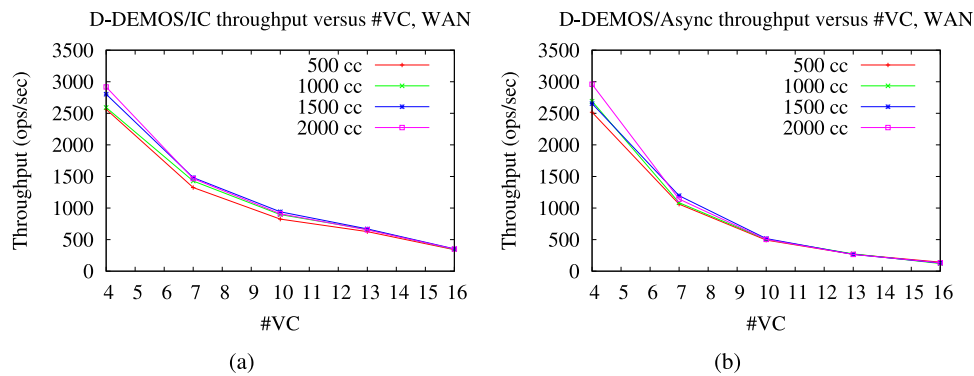

Fig. 18 – Vote Collection throughput of D-DEMOS/IC (a) and D-DEMOS/Async (b), versus the number of VC nodes, under a WAN setting. Election parameters are $n = 200,000$ and $m = 4$.


Fig. 19 – Vote Collection throughput of D-DEMOS/IC (a) and D-DEMOS/Async (b), versus the number of concurrent clients, under a WAN setting. Plots illustrate performance for different cardinalities of VC nodes, thus different fault tolerance settings. Election parameters are $n = 200,000$ and $m = 4$.

tion phase ends when all votes have been cast, at which point the vote set consensus phase starts, and so on. The "Push to BB and encrypted tally" phase is the time it takes for the vote collectors to push the final vote code set to the BB nodes, including all actions necessary by the BB to calculate and publish the encrypted result. The "Publish result" phase is the time it takes for *trustees* to calculate and push their share of the opening of the final tally to the BB, and for the BB to publish the final tally. Note that, in most voting procedures, the vote

collection phase would in reality last several hours and even days as stipulated by national law (see Estonia voting system). Thus, looking only at the post-election phases of the system, we see that the time it takes to publish the final tally on the BB is quite fast. Comparing the two versions of D-DEMOS, we observe D-DEMOS/IC is faster during both Vote Collection and Vote Set Consensus phases. This is expected, because of the extra communication round of D-DEMOS/Async during voting, as well as the more complex consensus-per-ballot
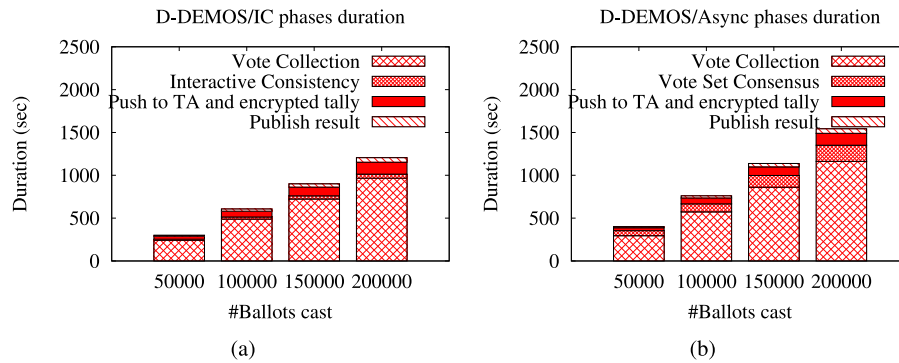
Fig. 20 – This figure illustrates the duration of all system phases. Results depicted are for 4 VCs, *n* = 200,000 and *m* = 4. All phases are disk based.

approach to achieving Vote Set Consensus. However, D-DEMOS/Async is more robust than D-DEMOS/IC, as it does not require any kind of synchronization between nodes.

Overall, although we introduced Byzantine Fault Tolerance across all phases of a voting system (besides setup), we demonstrate it achieves high performance, enough to run real-life elections of large electorate bodies.

7. Conclusion and future work

We have presented a suite of state-of-the-art, end-to-end verifiable, distributed internet voting systems with no single point of failure besides setup. Both systems allow voters to verify their vote was tallied-as-intended without the assistance of special software or trusted devices, and external auditors to verify the correctness of the election process. Additionally, the systems allows voters to delegate auditing to a third party auditor, without sacrificing their privacy. We have provided a model and security analysis of both voting systems. Finally, we have implemented prototypes of the integrated systems, measured their performance, and demonstrated their ability to handle large-scale elections.

We have used our system to conduct exit polls at three large voting sites for two national-level elections. We look forward to gaining more experience and feedback about our systems by exploring their use in election and decision-making procedures at all levels throughout the Greek university system, and studying their adoption for use by the General Confederation of Greek Workers, the largest civil union of workers in Greece. Finally, our systems currently support only *1-out-of-m* elections, in which voters choose one out of *m* options from their ballots. As future work, we will expand our systems to support *k-out-of-m* elections.

We hope our work contributes towards improving the election process, while also reducing the cost of performing elections in modern democracies.

Declarations of interest

None.

Acknowledgements

Appendix A. Security of D-Demos

In this section, we present at length the security properties that D-DEMOS achieves. Specifically, we show that D-DEMOS/IC and D-DEMOS/Async achieve liveness and safety, according to which every voter that submits her vote prior to a well-defined time threshold, will obtain a valid receipt (liveness) and her vote will be included in the election tally and published in the BB (safety contract). In addition, both versions achieve end-to-end verifiability and voter privacy at the same level as Kiayias et al. (2015)[2], thus allowing a top-tier integrity guarantee without compromising secrecy.

We use *m, n* to denote the number of options and voters respectively. We denote by λ the cryptographic security parameter and we write negl(λ) to denote that a function is negligible in λ, i.e., it is asymptotically smaller than the inverse of any polynomial in λ.

The remaining sections reference heavily the Cryptographic Tools Section 2.3, which includes the notions and claims about the security of the cryptographic tools we use in the two versions of D-DEMOS.

A.1. Liveness

To prove the liveness that D-DEMOS guarantees, we assume (i) an upper bound δ on the delay of the delivery of messages and (ii) an upper bound Δ on the drift of all clocks (see Assumptions B and C in Section 4.3). Furthermore, to express liveness rigorously, we formalize the behavior of honest voters regarding maximum waiting before vote resubmission as follows:

---

[2] In Kiayias et al. (2015), the authors use the term *voter privacy/receipt-freeness*, but they actually refer to the same property.

Step	Time upper bounds at each clock			
	Clock	Clock[V]	Clock[VC]	honest VC nodes' clocks
V is initialized	T	T	$T + \Delta$	$T + \Delta$
V submits her vote to VC	$T + T_{\text{comp}} + \Delta$	$T + T_{\text{comp}}$	$T + T_{\text{comp}} + 2\Delta$	$T + T_{\text{comp}} + 2\Delta$
VC receives V's ballot	$T + T_{\text{comp}} + \Delta + \delta$	$T + T_{\text{comp}} + 2\Delta + \delta$	$T + T_{\text{comp}} + 2\Delta + \delta$	$T + T_{\text{comp}} + 2\Delta + \delta$
VC verifies the validity of V's ballot and broadcasts its share	$T + 2T_{\text{comp}} + 3\Delta + \delta$	$T + 2T_{\text{comp}} + 4\Delta + \delta$	$T + 2T_{\text{comp}} + 2\Delta + \delta$	$T + 2T_{\text{comp}} + 4\Delta + \delta$
All the other honest VC nodes receive VC's share	$T + 2T_{\text{comp}} + 3\Delta + 2\delta$	$T + 2T_{\text{comp}} + 4\Delta + 2\delta$	$T + 2T_{\text{comp}} + 4\Delta + 2\delta$	$T + 2T_{\text{comp}} + 4\Delta + 2\delta$
All the other honest VC nodes verify the validity of V's share and broadcast their shares	$T + 3T_{\text{comp}} + 5\Delta + 2\delta$	$T + 3T_{\text{comp}} + 6\Delta + 2\delta$	$T + 3T_{\text{comp}} + 6\Delta + 2\delta$	$T + 3T_{\text{comp}} + 4\Delta + 2\delta$
VC receives all the $N_v - 1$ other honest VC nodes' shares	$T + 3T_{\text{comp}} + 5\Delta + 3\delta$	$T + 3T_{\text{comp}} + 6\Delta + 3\delta$	$T + 3T_{\text{comp}} + 6\Delta + 3\delta$	$T + 3T_{\text{comp}} + 6\Delta + 3\delta$
VC verifies the validity of all the $N_v - 1$ other honest VC nodes' shares	$T + (N_v + 2)T_{\text{comp}} + 7\Delta + 3\delta$	$T + (N_v + 2)T_{\text{comp}} + 8\Delta + 3\delta$	$T + (N_v + 2)T_{\text{comp}} + 6\Delta + 3\delta$	$T + (N_v + 2)T_{\text{comp}} + 8\Delta + 3\delta$
VC reconstructs and sends V's receipt	$T + (N_v + 3)T_{\text{comp}} + 7\Delta + 3\delta$	$T + (N_v + 3)T_{\text{comp}} + 8\Delta + 3\delta$	$T + (N_v + 3)T_{\text{comp}} + 6\Delta + 3\delta$	$T + (N_v + 3)T_{\text{comp}} + 8\Delta + 3\delta$
V obtains her receipt	$T + (N_v + 3)T_{\text{comp}} + 7\Delta + 4\delta$	$T + (N_v + 3)T_{\text{comp}} + 8\Delta + 4\delta$	$T + (N_v + 3)T_{\text{comp}} + 8\Delta + 4\delta$	$T + (N_v + 3)T_{\text{comp}} + 8\Delta + 4\delta$
V verifies the validity of her receipt	$T + (N_v + 4)T_{\text{comp}} + 7\Delta + 4\delta$	$T + (N_v + 4)T_{\text{comp}} + 8\Delta + 4\delta$	$T + (N_v + 4)T_{\text{comp}} + 8\Delta + 4\delta$	$T + (N_v + 4)T_{\text{comp}} + 8\Delta + 4\delta$

Fig. A.1 – Time upper bounds at Clock, Clock**[V]**, Clock**[VC] and other honest VC nodes' clocks at each step of the interaction of the voter V with responder VC during D-DEMOS/IC voting phase. The grayed cells indicate the reference point of the clock drifts at each step.**

Definition 1 ([*d*]- *Patience*). Let V be an honest voter that submits her vote at some VC node when Clock[V] = T. We say that V is [*d*]-*patient*, when the following condition holds:

If V does not obtain a valid receipt by the time that Clock[V] = T + d, then she will blacklist this VC node and submit the same vote to another randomly selected VC node.

A.1.1. Liveness of D-DEMOS/IC

Using Definition 1, we prove the liveness of D-DEMOS/IC in the following theorem. A crucial step in the proof, is to compute an upper bound on the time required for an honest responder VC node to issue a receipt to V. This bound will be derived by the upper time bounds that correspond to each step of the voting protocol, as described in Sections 4.5.1 and 4.6, taking also into account the Δ and δ upper bounds. In Fig. A.1, we provide upper bounds on the advance of the global clock and the internal clocks of V and the VC nodes, so that we illustrate the description of the computation described below.

