DRE-ip: A Verifiable E-Voting Scheme without Tallying Authorities[⋆]

Siamak F. Shahandashti and Feng Hao

School of Computing Science, Newcastle University, UK
{siamak.shahandashti,feng.hao}@ncl.ac.uk

Abstract. Nearly all verifiable e-voting schemes require trusted tallying authorities to guarantee voter privacy. An exception is the DRE-i system which removes this requirement by pre-computing all encrypted ballots before the election using related random factors that will later cancel out and allow the public to verify the tally after the election. While the removal of tallying authorities significantly simplifies election management, the pre-computation of ballots necessitates secure ballot storage, as leakage of precomputed ballots endangers voter privacy. In this paper, we address this problem and propose DRE-ip (DRE-i with enhanced privacy). Adopting a different design strategy, DRE-ip is able to encrypt ballots in real time in such a way that the election tally can be publicly verified without decrypting the cast ballots. As a result, DRE-ip achieves end-to-end verifiability without tallying authorities, similar to DRE-i, but with a significantly stronger guarantee on voter privacy. In the event that the voting machine is fully compromised, the assurance on tallying integrity remains intact and the information leakage is limited to the minimum: only the partial tally at the time of compromise is leaked.

1 Introduction

Direct-recording electronic (DRE) machines have been extensively used for voting at polling stations around the world. In a typical process, a registered voter obtains a token after being authenticated at the polling station. She then enters a private booth and presents the token to a DRE machine. The token is for one-time use and allows the voter to cast only one vote. Usually, the DRE machine has a touch screen to record the vote directly from the voter (hence the name DRE). The machine may tally the votes in real time, or store the votes and tally later. In either case, the machine works like a black box: if an attacker maliciously changes the votes (or the tally thereof), this is likely to go unnoticed.

Lack of assurance on tallying integrity is commonly regarded as a critical weakness of such DRE machines. To address this problem, several cryptographic protocols are proposed in the literature. The seminal work by Chaum in 2004 [16] involves using visual cryptography to allow voters to verify the integrity of an election. The assurance on the integrity includes guarantees that the votes are

[⋆] This is the full version of a paper by the same title to appear in ESORICS 2016.

cast as intended, recorded as cast, and tallied as recorded. The fulfilment of all three constitutes the widely-accepted notion of end-to-end (E2E) verifiability.

Chaum's solution inspired a class of voting systems providing E2E verifiability. Prominent examples include MarkPledge [30], Prêt à Voter [31], Scantegrity [14] (and its predecessor PunchScan [21]), Helios [1], and STAR-Vote [4]. These systems are based on different voting media including physical ballots, optical scanners, DREs and web browsers. They use different tallying techniques, based on mix-nets or homomorphic encryption. But all these schemes allow individual voters to verify if their votes have been cast as intended and recorded as cast, and any observer to verify if all votes have been tallied as recorded.

In this paper we limit our attention to DRE-based elections. We focus on DRE as it has already been widely deployed for national elections worldwide. Today, nearly all of the deployed DRE systems work like a black box and offer no guarantee on integrity; consequently, their use has been abandoned in several countries such as the Netherlands, Germany and Ireland. However, in many other countries, these (unverifiable) DRE machines continue to be extensively used. We believe there is an urgent need to address this real-world problem.

Apart from Chaum's system, other existing E2E verifiable schemes for DRE-based elections include MarkPledge [30], VoteBox [33], and STAR-Vote [4]. These systems may differ significantly in details, but they share some common features. They all offer integrity assurance by introducing a set of trustworthy tallying authorities (TAs). Instead of the DRE directly recording the vote, the machine encrypts the vote on the fly under the joint public key of the TAs. Each TA is responsible for safeguarding a share of the decryption key. When the voting is closed, a quorum of TAs jointly perform the tallying process which involves decryption of the ballots (or tally thereof) in a publicly-verifiable manner.