Theorem 1 (Liveness of D-Demos/IC). *Consider a D-DEMOS/IC run with n voters, m options and N_v VC nodes. Let $\mathcal{A}$ be an adversary against D-DEMOS/IC under the model described in Section 4.3 that corrupts up to $f_v < N_v/3$ VC nodes. Assume there is an upper bound Δ on clock synchronization loss and an upper bound δ on the delay of message delivery among honest VC nodes. Let T_{comp} be the worst-case running time of any procedure run by the VC nodes and the voters described in Sections 4.5.1 and 4.6, respectively, during the voting protocol.*

Let T_{end} denote the election end time. Define

$$T_{\text{wait}} := (N_v + 4)T_{\text{comp}} + 8\Delta + 4\delta.$$

Then, the following conditions hold:

1. *Every [T_{wait}]-patient voter V that is engaged in the voting protocol by the time that Clock[V] = $T_{\text{end}} - (f_v + 1) \cdot T_{\text{wait}}$, will obtain a valid receipt.*

2. *Every $[T_{\mathrm{wait}}]$-patient voter V that is engaged in the voting protocol by the time that $\mathrm{Clock}[V] = T_{\mathrm{end}} - y \cdot T_{\mathrm{wait}}$, where $y \in [f_v]$, will obtain a valid receipt with more than $1 - 3^{-y}$ probability.*

Proof. Let V be a $[T_{\mathrm{wait}}]$-patient voter initialized by the adversary $\mathcal{A}$ when $\mathrm{Clock} = \mathrm{Clock}[V] = T$. Upon initialization, V's internal clock is synchronized with the global clock at time $\mathrm{Clock} = \mathrm{Clock}[V] = T$. After at most T_{comp} steps, V submits her vote (serial$-$no, vote$-$code) at internal clock time: $\mathrm{Clock}[V] = T + T_{\mathrm{comp}}$, hence at global clock time: $\mathrm{Clock} \leq T + T_{\mathrm{comp}} + \Delta$. At that time, the internal clock of VC will be at most $T + T_{\mathrm{comp}} + 2\Delta$. Thus, VC will receive the vote of V at internal time $\mathrm{Clock}[VC] \leq T + T_{\mathrm{comp}} + 2\Delta + \delta$. Then, VC performs at most T_{comp} steps to verify the validity of the vote before it broadcasts its receipt share.

All the other honest VC nodes will receive VC's receipt share by global clock time:

$$\mathrm{Clock} \leq (T + T_{\mathrm{comp}} + 2\Delta + \delta) + (T_{\mathrm{comp}} + \Delta + \delta) = T + 2T_{\mathrm{comp}} + 3\Delta + 2\delta,$$

which implies that the time at their internal clocks is at most $T + 2T_{\mathrm{comp}} + 4\Delta + 2\delta$. Then, they will verify VC's share and broadcast their shares for V's vote after at most T_{comp} steps. The global clock at that point is no more than

$$\mathrm{Clock} \leq (T + 2T_{\mathrm{comp}} + 4\Delta + 2\delta) + T_{\mathrm{comp}} + \Delta = T + 3T_{\mathrm{comp}} + 5\Delta + 2\delta.$$

Therefore, VC will obtain the other honest VC nodes' shares at most when

$$\mathrm{Clock}[VC] \leq (T + 3T_{\mathrm{comp}} + 5\Delta + 2\delta) + \Delta + \delta = T + 3T_{\mathrm{comp}} + 6\Delta + 3\delta$$

and will process them in order to reconstruct the receipt for V. In order to collect $N_v - f_v - 1$ receipt shares that are sufficient for reconstruction, VC may have to perform up to $N_v - 1$ receipt-share verifications, as the f_v malicious VC nodes may also send invalid messages. This verification requires at most $(N_v - 1) \cdot T_{\mathrm{comp}}$ steps. Taking into account the T_{comp} steps for the reconstruction process, we conclude that VC will finish computation by global time

$$= (T + 3T_{\mathrm{comp}} + 6\Delta + 3\delta) + (N_v - 1)T_{\mathrm{comp}} + T_{\mathrm{comp}} + \Delta$$
$$= T + (N_v + 3)T_{\mathrm{comp}} + 7\Delta + 3\delta.$$

Finally, V will obtain the receipt after at most δ delay from the moment that VC finishes computation, and she needs T_{comp} steps to verify the validity of this receipt. Taking into consideration the drift on V's internal clock, we have that if V is honest and has not yet obtained a receipt by the time that

$$\mathrm{Clock}[V] = (T + (N_v + 3)T_{\mathrm{comp}} + 7\Delta + 3\delta) + T_{\mathrm{comp}} + \Delta + \delta = T + T_{\mathrm{wait}},$$

then, being $[T_{\mathrm{wait}}]$-patient, she can blacklist VC and resubmit her vote to another VC node. We will show that the latter fact implies conditions (1) and (2) in the statement of the theorem:

Condition (1): since there are at most f_v malicious VC nodes, V will certainly run into an honest VC node at her $(f_v + 1)$th attempt (if reached). Therefore, if V is engaged in the voting protocol by the time that $\mathrm{Clock}[V] = T_{\mathrm{end}} - (f_v + 1) \cdot T_{\mathrm{wait}}$, then she will obtain a receipt.

Condition (2): if V has waited for more than $y \cdot T_{\mathrm{wait}}$ time and has not yet received a receipt, then it has run at least y failed attempts in a row. At the jth attempt, V has $\frac{f_v - (j-1)}{N_v - (j-1)}$ probability to randomly select one of the remaining $f_v - (j-1)$ malicious VC nodes out of the $N_v - (j-1)$ non-blacklisted VC nodes. Thus, the probability that V runs at least y failed attempts in a row is

$$\prod_{j=1}^{y} \frac{f_v - (j-1)}{N_v - (j-1)} = \prod_{j=1}^{y} \frac{f_v - (j-1)}{3f_v + 1 - (j-1)} < 3^{-y}.$$

Therefore, if V is engaged in the voting protocol by the time that $\mathrm{Clock}[V] = T_{\mathrm{end}} - y \cdot T_{\mathrm{wait}}$, then the probability that she will obtain a receipt is more than $1 - 3^{-y}$. □

A.1.2. Liveness of D-DEMOS/Async
The proof of liveness in the asynchronous version of D-DEMOS differs from the one of D-DEMOS/IC in the computation of the T_{wait} upper bound, which now depends on the steps of the VC nodes presented in Section 4.5.2. The upper bounds on the advance of the global clock and the internal clocks of V and the VC nodes is analogously differentiated, as depicted in Fig. A.2.

Theorem 2 (Liveness of D-Demos/Async). *Consider a D-DEMOS/Async run with n voters, m options and N_v VC nodes. Let $\mathcal{A}$ be an adversary against D-DEMOS/Async with m options and n voters under the model described in Section 4.3 that corrupts up to $f_v < N_v/3$ VC nodes. Assume there is an upper bound Δ on clock synchronization loss and an upper bound δ on the delay of message delivery among honest VC nodes. Let T_{comp} be the worst-case running time of any procedure run by the VC nodes and the voters described in Sections 4.5.2 and 4.6, respectively, during the voting protocol.*

Let T_{end} denote the election end time. Define

$$T_{\mathrm{wait}} := (2N_v + 5)T_{\mathrm{comp}} + 12\Delta + 6\delta.$$

Then, the following conditions hold:

1. *Every $[T_{\mathrm{wait}}]$-patient voter that is engaged in the voting protocol by the time that $\mathrm{Clock}[V] = T_{\mathrm{end}} - (f_v + 1) \cdot T_{\mathrm{wait}}$, will obtain a valid receipt.*
2. *Every $[T_{\mathrm{wait}}]$-patient voter that is engaged in the voting protocol by the time that $\mathrm{Clock}[V] = T_{\mathrm{end}} - y \cdot T_{\mathrm{wait}}$, where $y \in [f_v]$, will obtain a valid receipt with more than $1 - 3^{-y}$ probability.*

Proof. The T_{wait} upper bound is computed according to the diagram in Fig. 10. Following the reasoning in the proof of Theorem 1, we get that

$$T_{\mathrm{wait}} := (2N_v + 5)T_{\mathrm{comp}} + 12\Delta + 6\delta.$$

Subsequently, we show that conditions (1) and (2) hold for any $[T_{\mathrm{wait}}]$-patient voter, exactly as in the proof of Theorem 1. □

A.2.   Safety

D-DEMOS's safety guarantee is expressed as a contract adhered by the VC subsystem, stated in Section 4.2. This contract is fulfilled by both D-DEMOS versions, though D-DEMOS/IC requires some additional assumptions to hold, as compared with D-DEMOS/Async that assumes only fault tolerance of the underlying subsystems (see Section 4.3). Moreover, the proofs of safety of the two versions diverge. Specifically, the safety

Step	Time upper bounds at each clock			
	Clock	Clock$[V]$	Clock$[VC]$	honest VC nodes' clocks
V is initialized	T	T	$T + \Delta$	$T + \Delta$
V submits her vote to VC	$T + T_{\text{comp}} + \Delta$	$T + T_{\text{comp}}$	$T + T_{\text{comp}} + 2\Delta$	$T + T_{\text{comp}} + 2\Delta$
VC receives V's ballot	$T + T_{\text{comp}} + \Delta + \delta$	$T + T_{\text{comp}} + 2\Delta + \delta$	$T + T_{\text{comp}} + 2\Delta + \delta$	$T + T_{\text{comp}} + 2\Delta + \delta$
VC verifies the validity of V's ballot and broadcasts an ENDORSE message	$T + 2T_{\text{comp}} + 3\Delta + \delta$	$T + 2T_{\text{comp}} + 4\Delta + \delta$	$T + 2T_{\text{comp}} + 2\Delta + \delta$	$T + 2T_{\text{comp}} + 4\Delta + \delta$
All the other honest VC nodes receive VC's ENDORSE message	$T + 2T_{\text{comp}} + 3\Delta + 2\delta$	$T + 2T_{\text{comp}} + 4\Delta + 2\delta$	$T + 2T_{\text{comp}} + 4\Delta + 2\delta$	$T + 2T_{\text{comp}} + 4\Delta + 2\delta$
All the other honest VC nodes verify the validity of the ENDORSE message and respond with an ENDORSEMENT message	$T + 3T_{\text{comp}} + 5\Delta + 2\delta$	$T + 3T_{\text{comp}} + 6\Delta + 2\delta$	$T + 3T_{\text{comp}} + 6\Delta + 2\delta$	$T + 3T_{\text{comp}} + 4\Delta + \delta$
VC receives the ENDORSEMENT messages of all the other honest VC nodes	$T + 3T_{\text{comp}} + 5\Delta + 3\delta$	$T + 3T_{\text{comp}} + 6\Delta + 3\delta$	$T + 3T_{\text{comp}} + 6\Delta + 3\delta$	$T + 3T_{\text{comp}} + 6\Delta + 3\delta$
VC verifies the validity of all the $N_v - 1$ received messages until it obtains $N_v - f_v$ valid ENDORSEMENT messages	$T + (N_v + 2)T_{\text{comp}} + 7\Delta + 3\delta$	$T + (N_v + 2)T_{\text{comp}} + 8\Delta + 3\delta$	$T + (N_v + 2)T_{\text{comp}} + 6\Delta + 3\delta$	$T + (N_v + 2)T_{\text{comp}} + 8\Delta + 3\delta$
VC forms UCERT certificate and broadcsts its share and UCERT	$T + (N_v + 3)T_{\text{comp}} + 7\Delta + 3\delta$	$T + (N_v + 3)T_{\text{comp}} + 8\Delta + 3\delta$	$T + (N_v + 3)T_{\text{comp}} + 6\Delta + 3\delta$	$T + (N_v + 3)T_{\text{comp}} + 8\Delta + 3\delta$
All the other honest VC nodes receive VC's broadcast share and UCERT	$T + (N_v + 3)T_{\text{comp}} + 7\Delta + 4\delta$	$T + (N_v + 3)T_{\text{comp}} + 8\Delta + 4\delta$	$T + (N_v + 3)T_{\text{comp}} + 8\Delta + 4\delta$	$T + (N_v + 3)T_{\text{comp}} + 8\Delta + 4\delta$
All the other honest VC nodes verify the validity of UCERT and V's share and broadcast their shares	$T + (N_v + 4)T_{\text{comp}} + 9\Delta + 4\delta$	$T + (N_v + 4)T_{\text{comp}} + 10\Delta + 4\delta$	$T + (N_v + 4)T_{\text{comp}} + 10\Delta + 4\delta$	$T + (N_v + 4)T_{\text{comp}} + 8\Delta + 4\delta$
VC receives all the other honest VC nodes' shares	$T + (N_v + 4)T_{\text{comp}} + 9\Delta + 5\delta$	$T + (N_v + 4)T_{\text{comp}} + 10\Delta + 5\delta$	$T + (N_v + 4)T_{\text{comp}} + 10\Delta + 5\delta$	$T + (N_v + 4)T_{\text{comp}} + 10\Delta + 5\delta$
VC verifies the validity of all the $N_v - 1$ received messages until it obtains $N_v - f_v$ valid shares	$T + (2N_v + 3)T_{\text{comp}} + 11\Delta + 5\delta$	$T + (2N_v + 3)T_{\text{comp}} + 12\Delta + 5\delta$	$T + (2N_v + 3)T_{\text{comp}} + 10\Delta + 5\delta$	$T + (2N_v + 3)T_{\text{comp}} + 12\Delta + 5\delta$
VC reconstructs and V's receipt and sends it to V	$T + (2N_v + 4)T_{\text{comp}} + 11\Delta + 5\delta$	$T + (2N_v + 4)T_{\text{comp}} + 12\Delta + 5\delta$	$T + (2N_v + 4)T_{\text{comp}} + 10\Delta + 5\delta$	$T + (2N_v + 4)T_{\text{comp}} + 12\Delta + 5\delta$
V obtains her receipt	$T + (2N_v + 4)T_{\text{comp}} + 11\Delta + 6\delta$	$T + (2N_v + 4)T_{\text{comp}} + 12\Delta + 6\delta$	$T + (2N_v + 4)T_{\text{comp}} + 12\Delta + 6\delta$	$T + (2N_v + 4)T_{\text{comp}} + 12\Delta + 6\delta$
V verifies the validity of her receipt	$T + (2N_v + 5)T_{\text{comp}} + 11\Delta + 6\delta$	$T + (2N_v + 5)T_{\text{comp}} + 12\Delta + 6\delta$	$T + (2N_v + 5)T_{\text{comp}} + 12\Delta + 6\delta$	$T + (2N_v + 5)T_{\text{comp}} + 12\Delta + 6\delta$

Fig. A.2 – Time upper bounds at Clock, Clock**[V]**, Clock**[VC] and other honest VC nodes' clocks at each step of the interaction of the voter V with responder VC during D-DEMOS/Async voting phase. The grayed cells indicate the reference point of the clock drifts at each step.**

of D-DEMOS/IC relies on the security of the fixed SHA-256 hash function and the AES-128-CBC$ symmetric encryption scheme. Therefore, the safety statement is with respect to specific security parameters. On the contrary, the safety of D-DEMOS/Async depends on the RSA signature scheme, therefore our analysis follows an asymptotic approach.

A.2.1. Safety of D-DEMOS/IC

As in liveness, we assume the upper bounds δ, Δ on the delay of message delivery and the drifts of all nodes' clocks to implement T_{end} and T_{barrier} as the starting point and the barrier of the IC protocol. We consider 128-bit security of the

commitment scheme assuming that every adversary running in less than 2^{64} steps has no more than 2^{-128} probability of obtaining any information about a single committed value (i.e., we set $c = 6/7$, where c is mentioned in Section 2.3.1).

Theorem 3 (Safety of D-Demos/IC). *Consider a D-DEMOS/IC run with n voters, m options, vote-codes of bit-length κ, N_v VC nodes, N_b BB nodes and N_t trustees, under the restrictions that (i) $m \cdot n \leq 2^{41}$ and (ii) $\kappa \geq 105$. Let $\mathcal{A}$ be an adversary against D-DEMOS under the model described in Section 4.3 that corrupts up to $f_v < N_v/3$ VC nodes, up to $f_b < N_b/2$ BB nodes and up to $N_t - h_t$ out-of N_t trustees. Assume there is an upper bound Δ on clock synchronization loss and*

an upper bound δ on the delay of message delivery. Let T_{end} be the end of the voting phase and $T_{barrier}$ be the end of the value dissemination phase of the interactive consistency protocol, as described in Section 4.3. Then, all honest voters who received a valid receipt from a VC node, are assured that their vote will be published on the honest BB nodes and included in the election tally, with probability at least

$$1 - \frac{nf_v}{2^{64} - f_v} - \left(3(mn)^3 + 2^{25}(mn)^2 + 2^{64}mn\right) \cdot 2^{-125}.$$

Proof. A crucial step for proving the safety of D-DEMOS/IC is to ensure it is hard for the adversary to compute non-submitted valid vote codes from the ballots of honest voters. This is done in the following claim.

CLAIM 3.1: *The probability that $\mathcal{A}$ outputs a vote code from the ballot of some honest voter V which was not cast by V is less than $\left(3(mn)^3 + 2^{25}(mn)^2 + 2^{64}mn\right) \cdot 2^{-125}$.*

Proof of Claim 3.1: Let **C** be the set of all vote codes generated by the EA. An arbitrary execution of $\mathcal{A}$ determines the following subsets of **C**: (i) the set of vote codes $\mathbf{C}_1$ that all honest voters submitted at the election phase, (ii) the set of the vote codes $\mathbf{C}_2$ located in unused ballots of honest voters that did not engage in the voting protocol and (iii) the set of vote codes $\mathbf{C}_3$ in the ballots of corrupted voters.

Since every vote code is a random κ-bit string, the event that $\mathcal{A}$ guesses some of the $2mn$ vote codes can happen with no more than $2mn(2^{-\kappa}) = 2^{-(\kappa-1)}mn$ probability. Furthermore, $\mathcal{A}$ is restricted by the fault tolerance thresholds of the VC, BB and trustees subsystems. Hence, by (i) the random vote code generation, (ii) the fault tolerance thresholds, (iii) the hiding property of the commitment scheme and (iv) the perfect simulatability of the zero-knowledge proofs, we assume that except for some probability bounded by $2^{-(\kappa-1)}mn + 0 + 2^{-(\kappa-1)}mn + 0 = 2^{-(\kappa-2)}mn$, the information associated with the vote codes that $\mathcal{A}$ obtains is,

  (i) The VC initialization data (for every VC node that $\mathcal{A}$ corrupts).
  (ii) All the BB initialization data. The part of these data that is associated with the vote codes is the list of all AES-128-CBC\$ vote code encryptions under msk.
  (iii) The set $\mathbf{C}_1 \cup \mathbf{C}_2 \cup \mathbf{C}_3$.

Reduction to IND-CPA security of AES-128-CBC\$. Given the code of $\mathcal{A}$, we construct an algorithm $\mathcal{B}$ against the $(t, q, (2t + 258 \cdot q + 3q^2) \cdot 2^{-128})$-IND-CPA security of the underlying AES-128-CBC\$ (see Section 2.3.4). Namely, $\mathcal{B}$ invokes $\mathcal{A}$ and attempts to simulate a setup and run of D-DEMOS/IC as follows:

1. $\mathcal{B}$ chooses a random triple $(j^*, \ell^*, X^*) \in [m] \times [n] \times \{A, B\}$.
2. For every $(j, \ell, X) \in [m] \times [n] \times \{A, B\} \setminus \{(j^*, \ell^*, X^*)\}$, $\mathcal{B}$ executes the following steps:
   (a) $\mathcal{B}$ chooses a random 64-bit $\text{vote-code}_{\ell, j}^X$ and associates it with $\text{option}_{\ell, j}^X$.
   (b) $\mathcal{B}$ makes an encryption query $(m_{0, \ell, j}^X, m_{1, \ell, j}^X) = (\text{vote-code}_{\ell, j}^X, \text{vote-code}_{\ell, j}^X)^X$ and receives an AES-128-CBC\$ encryption of $\text{vote-code}_{\ell, j}^X$.
   (c) $\mathcal{B}$ chooses a random $\text{salt}_{\ell, j}^X$ and computes $H_{\ell, j}^X \leftarrow SHA256(\text{vote-code}_{\ell, j}^X, \text{salt}_{\ell, j}^X)$.

   (d) $\mathcal{B}$ generates the cryptographic payload $\text{payload}_{\ell, \pi_\ell^X(j)}$ associated with $\text{option}_{\ell, j}^X$.
3. $\mathcal{B}$ chooses random values $\text{vote-code}_0^*, \text{vote-code}_1^* \in \{0, 1\}^{64}$, $\text{salt}^* \in \{0, 1\}^{64}$.
4. $\mathcal{B}$ makes the encryption query challenge $\text{vote-code}_0^*, \text{vote-code}_1^*$ and receives the AES-128-CBC\$ encryption y^* of vote-code_b^*, where b is the outcome of a coin-flip.
5. $\mathcal{B}$ tabulates BB initialization data as EA does, by using vote-code_0^* as the vote code associated with $\text{option}_{\ell^*, j^*}$, the hash $SHA256(\text{vote-code}_0^*, \text{salt}^*)$ as H_{ℓ^*, j^*}^X and y^* as the AES-128-CBC\$ ciphertext that corresponds to vote-code_0^*.
6. $\mathcal{B}$ interacts with $\mathcal{A}$ according to the model described in Section 4.3.
7. If $\mathcal{A}$ outputs vote-code_0^*, then $\mathcal{B}$ outputs 0. Otherwise, $\mathcal{B}$ outputs 1.