The addition of external TAs however introduces difficulties in the implementation. In theory, the TAs should be selected from parties with conflicting interests. They should have the expertise to be able to independently manage their own key shares and perform cryptographic operations – if they delegate their key management tasks, the delegates need to be trusted as well. A comparatively high level of cryptographic and computing skills is expected from the TAs. Furthermore, the quorum should be set sufficiently large such that collusion among TAs is infeasible, but at the same time, sufficiently small such that the process is error-tolerant, since non-availability of TA keys will render the election result non-computable. Reconciling the two is not an easy task. As reported by real-world experience of building E2E verifiable voting based on Helios, the implementation of the TAs proved to be "one particularly difficult issue" [2].

Hao et al. investigated if it was possible to achieve E2E verifiability for a DRE-based election without involving any TAs [24]. They proposed a TA-free E2E voting system, called DRE-i (DRE with integrity). In DRE-i, the machine directly records the voter's choice as in the existing practice of current DRE-based elections. However, the machine is required to publish additional audit data on a public bulletin board, to enable every voter to verify the integrity of the voting process. In DRE-i, the encryption of votes is based on a variant of the ElGamal encryption scheme: instead of using a fixed public key for encryption as in classic ElGamal, DRE-i uses a dynamically constructed public key for encrypting ballots. The system removes the need for TAs by pre-computing encrypted ballots in a structured manner such that after the election, multiplication of all the published ciphertexts cancels out the random factors that were introduced during the encryption process, and permits anyone to verify the tally.

DRE-i demonstrates that the role of the TAs is not indispensable in achieving E2E verifiability in a DRE-based election. However, its pre-computation strategy inevitably introduces the requirement of ensuring that the pre-computed data is securely stored and accessed during the voting phase. Furthermore, it means that it is possible for an adversary that breaks into the secure storage module to potentially compromise the privacy of all ballots. The authors of DRE-i [24] suggest to use tamper-resistant hardware to protect the pre-computed data in sensitive elections. However, the use of tamper-resistant hardware may significantly drive up the cost for each DRE machine. Furthermore, designing secure API for tamper-resistant hardware is a challenging problem on its own.

It remains an open problem as whether it is possible to achieve the best of both worlds, i.e., strong assurance on the integrity of a DRE-based election without involving any TAs, and simultaneously, a strong guarantee on the privacy of votes without depending on tamper-resistant hardware.

In this paper, we provide a positive answer to this question and present a new E2E voting system, which we call *DRE-ip* (DRE-i with enhanced privacy). Instead of pre-computing ciphertexts, DRE-ip adopts a more conventional approach, as in other existing DRE-based verifiable systems, to encrypt the vote on the fly during voting. DRE-ip achieves E2E verifiability without TAs, but at the same time provides a significantly stronger privacy guarantee than DRE-i.

Our Contributions. We present DRE-ip, an end-to-end verifiable DRE-based voting system that encrypts ballots in real-time, but requires no TAs to decrypt ballots in the tallying phase. We consider two types of attacks, which we call non-intrusive and intrusive, based on whether the adversary can compromise the DRE machine or not. We prove that DRE-ip provides indistinguishability of elections with the same tally against non-intrusive attacks based on the decision Diffie-Hellman assumption. In the event of an intrusive attack, we prove that only the privacy of the ballots cast during the attack period is lost – a loss which is inevitable – and the ballots cast outside the attack period are guaranteed to remain private under the Square Diffie-Hellman assumption. Thus, DRE-ip constitutes the first verifiable DRE-based system that removes the need for tallying authorities without introducing new assumptions.

Related Work. In his seminal work on anonymous communications, Chaum put forward e-voting as an application of his technique [15]. This prompted considerable research on e-voting, among which is the work of Benaloh [10] that proposed a formal definition of *ballot secrecy.* Later, Benaloh and Tuinstra argued for *receipt-freeness* [9], and Juels, Catalano, and Jakobsson put forward *coercion-resistance* [25] as progressively stronger notions of privacy. On the other hand, verifiability has evolved as a property guaranteeing the integrity of e-voting systems. Earlier works considered *individual verifiability.* The notion of *universal verifiability* emerged in later works and Sako and Kilian explicitly formalized it [32]. Finally, through the works of Chaum [16] and Neff [30], the notions of verifiability were refined into the now widely-accepted notion of *end-to-end verifiability*, which includes guarantees that the votes are cast as intended, recorded as cast, and tallied as recorded. End-to-end verifiability has become a de facto standard for any e-voting scheme. Accordingly, in this paper, we limit our attention to end-to-end verifiable voting schemes.