Let G be the event that $\mathcal{A}$ outputs some $\text{vote-code} \in \mathbf{C} \setminus (\mathbf{C}_1 \cup \mathbf{C}_2 \cup \mathbf{C}_3)$. By the construction of $\mathcal{B}$, if the IND-CPA challenge bit b is 0, then $\mathcal{B}$ simulates a D-DEMOS/IC election perfectly. Furthermore, if $b = 0$ and vote-code corresponds to the randomly chosen position $(j^*, \ell^*, X^*) \in [m] \times [n] \times \{A, B\}$, then it outputs 0 ($\text{vote-code} = \text{vote-code}_0^*$). Since $\mathcal{B}$ randomly guesses the triple (ℓ^*, j^*, X^*), we have that

$$\Pr[\mathcal{B} \text{ outputs } 1 \mid b = 0] = 1 - \Pr[\mathcal{B} \text{ outputs } 0 \mid b = 0]$$

$$= 1 - \frac{\Pr[G \mid b = 0]}{2mn}. \tag{A.1}$$

On the other hand, if $b = 1$, then vote-code_0^* is the preimage of $SHA256(\text{vote-code}_0^*, \text{salt}^*)$, while y^* is the encryption of an independently generated vote code. Based on this observation, we construct an algorithm $\mathcal{C}$ that acts as an attacker against the $(t, t^2 \cdot 2^{-256})$-collision resistance of SHA-256 (see Section 2.3.3). Namely, on input some hash value H, $\mathcal{C}$ executes the following steps:

1. $\mathcal{C}$ chooses a random triple $(j^*, \ell^*, X^*) \in [m] \times [n] \times \{A, B\}$.
2. For every $(j, \ell, X) \in [m] \times [n] \times \{A, B\}$, $\mathcal{C}$ chooses random values $\text{vote-code}_{\ell, j}^X \in \{0, 1\}^{160}$, $\text{salt}_{\ell, j}^X \in \{0, 1\}^{64}$.
3. $\mathcal{C}$ tabulates all election information normally except that for (ℓ^*, j^*, X^*) it provides H instead of the hash value $SHA256(\text{vote-code}_{\ell, j}^{X^*}, \text{salt}_{\ell, j}^{X^*})$.
4. $\mathcal{C}$ interacts with $\mathcal{A}$ according to the model described in Section 4.3.
5. $\mathcal{C}$ receives the output of $\mathcal{A}$, labeled by z.
6. $\mathcal{C}$ searches for a $w \in \{0, 1\}^{64}$ s.t. $h(z, w) = H$. If $\mathcal{C}$ finds such a w, then it outputs $z \| w$. Otherwise, it aborts.

For simplicity and w.l.o.g., we can assume that for each $(j, \ell, X) \in [m] \times [n] \times \{A, B\}$, the time complexity for information preparation is on the order of 256^3 (cube of the string length, set to 256 bits). The running time of $\mathcal{A}$ is 2^{64}. Assuming linear complexity for hashing and checking a random value, the brute force search for the correct w in step 6. takes $2^{64} \cdot 256 = 2^{72}$ steps. Therefore, given that $mn \leq 2^{41}$, we conclude the $\mathcal{C}$ runs in steps bounded by $2mn \cdot 256^3 + 2^{64} + 2^{64} \cdot 256 \leq mn2^{25} + 2^{64} + 2^{72} < 2^{73}$.

By the $(t, t^2 \cdot 2^{-256})$-collision resistance of $h(\cdot)$ (see Section 2.3.3), the probability that $\mathcal{C}$ finds a preimage of H is less than $2^{146} \cdot 2^{-256} < 2^{-110}$. By the construction of $\mathcal{C}$,

if $\mathcal{A}$ outputs the vote code that corresponds to position $(\ell^*, j^*, X^*) \in [n] \times [m] \times \{A, B\}$, then $\mathcal{C}$ certainly wins. Therefore, we have that

$$\Pr[\mathcal{B} \text{ outputs } 1 \mid b = 1] = 1 - \Pr[\mathcal{B} \text{ outputs } 0 \mid b = 1]$$

$$= 1 - \frac{\Pr[G \mid b = 1]}{2mn} - 2^{-(\kappa-2)}mn \geq$$

$$\geq 1 - \Pr[\mathcal{C} \text{ returns the preimage of SHA-256}]$$

$$> 1 - 2^{-110} - 2^{-(\kappa-2)}mn. \tag{A.2}$$

Hence, by Eq. (A.1) and (A.2), we conclude that

$$\mathbf{Adv}^{\text{IND–CPA}}_{128-\text{AES–CBC\$}}(\mathcal{B}) > \frac{\Pr[G \mid b = 0]}{2mn} - 2^{-110} - 2^{-(\kappa-2)}mn. \tag{A.3}$$

Along the lines of the time complexity analysis of $\mathcal{C}$, the time complexity of $\mathcal{B}$ is bounded by $2mn \cdot 256^3 + 2^{64} = 2^{25}mn + 2^{64} < 2^{66}$, where we used that $mn \leq 2^{41}$. In addition, $\mathcal{B}$ makes at most $2 \cdot m \cdot n$ queries. Hence, by the $(t, q, (2t+258 \cdot q+3q^2) \cdot 2^{-128})$-IND-CPA security of AES-CBC\$ (see Section 2.3.4) and (A.3), we conclude that for every vote-code length $\kappa \geq 105$,

$$\frac{\Pr[G \mid b = 0]}{2mn} - 2^{-110} - 2^{-(\kappa-2)}mn$$

$$< (2^{26}mn + 2^{65} + 516mn + 12(mn)^2) \cdot 2^{-128}$$

$$\Rightarrow \Pr[G \mid b = 0] - 2^{-109}mn - 2^{-(\kappa-3)}(mn)^2$$

$$< (2^{27}(mn)^2 + 2^{66}mn + 2^9(mn)^2 + 24(mn)^3) \cdot 2^{-128} \Leftrightarrow \Pr[G \mid b = 0]$$

$$< ((2^{66}+2^{19}) \cdot mn + (2^{27}+2^9+2^{(131-\kappa)}) \cdot (mn)^2 + 24 \cdot (mn)^3) \cdot 2^{-128}$$

$$\Rightarrow \Pr[G \mid b = 0] < (2^{67}mn + 2^{28}(mn)^2 + 24(mn)^3) \cdot 2^{-128}$$

$$\Leftrightarrow \Pr[G \mid b = 0] < (3(mn)^3 + 2^{25}(mn)^2 + 2^{64}mn) \cdot 2^{-125}, \tag{A.4}$$

which completes the proof of the claim, as the election simulation for $b = 0$ is perfect.

(End of Claim 3.1) ⊣

Given Claim 3.1, the proof is completed in two stages.

1. *Vote set consensus.* By the upper bound restriction on all clock drifts, all honest VC nodes will enter the Value Dissemination phase at T_{end} and the Result Consensus phase of the Interactive Consistency protocol at T_{barrier} within some distance Δ from the global clock. The agreement property of interactive consistency ensures that all honest VC nodes will contain the same vector $\langle VS_1, \ldots, VS_n \rangle$ of all nodes' sets of voted and pending ballots. Subsequently, all honest VC nodes, execute the same deterministic algorithm of Fig. 9, and will agree on the same set of votes denoted by Votes. This will be the set of votes that are marked to be tallied by the honest VC nodes.

2. *Protocol contract.* Let V_ℓ be an honest voter that has obtained a receipt for his vote ⟨serial−no, vote−code⟩, but his vote is not included in Votes. By the vote consensus property proved previously, we have that some honest VC node VC, decided to discard V_ℓ's vote. According to the algorithm described in Fig. 9 that determines Votes, the latter can happen only because either *Case (i)*: $\mathcal{A}$ succeeds in guessing the valid receipt of V_ℓ, or *Case (ii)*: a vote−code−2 different than vote−code appears in the list for the ballot indexed by serial−no or *Case (iii)*: vote−code appears less than $N_v - 2f_v$ times in the list for the ballot indexed by serial−no. We study all Cases (i),(ii),(iii):

Case (i). If $\mathcal{A}$ succeeds in guessing a valid receipt, then it can force the VC subsystem to consider V's ballot not voted by not participating in the receipt reconstruction. By the information theoretic security of the VSS scheme, given that $\mathcal{A}$ is restricted by the fault tolerance thresholds, its guess of the receipt must be at random. Since there are at most f_v malicious VC nodes, the adversary has at most f_v attempts to guess the receipt. Moreover, the receipt is a randomly generated 64-bit string, so after i attempts, $\mathcal{A}$ has to guess among $(2^{64} - i)$ possible choices. Taking a union bound for n voters, the probability that $\mathcal{A}$ succeeds for any of the obtained receipts is no more than

$$\sum_{\ell=1}^{n} \left(\sum_{i=0}^{f_v - 1} \frac{1}{2^{64} - i} \right) \leq \frac{n f_v}{2^{64} - f_v}.$$

Case (ii). V_ℓ is honest, hence it has submitted the same vote in every possible attempt to vote prior to the one she obtained her receipt. Therefore, Case (ii) may occur only if the adversary $\mathcal{A}$ manages to produce vote−code−2 by the vote code related election information it has access to. Namely, (a) the set of vote codes that all honest voters submitted at the election phase, (b) the set of the vote codes that were located in unused ballots and (c) the set of vote codes in the ballots of corrupted voters. By assumption, vote−code−2 is in neither of these three sets. Hence, by Claim 3.1, the probability that $\mathcal{A}$ computes vote−code−2 is less than $(3(mn)^3 + 2^{25}(mn)^2 + 2^{64}mn) \cdot 2^{-125}$.

Case (iii). In order for V_ℓ to obtain a receipt, at least $N_v - f_v$ VC nodes must collaborate by providing their shares. The faulty VC nodes are at most f_v, so at least $N_v - 2f_v$ honest VC nodes will include ⟨serial−no, vote−code⟩ in their set of voted and pending ballots. Thus, Case (iii) cannot occur.

Consequently, all the honest VC nodes will forward the agreed set of votes (hence, also V_ℓ's vote) to the BB nodes. By the fault tolerance threshold for the BB subsystem, the f_b honest BB nodes will publish V_ℓ's vote. Finally, the h_t out-of N_t honest trustees will read V's vote from the majority of BB nodes and include it in the election tally. Therefore, the probability that $\mathcal{A}$ achieves in excluding the vote of at least one honest voter that obtained a valid receipt from the BB or the election tally is less than $\frac{n f_v}{2^{64} - f_v} - (3(mn)^3 + 2^{25}(mn)^2 + 2^{64}mn) \cdot 2^{-125}$, which completes the proof. □

A.2.2. Safety of D-DEMOS/Async

The safety of D-DEMOS/Async is founded on the certificate generation mechanism among the VC nodes, which in turn exploits the security of the underlying signature scheme.

Theorem 4 (Safety of D-Demos/Async). *Let $\mathcal{A}$ be an adversary against D-DEMOS under the model described in Section 4.3 that corrupts up to $f_v < N_v/3$ VC nodes, up to $f_b < N_b/2$ BB nodes and up to $N_t - h_t$ out-of N_t trustees. Then, all honest voters who received a valid receipt from a VC node, are assured that their vote will be published on the honest BB nodes and included in the election tally, with probability at least*

$$1 - \frac{n f_v}{2^{64} - f_v} - \text{negl}(\lambda).$$

Proof. Let V_ℓ be an honest voter. Then, $\mathcal{A}$'s strategy on attacking safety (i.e., provide a valid receipt to V_ℓ but force the VC

subsystem to discard V's ballot), is captured by either one of the two following cases: *Case (i):* $\mathcal{A}$ produces the receipt without being involved in a complete interaction with the VC subsystem (i.e., with at least $f_v + 1$ honest VC nodes). *Case (ii):* $\mathcal{A}$ provides a properly reconstructed receipt via a complete interaction with the VC subsystem (in both cases we assume $\mathcal{A}$ controls the *responder* VC node).

Let E_1 (resp. E_2) be the event that Case 1 (resp. Case 2) happens. We study both cases:

Case (i). In this case, $\mathcal{A}$ must produce a receipt that matches V's ballot with less than $N_v - f_v$ shares. $\mathcal{A}$ may achieve this by either one of the following ways:

1. $\mathcal{A}$ attempts to guess the valid receipt; If $\mathcal{A}$ succeeds, then it can force the VC subsystem to consider V's ballot not voted as no valid UCERT certificate will be generated for V's ballot (malicious *responder* does not send an ENDORSE message). As shown in the proof of Theorem 3, the probability of a successful guess for $\mathcal{A}$ is less than $\frac{nf_v}{2^{64}-f_v}$.
2. $\mathcal{A}$ attempts to produce fake UCERT certificates by forging digital signatures of other nodes. By the security of the digital signature scheme, this attack has $\mathrm{negl}(\lambda)$ success probability.

By the above, we have that $\Pr[\mathcal{A} \text{ wins } |E_1] \leq \frac{nf_v}{2^{64}-f_v} + \mathrm{negl}(\lambda)$.

Case (ii). In this case, by the security arguments stated in Section 4.5 (steps 1–5), every honest VC node will include the vote of V_ℓ in the set of voted tuples. This is because (a) it locally knows the valid (certified) vote code for V_ℓ which is accompanied by UCERT or (b) it has obtained the valid vote code via a RECOVER-REQUEST message. Recall that unless there are fake certificates (which happens with negligible probability) there can be only one valid vote code for V_ℓ.

Consequently, all the honest VC nodes will forward the agreed set of votes (hence, also V_ℓ's vote) to the BB nodes. By the fault tolerance threshold for the BB subsystem, the f_b honest BB nodes will publish V's vote. Finally, the h_t out-of N_t honest trustees will read V_ℓ's vote from the majority of BB nodes and include it in the election tally. Thus, we have that $\Pr[\mathcal{A} \text{ wins } |E_2] = \mathrm{negl}(\lambda)$.

Therefore, all the votes of honest voters that obtained a valid receipt, will be published on the honest BB nodes and included in the election tally, with probability at least

$$1 - \Pr[\mathcal{A} \text{ wins }] \geq 1 - \Pr[\mathcal{A} \text{ wins } |E_1] - \Pr[\mathcal{A} \text{ wins } |E_2]$$
$$\geq 1 - \frac{nf_v}{2^{64}-f_v} - \mathrm{negl}(\lambda).$$

□

A.2.3. Usability vs. security trade-off

Theorem 3 statement and proof shed light on the limitations of D-DEMOS regarding the usability vs. security trade off. In particular, the specifications of the underlying cryptographic tools (SHA-256 and AES-CBC$), as formally expressed in Eq. (A.4), dictate the use of vote-codes with at least 105 bit-length. The latter implies that voters are provided with vote-codes of 14 characters in alphanumeric form, encoded in Base64, a size that lies between the length of a credit card number and a Microsoft product key. Besides, the restriction that $mn \leq 2^{41}$ is expected to be met in most election scenarios

(e.g., up to 10^3 options and 10^9 voters), hence D-DEMOS is fully scalable from a safety perspective.

In an alternative D-DEMOS specification where ciphers of bigger block length are applied (e.g. an 192 block-length cipher from the Rijndael family), the right hand of the inequality in Eq. (A.4) that refers to the symmetric encryption security error becomes very small. As a consequence, the term $2^{-(\kappa-3)}(mn)^2$ denoting the probability the adversary guesses a valid vote-code of length κ dominates over the cryptographic error in the total upper bound of the adversary's success probability. The parameters m, n, κ can be fixed so that $2^{-(\kappa-3)}(mn)^2$ is sufficiently low even for smaller vote-code lengths, thus increasing D-DEMOS's usability, especially in mid-scale election scenarios. Indicatively, if we require that $2^{-(\kappa-3)}(mn)^2 < \frac{2^{-10}}{2mn}$ (thus, by the union bound, the adversary has less than 0.1% probability to guess even a single vote-code), and by fixing $mn = 2^{20}$ (e.g., 10 options and 10^6 voters), we get that $\kappa \geq 75$. In Base64 encoding, this means that vote-codes of 10 characters are required; note that many Internet sites recommend password lengths in the 8–12 character range.

A.3. End-to-end verifiability

We adopt the end-to-end (E2E) verifiability definition in Kiayias et al. (2015), modified accordingly to our setting. Namely, we encode the options set $\{\mathrm{option}_1, \ldots, \mathrm{option}_m\}$, where the encoding of option_i is an m-bit string which is 1 only in the ith position. Let F be the *election evaluation function* such that $F(\mathrm{option}_{i_1} \ldots, \mathrm{option}_{i_n})$ is equal to an m-vector whose ith location is equal to the number of times option_i was voted. Then, we use the metric d_1 derived by the L1-norm scaled to half, i.e., $d_1(R, R') = \frac{1}{2} \cdot \sum_{i=1}^{n} |R_i - R'_i|$, where R_i, R'_i is the ith coordinate of R, R', respectively, to measure the success probability of the adversary with respect to the amount of tally deviation d and the number of voters that perform audit θ. In addition, we make use of a *vote extractor* algorithm $\mathcal{E}$ (not necessarily running in polynomial-time) that extracts the non-honestly cast votes.

We define E2E verifiability via an attack game between a challenger and an adversary specified in detail in Fig. A.3.

Definition 2 (E2E VERIFIABILITY). Let $0 < \epsilon < 1$ and $n, m, N_v, N_b, N_t \in \mathbb{N}$ polynomial in the security parameter λ with $\theta \leq n$. Let Π be an e-voting system with n voters, N_v VC nodes, N_b BB nodes and N_t trustees. We say that Π achieves *end-to-end verifiability* with error ϵ, w.r.t. the election function F, a number of θ honest successful voters and tally deviation d if there exists a (not necessarily polynomial-time) vote extractor $\mathcal{E}$ such that for any PPT adversary $\mathcal{A}$ it holds that

$$\Pr[G_{e2e-ver}^{\mathcal{A},\mathcal{E},d,\theta}(1^\lambda, m, n, N_v, N_b, N_t) = 1] \leq \epsilon.$$

To prove E2E verifiability of D-DEMOS, we need a min-entropy variant of the Schwartz-Zippel lemma, to check the equality of two univariate polynomials p_1, p_2, i.e., test $p_1(x) - p_2(x) = 0$ for random $x \xleftarrow{D} \mathbb{Z}_q$, where q is prime. The probability that the test passes is at most $\frac{\max(d_1, d_2)}{2^\kappa}$ if $p_1 \neq p_2$, where d_i is the degree of p_i for $i \in \{1, 2\}$. We leverage Lemma 1 from Kiayias et al. (2015).

E2E Verifiability Game $G_{\text{e2e-ver}}^{\mathcal{A},\mathcal{E},d,\theta}(1^\lambda, m, n, N_v, N_b, N_t)$

(i). $\mathcal{A}$ on input $1^\lambda, n, m, N_v, N_b, N_t$, chooses a list of options $\{\text{option}_1, \dots, \text{option}_m\}$, a set of voters $\mathcal{V} = \{V_1, \dots, V_n\}$, a set of VC nodes $\mathcal{VC} = \{\text{VC}_1, \dots, \text{VC}_{N_v}\}$, a set of BB nodes $\mathcal{BB} = \{\text{BB}_1, \dots, \text{BB}_{N_b}\}$, and a set of trustees $\mathcal{T} = \{T_1, \dots, T_{N_t}\}$. It provides the challenger Ch with all the above sets. Throughout the game, $\mathcal{A}$ controls the EA, all the VC nodes and all the trustees. In addition, $\mathcal{A}$ may corrupt a fixed set of less than $\lfloor N_b/2 \rfloor$ BB nodes, denoted by $\mathcal{BB}_{\text{succ}}$ (i.e., the majority of the BB nodes remain honest). On the other hand, Ch plays the role of all the honest BB nodes.

(ii). $\mathcal{A}$ and $\mathcal{C}$ engage in an interaction where $\mathcal{A}$ schedules the vote casting executions of all voters. For each voter V_ℓ, $\mathcal{A}$ can either completely control the voter or allow $\mathcal{C}$ to operate on V_ℓ's behalf, in which case $\mathcal{A}$ provides $\mathcal{C}$ with an option selection option_{i_ℓ}. Then, $\mathcal{C}$ casts a vote for option_{i_ℓ}, and, provided the voting execution terminates successfully, $\mathcal{C}$ obtains the audit information audit_ℓ on behalf of V_ℓ.

(iii). Finally, $\mathcal{A}$ posts a version of the election transcript info_j in every honest BB node $\text{BB}_j \notin \mathcal{BB}_{\text{corr}}$.

Let $\mathcal{V}_{\text{succ}}$ be the set of honest voters (i.e., those controlled by $\mathcal{C}$) that terminated successfully. The game returns a bit which is 1 if and only if the following conditions hold true:

(1) $\forall \text{BB}_j, \text{BB}_{j'} \notin \mathcal{BB}_{\text{corr}} : \text{info}_j = \text{info}_{j'} := \text{info}$

(2) $|\mathcal{V}_{\text{succ}}| \geq \theta$ (i.e., at least θ honest voters terminated).

(3) $\forall \ell \in [n]$: if $V_\ell \in \mathcal{V}_{\text{succ}}$ then V_ℓ verifies successfully, when given $(\text{info}, \text{audit}_\ell)$ as input.

and either one of the following two conditions:

(4)   (a) if $\bot \neq \langle \text{option}_{i_\ell} \rangle_{V_\ell \notin \mathcal{V}_{\text{succ}}} \leftarrow \mathcal{E}(\text{info}, \{\text{audit}_\ell\}_{V_\ell \in \mathcal{V}_{\text{succ}}})$ then

$$d_1\left(\mathbf{Result}(\text{info}), F(\text{option}_{i_1} \dots, \text{option}_{i_n})\right) \geq d .$$

(b) $\bot \leftarrow \mathcal{E}(\text{info}, \{\text{audit}_\ell\}_{V_\ell \in \mathcal{V}_{\text{succ}}})$.

Fig. A.3 – The E2E Verifiability Game between the challenger $\mathcal{C}$ and the adversary $\mathcal{A}$ using the vote extractor $\mathcal{E}$.

Lemma 1 (Min-entropy Schwartz-Zippel Kiayias et al. (2015)). *Let q be a prime and $p(x)$ be a non-zero univariate polynomial of degree d over $\mathbb{Z}_q$. Let D be a probability distribution on $\mathbb{Z}_q$ such that $H_\infty(D) \geq \kappa$. The probability of $p(x) = 0$ for a randomly chosen $x \xleftarrow{D} \mathbb{Z}_q$ is at most $\frac{d}{2^\kappa}$.*

We now analyse the soundness of the zero knowledge proof for each option encoding commitment. Note that a correct option encoding is an m-vector, where one of the m elements is 1 and the rest elements are 0 (a.k.a. unit vector). Our zero knowledge proof utilizes the Chaum–Pedersen DDH-tuple proofs (Chaum and Pedersen, 1993) in conjunction with the Sigma OR-composition technique (Cramer et al., 1994) to show each (lifted) ElGamal ciphertext encrypts either 0 or 1 and the product of all the m ElGamal ciphertexts encrypts 1. We adopt the soundness amplification technique from Kiayias et al. (2015); namely, if the voters' coins $\mathbf{c}$ are longer than $\lfloor \log q \rfloor$ then we divide it into κ blocks, $(\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_\kappa)$ such that each block has less than $\lfloor \log q \rfloor$ coins, where q is the order of the underlying group used in the ElGamal encryption. Given a statement x, for each $\mathbf{c}_i$, $i \in [\kappa]$, the prover needs to produce the zero knowledge transcript $(x, \phi_{1,i}, \mathbf{c}_i, \phi_{2,i})$ in order. The verifier ac-

cepts the proof if and only if for all $i \in [\kappa]$, $\text{Verify}(x, \phi_{1,i}, \mathbf{c}_i, \phi_{2,i}) = \text{accept}$. Hence, we have the following Lemma 2.

Lemma 2. *Denote $\mathbf{c} = (\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_\kappa)$. If $H_\infty(\mathbf{c}) = \theta$, we have for all adversaries $\mathcal{A}$:*

$$\varepsilon(m, n, \theta, \kappa) = \Pr\left[\begin{array}{l} (x, \{\phi_{1,i}\}_{i\in[\kappa]}) \leftarrow \mathcal{A}(1^\lambda); \\ \{\phi_{2,i}\}_{i\in[\kappa]} \leftarrow \mathcal{A}(\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_\kappa) : \\ x \text{ is not a valid option encoding commitment} \\ \wedge \forall i \in [\kappa], \text{Verify}(x, \phi_{1,i}, \mathbf{c}_i, \phi_{2,i}) = \text{accept} \end{array}\right]$$

$$\leq 2^{-\theta} .$$

Proof. For $i \in \kappa$, denote $H_\infty(\mathbf{c}_i) = \theta_i$, and $\sum_{i=1}^\kappa \theta_i = \theta$. Chaum–Pedersen DDH-tuple proof (Chaum and Pedersen, 1993) internally constructs and checks a degree-1 polynomial; therefore according to Lemma 1, the probability that the adversary $\mathcal{A}$ to cheat a single DDH-tuple zero knowledge proof is at most $2^{-\theta'}$, where θ' is the min-entropy of the challenge. Moreover, Sigma OR-composition technique (Cramer et al., 1994) perfectly maintains the soundness, so the probability that the adversary $\mathcal{A}$ to cheat the zero knowledge proofs for each (lifted) ElGamal ciphertext encrypts 0/1 is at most $2^{-\theta'}$. Note that the zero knowledge proofs of the

option encoding commitment is AND-composition of all the elementary zero knowledge proofs, the probability that x is invalid and $\mathsf{Verify}(x, \phi_{1,i}, \mathbf{c}_i, \phi_{2,i}) = \mathsf{accept}$ is at most $2^{-\theta_i}$. Hence, the probability that $\forall i \in [\kappa]$, $\mathsf{Verify}(x, \phi_{1,i}, \mathbf{c}_i, \phi_{2,i}) = \mathsf{accept}$ is $\varepsilon(m, n, \theta, \kappa) = \prod_{i=1}^{\kappa} 2^{-\theta_i} = 2^{-\sum_{i=1}^{\kappa} \theta_i} = 2^{-\theta}$. □

Applying Lemma 2, we prove that D-DEMOS (both the IC and the Async version) achieves E2E verifiability according to Definition 2.

Proof. Without loss of generality, we can assume that every party can read consistently the data published in the majority of the BB nodes, as otherwise the adversary fails to satisfy condition 1 of the E2E verifiability game.

We first construct a vote extractor $\mathcal{E}$ for D-DEMOS as follows:

- $\mathcal{E}$ takes input as the election transcript, info and a set of audit information $\{\mathsf{audit}_\ell\}_{V_\ell \in \mathcal{V}_{\mathsf{succ}}}$. If info is not meaningful, then $\mathcal{E}$ outputs $\bot$.
- Let $B \leq |\tilde{\mathcal{V}}|$ be the number of different serial numbers that appear in $\{\mathsf{audit}_\ell\}_{V_\ell \in \tilde{\mathcal{V}}}$. $\mathcal{E}$ (arbitrarily) arranges the voters in $V_\ell \in \mathcal{V}_{\mathsf{succ}}$ and the serial numbers not included in $\{\mathsf{audit}_\ell\}_{V_\ell \in \mathcal{V}_{\mathsf{succ}}}$ as $\langle V_\ell^{\mathcal{E}} \rangle_{\ell \in [n - |\mathcal{V}_{\mathsf{succ}}|]}$ and $\langle \mathsf{tag}_\ell^{\mathcal{E}} \rangle_{\ell \in [n-B]}$ respectively.
- For every $\ell \in [n - |\mathcal{V}_{\mathsf{succ}}|]$, $\mathcal{E}$ extracts option_{i_ℓ} by brute force opening and decrypting (in superpolynomial time) all the committed and encrypted BB data, or sets option_{i_ℓ} as the zero vector, in case V_ℓ's vote is not published in the BB.
- If there is an invalid option-commitment (i.e., it is not a commitment to some candidate encoding), then $\mathcal{E}$ outputs $\bot$. Otherwise, it outputs $\langle \mathsf{option}_{i_\ell} \rangle_{V_\ell \notin \mathcal{V}_{\mathsf{succ}}}$.

We will prove the E2E verifiability of D-DEMOS based on $\mathcal{E}$. Assume an adversary $\mathcal{A}$ that wins the game $G_{\mathsf{e2e-ver}}^{\mathcal{A},\mathcal{E},d,\theta}(1^\lambda, m, n, N_v, N_b, N_t)$. Namely, $\mathcal{A}$ breaks E2E verifiability by allowing at least θ honest successful voters and achieving tally deviation d.

Let Z be the event that $\mathcal{A}$ attacks by making at least one of the option-encoding commitments associated with some cast vote code invalid (i.e., it is in tally set $\mathbf{E}_{\mathsf{tally}}$ but it is not a commitment to some candidate encoding). By condition 2, there are at least θ honest and successful voters, hence the min-entropy of the collected voters' coins is at least θ. By Lemma 2, the zero-knowledge proofs used in D-DEMOS for committed ballot correctness in the BB is sound except for some probability error $2^{-\theta}$. Since $\theta \geq 1$ and condition 3 holds, there is at least one honest voter that verifies, thus we have that $\Pr[G_{\mathsf{e2e-ver}}^{\mathcal{A},\mathcal{E},d,\theta}(1^\lambda, m, n, N_v, N_b, N_t) = 1 \wedge Z] \leq 2^{-\theta}$.

Now assume that Z does not occur. In this case, the vote extractor $\mathcal{E}$ will output the intended adversarial votes up to permutation. Thus, the deviation from the intended result that $\mathcal{A}$ achieves, derives only by miscounting the honest votes. This may be achieved by $\mathcal{A}$ in two different possible ways:

- *Modification attacks.* When committing to the information of some honest voter's ballot part $\mathcal{A}$ changes the vote code and option correspondence that is printed in the ballot. This attack will be detected if the voter does choose to au-

dit with the modified ballot part (it uses the other part to vote). The maximum deviation achieved by this attack is 1 (the vote will count for another candidate).

- *Clash attacks.* $\mathcal{A}$ provides y honest voters with ballots that have the same serial number, so that the adversary can inject $y - 1$ votes of his preference in the $y - 1$ "empty" audit locations in the BB. This attack is successful only if all the y voters verify the same ballot on the BB and hence miss the injected votes that produce the tally deviation. The maximum deviation achieved by this attack is $y - 1$.

We stress that if Z does not occur, then the above two attacks are the only meaningful[3] for $\mathcal{A}$ to follow. Indeed, if (i) all zero knowledge proofs are valid, (ii) all the honest voters are pointed to a unique audit BB location indexed by the serial number on their ballots, and (iii) the information committed in this BB location matches the vote code and option association in the voters' unused ballot parts, then by the binding property of the commitments, all the tally computed by the commitments included in $\mathbf{E}_{\mathsf{tally}}$ will decrypt to the actual intended result.

Since the honest voters choose the used ballot parts at random, the success probability of x deviation via the modification attack is $(1/2)^x$. In addition, the success probability to clash y honest voters is $(1/2)^{y-1}$ (all y honest voters choose the same version to vote). As a result, by combinations of modification and clash attacks, $\mathcal{A}$'s success probability reduces by a factor $1/2$ for every unit increase of tally deviation. Therefore, the upper bound of the success probability of $\mathcal{A}$ when Z does not occur is $\Pr[G_{\mathsf{e2e-ver}}^{\mathcal{A},\mathcal{E},d,\theta}(1^\lambda, m, n, N_v, N_b, N_t) = 1 \mid \neg Z] \leq 2^{-d}$.

Hence, we conclude that $\Pr[G_{\mathsf{e2e-ver}}^{\mathcal{A},\mathcal{E},d,\theta}(1^\lambda, m, n, N_v, N_b, N_t) = 1] \leq 2^{-\theta} + 2^{-d}$. □

Applying Lemma 2, the following theorem states that D-DEMOS (both the IC and the Async version) achieves E2E verifiability according to Definition 2.

Theorem 5 (E2E Verifiability of D-Demos). *Let $n, m, N_v, N_b, N_t, \theta, d \in \mathbb{N}$ where $1 \leq \theta \leq n$. Then, D-DEMOS run with n voters, m options, N_v VC nodes, N_b BB nodes and N_t trustees achieves end-to-end with error $2^{-\theta} + 2^{-d}$, w.r.t. the election function F, a number of θ honest successful voters and tally deviation d.*

Proof. (*Sketch*). Without loss of generality, we can assume that every party can read consistently the data published in the majority of the BB nodes, as otherwise the adversary fails to satisfy condition 1 of the E2E verifiability game. Via brute force search, the vote extractor $\mathcal{E}$ for D-DEMOS either (i) decrypts the adversarial votes (up to permutation) if all respective option-encoding commitments are valid, or (ii) aborts otherwise. We analyze the two cases

(i) If all option-encoding commitments are valid, then the output of $\mathcal{E}$ implies that the tally deviation that the adversary $\mathcal{A}$ can achieve may derive only by attacking the honest voter. Namely, by pointing the honest voter to audit in a BB location where the audit data is inconsistent with the respective information in at least one part of the voter's ballot. As in

---

[3] By meaningful we mean that the attack is not trivially detected. For example, the adversary may post malformed information in the BB nodes but if so, it will certainly fail at verification.

Voter Privacy Game $G_{\text{priv}}^{\mathcal{A},\mathcal{S},\phi}(1^\lambda, n, m, N_v, N_b, N_t)$

(i). $\mathcal{A}$ on input $1^\lambda, n, m, N_v, N_b, N_t$, chooses a list of options $\mathcal{P} = \{P_1, \ldots, P_m\}$, a set of voters $\mathcal{V} = \{V_1, \ldots, V_n\}$, a set of trustees $\mathcal{T} = \{T_1, \ldots, V_{N_t}\}$, a set of VC nodes $\{\mathsf{VC}_1, \ldots, \mathsf{VC}_{N_v}\}$ a set of BB nodes $\{\mathsf{BB}_1, \ldots, \mathsf{BB}_{N_b}\}$. It provides Ch with all the above sets. Throughout the game, $\mathcal{A}$ corrupts all the VC nodes a fixed set of $f_b < N_b/3$ BB nodes and a fixed set of $f_t < N_t/3$ trustees. On the other hand, Ch plays the role of the EA and all the non-corrupted nodes.

(ii). Ch engages with $\mathcal{A}$ in an election preparation interaction following the *Election Authority* protocol.

(iii). Ch chooses a bit value $b \in \{0, 1\}$.

(iv). The adversary $\mathcal{A}$ and the challenger Ch engage in an interaction where $\mathcal{A}$ schedules the voters which may run concurrently. For each voter $V_\ell \in \mathcal{V}$, the adversary chooses whether V_ℓ is corrupted:

- If V_ℓ is corrupted, then Ch provides the credential s_ℓ to $\mathcal{A}$, who will play the role of V_ℓ to cast the ballot.
- If V_ℓ is not corrupted, then $\mathcal{A}$ provides two option selections $\langle \text{option}_\ell^0, \text{option}_\ell^1 \rangle$ to the challenger Ch which operates on V_ℓ's behalf, voting for option option_ℓ^b. The adversary $\mathcal{A}$ is allowed to observe the network trace. After a ballot cast, the challenger Ch provides to $\mathcal{A}$: (a) the audit information α_ℓ that V_ℓ obtains from the protocol, and (b) <u>if $b = 0$</u>, the current view of the internal state of the voter V_ℓ, $view_\ell$, that the challenger obtains during voting, or <u>if $b = 1$</u>, a simulated view of the internal state of V_ℓ produced by $\mathcal{S}(view_\ell)$.

(v). The adversary $\mathcal{A}$ and the challenger Ch produce the election tally, running the *Trustee* protocol. $\mathcal{A}$ is allowed to observe the network trace of that protocol.

(vi). Finally, $\mathcal{A}$ using all information collected above (including the contents of the BB) outputs a bit b^*.

Denote the set of corrupted voters as $\mathcal{V}_{\text{corr}}$ and the set of honest voters as $\tilde{\mathcal{V}} = \mathcal{V} \setminus \mathcal{V}_{\text{corr}}$. The game returns a bit which is 1 if and only if the following hold true:

(1) $b = b^*$ (i.e., the adversary guesses b correctly).

(2) $|\mathcal{V}_{\text{corr}}| \leq \phi$ (i.e., the number of corrupted voters is bounded by ϕ).

(3) $f(\langle \text{option}_\ell^0 \rangle_{V_\ell \in \tilde{\mathcal{V}}}) = f(\langle \text{option}_\ell^1 \rangle_{V_\ell \in \tilde{\mathcal{V}}})$ (i.e., the election result w.r.t. the set of voters in $\tilde{\mathcal{V}}$ does not leak b).

Fig. A.4 – The Voter privacy Game between the adversary $\mathcal{A}$ and the challenger Ch using the simuator $\mathcal{S}$.

Kiayias et al. (2015, Theorem 4), we can show that every such single attack has 1/2 success probability (the voter had chosen to vote with the inconsistent ballot part) and in case of success, adds 1 to the tally deviation. Thus, in this case, the probability that $\mathcal{A}$ causes tally deviation d is no more than 2^{-d}.

(ii) If there is an invalid option-encoding commitment ($\mathcal{E}$ aborts), then the min entropy provided by at least θ honest successful voters is at least θ. Thus, by Lemma 2, the Sigma protocol verification will fail except from some soundness error $2^{-\theta}$. The proof is completed by taking the union bound on the two cases. □

A.4. Voter privacy

Our privacy definition extends the one used in Kiayias et al. (2015) (which is referred there as Voterprivacy/receipt-freeness) to the distributed setting of D-DEMOS. Similarly,

voter privacy is defined via a *Voter Privacy* indistinguishability game as depicted in Fig. A.4. Note that, our system achieves computational weak unlinkability among the privacy classes modeled by Bohli and Pashalidis (2011).

Definition 3 (*Voter privacy*). Let $0 < \epsilon < 1$ and $n, m, N_v, N_b, N_t \in \mathbb{N}$. Let Π be an e-voting system with n voters, m options awith n voters, N_v VC nodes, N_b BB nodes and N_t trustees w.r.t. the election function f. We say that Π achieves *voter privacy* with error ϵ for at most ϕ corrupted voters, if there is a PPT voter simulator $\mathcal{S}$ such that for any PPT adversary $\mathcal{A}$:

$$\left| \Pr[G_{\text{priv}}^{\mathcal{A},\mathcal{S},\phi}(1^\lambda, n, m, N_v, N_b, N_t) = 1] - 1/2 \right| = \text{negl}(\lambda).$$

In the following theorem, we prove that D-DEMOS (both the IC and the Async version) achieves voter privacy according to Definition 3.

Theorem 6 (*Voter privacy of D-DEMOS*). *Assume there is a constant* $c \in (0, 1)$ *such that for any* 2^{λ^c}-*time adversary* $\mathcal{A}$, *the advantage of breaking the hiding property of the underlying commitment scheme is* $\mathsf{Adv}_{\mathsf{hide}}(\mathcal{A}) = \mathsf{negl}(\lambda)$. *Let* $c' < c$ *be a constant and set* $\phi = \lambda^{c'}$. *Then, D-DEMOS run with n voters, m options,* N_v *VC nodes,* N_b *BB nodes and* N_t *trustees achieves voter privacy for at most* ϕ *corrupted voters.*

Proof. To prove voter privacy, we explicitly construct a simulator $\mathcal{S}$ such that we can convert any adversary $\mathcal{A}$ who can win the privacy game $G_{\mathsf{priv}}^{\mathcal{A},\mathcal{S},\phi}(1^\lambda, n, m, N_v, N_b, N_t)$ with a non-negligible probability into an adversary $\mathcal{B}$ who can break the hiding assumption of the underlying commitment scheme within $poly(\lambda) \cdot 2^{\lambda^{c'}} << 2^{\lambda^c}$ time.

Note that the challenger Ch is maintaining a coin $b \in \{0, 1\}$ and always uses the option option_ℓ^b to cast the honest voters' ballots. When $n - \phi < 2$, the simulator $\mathcal{S}$ simply outputs the real voters' views. When $n - \phi \geq 2$, consider the following simulator $\mathcal{S}$: At the beginning of the experiment, $\mathcal{S}$ flips a coin $b' \leftarrow \{0, 1\}$. Then, for each honest voter V_ℓ, $\mathcal{S}$ switches the vote codes for option option_ℓ^b and $\mathsf{option}_\ell^{b'}$.

Due to full VC corruption, $\mathcal{A}$ learns all the vote codes. However, it does not help the adversary to distinguish the simulated view from real view as the simulator only permutes vote codes. We now can show that if $\mathcal{A}$ can win $G_{\mathsf{priv}}^{\mathcal{A},\mathcal{S},\phi}(1^\lambda, n, m, N_v, N_b, N_t)$, then we can construct an adversary $\mathcal{B}$ that invokes $\mathcal{A}$ to win the IND-CPA game of the underlying ElGamal encryption. In the IND-CPA game, $\mathcal{B}$ receives as input a public key pk and executes the following steps:

1. It submits challenge messages $M_0 = 0, M_1 = 1$ and receives challenge ciphertext $C = \mathsf{Enc}_{\mathsf{pk}}(M_{b^*})$, where b^* is the IND-CPA challenge bit for $\mathcal{B}$.
2. It invokes $\mathcal{A}$ and simulates $G_{\mathsf{priv}}^{\mathcal{A},\mathcal{S},\phi}(1^\lambda, n, m, N_v, N_b, N_t)$, itself being the challenger.
3. $\mathcal{B}$ flips a coin $b \in \{0, 1\}$ and uses the received public key pk as the election commitment key.
4. At the beginning, $\mathcal{B}$ generates/guesses all the voters coins, $\mathbf{c} = (c_1, c_2, \ldots, c_n)$, and uses the coin c_ℓ for all the uncorrupted voter V_ℓ; if some corrupted voters' coins do not match the guessed ones, start over again. This requires 2^ϕ expected attempts to guess all the coins correctly.
5. $\mathcal{B}$ guesses the election tally $T = (T_1, T_2, \ldots, T_m)$, and starts over again if the guess is incorrect. This requires less than $(n + 1)^m$ expected attempts.
6. $\mathcal{B}$ simulates all the zero knowledge proofs using the guessed voters' coins.
7. $\mathcal{B}$ guesses/chooses an uncorrupted voter $V_{\ell'}$; the option encoding commitment of $V_{\ell'}$'s ballot for the ith option is set as:
   $(\mathsf{Enc}_{\mathsf{pk}}(T_1) \cdot C^{-T_1}, \ldots, \mathsf{Enc}_{\mathsf{pk}}(T_{i-1}) \cdot C^{-T_{i-1}}, \underline{\mathsf{Enc}_{\mathsf{pk}}(T_i) \cdot C^{-(T_i-1)}},$
   $\mathsf{Enc}_{\mathsf{pk}}(T_{i+1}) \cdot C^{-T_{i+1}}, \ldots, \mathsf{Enc}_{\mathsf{pk}}(T_m) \cdot C^{-T_m}).$
   For the rest of the voters, it commits the ith option as:
   $(\mathsf{Enc}_{\mathsf{pk}}(0), \ldots, \underline{C \cdot \mathsf{Enc}_{\mathsf{pk}}(0)}, \ldots, \mathsf{Enc}_{\mathsf{pk}}(0)).$
8. If V_ℓ is corrupted, then $\mathcal{B}$ provides the credential s_ℓ to $\mathcal{A}$.
9. If V_ℓ is not corrupted, then $\mathcal{B}$ receives two option selections $\langle \mathsf{option}_\ell^0, \mathsf{option}_\ell^1 \rangle$ from $\mathcal{A}$. It then casts the vote by submitting the vote code corresponding to option_ℓ^b.
10. $\mathcal{B}$ finishes the election according to the protocol and returns $b^* = 1$ if $\mathcal{A}$ guesses b correctly.

Note that if C encrypts 1, the commitments on the BB are the same as the ones in a real election; whereas, if C encrypts 0, the commitments of all the voters are commitments of 0's except one honest voter's commitment is the tally results. In the latter case, the adversary $\mathcal{A}$'s winning probability is exactly 1/2. Since the zero knowledge proofs are perfectly simulatable, it is easy to see that the advantage of $\mathcal{B}$ is the same as the advantage of $\mathcal{A}$. Moreover, the running time of $\mathcal{B}$ is $poly(\lambda) \cdot (n + 1)^m \cdot 2^\phi = O(2^{\lambda^{c'}})$ steps. By exploiting the distinguishing advantage of $\mathcal{A}$, $\mathcal{B}$ can break the hiding property of the option-encoding commitment scheme in $O(2^{\lambda^{c'}}) = o(2^{\lambda^c})$ steps, thus leading to contradiction. $\square$

REFERENCES

Adida B. Helios: web-based open-audit voting. In: Proceedings of the seventh USENIX security symposium, July 28-August 1, 2008, San Jose, CA, USA; 2008. p. 335–48.

Aublin P, Mokhtar SB, Quéma V. RBFT: redundant byzantine fault tolerance. In: IEEE thirty-third international conference on distributed computing systems, ICDCS 2013, 8-11 July, 2013, Philadelphia, Pennsylvania, USA; 2013. p. 297–306.

Aublin PL, Guerraoui R, Knežević N, Quéma V, Vukolić M. The next 700 bft protocols. ACM Trans Comput Syst (TOCS) 2015;32(4):12.

Bellare M, Rogaway P. Introduction to modern cryptography. UCSD CSE 207 Course Notes, 2005.

Benaloh J, Byrne MD, Eakin B, Kortum PT, McBurnett N, Pereira O, Stark PB, Wallach DS, Fisher G, Montoya J, Parker M, Winn M. Star-vote: a secure, transparent, auditable, and reliable voting system. Proceedings of the 2013 electronic voting technology workshop/workshop on trustworthy elections, EVT/WOTE '13, Washington, D.C., USA, August 12-13, 2013, 2013.

Bohli JM, Pashalidis A. Relations among privacy notions. ACM Trans Inf Syst Secur (TISSEC) 2011;14(1) 4:1–4:24. doi:10.1145/1952982.1952986.

Castro M, Liskov B. Practical byzantine fault tolerance and proactive recovery. ACM Trans Comput Syst (TOCS) 2002;20(4):398–461.

Chaum D. Surevote: tchnical overview. Proceedings of the workshop on trustworthy elections, 2001. WOTE

Chaum D, Essex A, Carback R, Clark J, Popoveniuc S, Sherman AT, Vora PL. Scantegrity: end-to-end voter-verifiable optical-scan voting. IEEE Secur Priv 2008;6(3):40–6.

Chaum D, Pedersen TP. Wallet databases with observers. In: CRYPTO '92. Springer-Verlag; 1993. p. 89–105.

Chaum D, Ryan PYA, Schneider SA. A practical voter-verifiable election scheme. In: ESORICS 2005; 2005. p. 118–39.

Chondros N, Zhang B, Zacharias T, Diamantopoulos P, Maneas S, Patsonakis C, Delis A, Kiayias A, Roussopoulos M. D-DEMOS: a distributed, end-to-end verifiable, internet voting system. In: Proceedings of the thirty-sixth IEEE international conference on distributed computing systems, ICDCS 2016, Nara, Japan, June 27-30, 2016; 2016. p. 711–20.

Clarkson MR, Chong S, Myers AC. Civitas: Toward a secure voting system. In: Proceedings of the IEEE symposium on security and privacy (S&P 2008), Oakland, California, USA; 2008. p. 354–68.

Clement A, Kapritsos M, Lee S, Wang Y, Alvisi L, Dahlin M, Riche T. Upright cluster services. In: Proceedings of the Twenty-second ACM symposium on operating systems principles 2009, SOSP 2009, Big Sky, Montana, USA, October 11-14, 2009; 2009a. p. 277–90.

Clement A, Wong EL, Alvisi L, Dahlin M, Marchetti M. Making byzantine fault tolerant systems tolerate byzantine faults, 9; 2009b. p. 153–68.

Community N. Netty, an asynchronous event-driven network application framework, 2015a. http://netty.io/.

Community P. Postgresql rdbms, 2015b. http://www.postgresql.org/.

Cowling JA, Myers DS, Liskov B, Rodrigues R, Shrira L. HQ replication: A hybrid quorum protocol for byzantine fault tolerance. In: Proceedings of the seventh symposium on operating systems design and implementation (OSDI '06), November 6-8, Seattle, WA, USA; 2006. p. 177–90.

Cramer R, Damgård I, Schoenmakers B. Proofs of partial knowledge and simplified design of witness hiding protocols. In: CRYPTO '94. Springer Berlin Heidelberg; 1994. p. 174–87.

Cramer R, Gennaro R, Schoenmakers B. A secure and optimally efficient multi-authority election scheme. In: Proceedings of the advances in cryptology – EUROCRYPT '97, international conference on the theory and application of cryptographic techniques, Konstanz, Germany, May 11-15, 1997; 1997. p. 103–18.

Culnane C, Ryan PYA, Schneider S, Teague V. vvote: a verifiable voting system. ACM Trans Inf Syst Secur (TISSEC) 2015;18(1) 3:1–3:30. doi:10.1145/2746338.

Culnane C, Schneider S. A peered bulletin board for robust use in verifiable voting systems. In: Proceedings of the 2014 IEEE twenty-seventh computer security foundations symposium (CSF). IEEE; 2014. p. 169–83.

Diamantopoulos P, Maneas S, Patsonakis C, Chondros N, Roussopoulos M. Interactive consistency in practical, mostly-asynchronous systems. In: Proceedings of the 2015 IEEE twenty-first international conference on parallel and distributed systems (ICPADS); 2015. p. 752–9.

Dini G. A secure and available electronic voting service for a large-scale distributed system. Future Gener Comput Syst 2003;19(1):69–85.

El Gamal T. A public key cryptosystem and a signature scheme based on discrete logarithms. In: CRYPTO. Springer-Verlag; 1985. p. 10–18.

ElGamal T. A public key cryptosystem and a signature scheme based on discrete logarithms. IEEE Trans Inf Theory 1985;31(4):469–72.

Feige U, Fiat A, Shamir A. Zero-knowledge proofs of identity. J Cryptol 1988;1(2):77–94.

Fisher K, Carback R, Sherman A. Punchscan: introduction and system definition of a high-integrity election system. WOTE, 2006.

Gjøsteen K. The norwegian internet voting protocol. IACR Cryptol ePrint Arch 2013;2013:473.

Grigorik I. High performance browser networking: what every web developer should know about networking and web performance, 2013. http://chimera.labs.oreilly.com/books/1230000000545/ch01.html#PROPAGATION_LATENCY.

Gritzalis DA. Principles and requirements for a secure e-voting system. Comput Secur 2002;21(6):539–56. doi:10.1016/S0167-4048(02)01014-3.

Hemminger S, et al. Network emulation with netem. In: Linux conf au. Citeseer; 2005. p. 18–23.

Huang Y, Geng X, Whinston AB. Defeating DDoS attacks by fixing the incentive chain. ACM Trans Internet Technol 2007;7(1).

Inc. G. Google protocol buffers, 2015. https://code.google.com/p/protobuf/.

Internet Policy Institue. Report of the national workshop on internet voting: issues and research agenda, 2001.

Joaquim R, Ferreira P, Ribeiro C. Eviv: an end-to-end verifiable internet voting system. Comput Secur 2013;32:170–91. doi:10.1016/j.cose.2012.10.001.

Kiayias A, Zacharias T, Zhang B. End-to-end verifiable elections in the standard model. In: EUROCRYPT 2015; 2015. p. 468–98.

Kotla R, Alvisi L, Dahlin M, Clement A, Wong EL. Zyzzyva: speculative byzantine fault tolerance. In: Proceedings of the twenty-first acm symposium on operating systems principles 2007, SOSP 2007, Stevenson, Washington, USA, October 14-17, 2007; 2007. p. 45–58.

Kulyk O, Neumann S, Marky K, Budurushi J, Volkamer M. Coercion-resistant proxy voting. Comput Secur 2017;71:88–99. doi:10.1016/j.cose.2017.06.007.

Kumaraguru P, Sheng S, Acquisti A, Cranor LF, Hong J. Teaching johnny not to fall for phish. ACM Trans Internet Technol 2010;10(2) 7:1–7:31. doi:10.1145/1754393.1754396.

Kutylowski M, Zagórski F. Scratch, click & vote: E2E voting over the internet. Lecture Notes in Computer Science, 6000. Springer; 2010. p. 343–56.

Laisi A. Asynchronous postgresql java driver, 2015. https://github.com/alaisi/postgres-async-driver/.

Lynch N. Distributed algorithms. Morgan Kaufmann; 1996.

Maneas S. Implementation and evaluation of a distributed, end-to-end verifiable, internet voting system. University of Athens; 2015. (Msc.).

MIRACL. Miracl multi-precision integer and rational arithmetic c/c++ library, 2015. http://www.certivox.com/miracl/.

Moran T, Naor M. Split-ballot voting: everlasting privacy with distributed trust. ACM Trans Inf Syst Secur (TISSEC) 2010;13(2) 16:1–16:43. doi:10.1145/1698750.1698756.

Neumann P. Security criteria for electronic voting. In: Proceedings of the national computer security conference; 1993. p. 478–81.

Pease M, Shostak R, Lamport L. Reaching agreement in the presence of faults. J ACM (JACM) 1980;27(2):228–34.

Pedersen T. Non-interactive and information-theoretic secure verifiable secret sharing. Proceedings of the advances in cryptology CRYPTO, 1991.

Quisquater JJ, Quisquater M, Quisquater M, Quisquater M, Guillou L, Guillou M, Guillou G, Guillou A, Guillou G, Guillou S. How to explain zero-knowledge protocols to your children. In: Proceedings of the advances in cryptology-CRYPTO89. Springer; 1990. p. 628–31.

Schneier B. Applied cryptography. John Wiley & Sons; 1996.

Tiemens T. Shamir's secret share in java, 2015. https://github.com/timtiemens/secretshare.

Yue C, Wang H. Bogusbiter: A transparent protection against phishing attacks. ACM Trans Internet Technol 2010;10(2) 6:1–6:31. doi:10.1145/1754393.1754395.

Zagórski F, Carback RT, Chaum D, Clark J, Essex A, Vora PL. Remotegrity: design and use of an end-to-end verifiable remote voting system. In: Proceedings of the eleventh international conference on applied cryptography and network security, ACNS 2013, Banff, AB, Canada, June 25-28, 2013; 2013. p. 441–57.

Nikos Chondros is currently a Postdoc researcher in the Distributed Systems group, at the Department of Informatics and Telecommunications, University of Athens, Greece. His current research interests include distributed systems, scalable systems, and reliable systems.

Bingsheng Zhang is currently a lecturer of Computer Science at the School of Computing and Telecommunications, Lancaster University, UK. His research focus is in the areas of IoT security and cryptography. In particular, he is interested in the following subjects: E-voting Security, Blockchain Security, IoT Security.

Thomas Zacharias is currently a Postdoc researcher in the Laboratory for Foundations of Computer Science, at the School of Informatics, University of Edinburgh, UK. His research interests include voting systems and secure communication

Panos Diamantopoulos is currently a Ph.D. candidate in the Distributed Systems group, in the Department of Informatics and Telecommunications, University of Athens, Greece. His current research interests include distributed systems and IoT applications.

Stathis Maneas is currently a Ph.D. candidate at the Department of Computer Science, University of Toronto, Canada and a member of the Computer Systems and Networks group. His current main research interests include the design and implementation of computer systems and storage systems, while focusing on system reliability.

Christos Patsonakis is currently a Ph.D. candidate in the Distributed Systems group, in the Department of Informatics and Telecommunications, University of Athens, Greece. His current research interests include distributed systems and block-chain protocols.

Alex Delis is a professor of Computer Science at the Department of Informatics and Telecommunication, University of Athens, Greece. His research interests include: Distributed and Virtualized Systems, Management of Big Data and Software Systems.

Aggelos Kiayias is chair in Cyber Security and Privacy and director of the Blockchain Technology Laboratory at the University of Edinburgh. He is also the Chief Scientist at blockchain technology company IOHK. His research interests are in computer security, information security, applied cryptography and foundations of cryptography with a particular emphasis in blockchain technologies and distributed systems, e-voting and secure multiparty protocols as well as privacy and identity management.

Mema Roussopoulos is an Associate Professor of Computer Science at the University of Athens. She received her Ph.D. in Computer Science from Stanford University. Her interests are in the areas of distributed systems and networking. She is a recipient of the NSF CAREER Award (while on the faculty at Harvard University), the ERC Starting Grant Award, and the Best Paper Award at ACM SOSP 2003.