There has been a renewed interest in academic research on e-voting in the past fifteen years and a number of end-to-end verifiable schemes have been designed and used in practice. Among the more influential schemes are Votegrity, proposed by Chaum [16], and MarkPledge, proposed by Neff [30], which are the first end-to-end verifiable schemes. Many other schemes follow similar approaches, including Prêt à Voter [31], a tailored variant of which has been recently used in state elections in Victoria, Australia [18], Scantegrity [14], which was trialled in local elections in Takoma Park, Maryland, USA [13], and STAR-Vote [4], which is scheduled for deployment in elections in Travis County, Texas, USA [27]. Other schemes that have been used in internal university or party elections include PunchScan [21], Bingo Voting [11], Helios [1], Wombat [7], and DRE-i [24].

2 Preliminaries

In this section, we review the preliminaries required for the description of DRE-ip, including the notation, the cryptographic setting, and the DRE-i system.

Notation. Following the notation introduced by Camenisch and Stadler [12], we use $\mathrm{P_K}\{\lambda : \mathit{\Gamma} = \gamma^\lambda\}$ to denote a non-interactive *proof of knowledge* of (a secret) λ such that (for publicly-known $\mathit{\Gamma}$ and γ): $\mathit{\Gamma} = \gamma^\lambda$. Where the context is clear, we shorten the notation to $\mathrm{P_K}\{\lambda\}$. We use $\mathrm{P_{WF}}\{A : X, Y, Z\}$ to denote a *proof of well-formedness* of A with respect to X, Y, and Z. Where the context is clear, we shorten the notation to $\mathrm{P_{WF}}\{A\}$.

2.1 Cryptographic Setting

We assume a DSA-like multiplicative cyclic group setting, where p and q are large primes that satisfy $q \mid p-1$. We work in the subgroup $\mathbb{G}_q$ of order q of the group $\mathbb{Z}_p^\star$ and assume that g is a generator of $\mathbb{G}_q$. Alternatively, our proposed system can be implemented over an elliptic curve in an ECDSA-like group setting.

The decision Diffie-Hellman (DDH) assumption [19] is defined as follows:

Assumption 1. (DDH) *For randomly chosen $a, b \in \mathbb{Z}_q^\star$ and $R \in \mathbb{G}_q$, given $(g, g^a, g^b, \mathit{\Omega})$ where $\mathit{\Omega} \in \{g^{ab}, R\}$, it is hard to decide whether $\mathit{\Omega} = g^{ab}$ or $\mathit{\Omega} = R$.*

The Square DDH assumption [28] is defined as follows:

Assumption 2. (Square DDH) For randomly chosen $a \in \mathbb{Z}_q^\star$ and $R \in \mathbb{G}_q$, given (g, g^a, Ω) where $\Omega \in \{g^{a^2}, R\}$, it is hard to decide whether $\Omega = g^{a^2}$ or $\Omega = R$.

Clearly, if one can break DDH, then Square DDH can be broken as well. Hence, Square DDH is a stronger assumption than DDH. Furthermore, there is evidence that Square DDH is strictly stronger [35, 26].

Zero knowledge proofs, first proposed by Goldwasser, Micali, and Rackoff [22], prove the truth of a statement without conveying any other information, i.e., they guarantee that whatever the verifier can feasibly compute after seeing a proof, they could have computed on their own. Subsequent work by Bellare and Goldreich [5] refined the definition of zero knowledge proofs to distinguish them from proofs of knowledge. Intuitively speaking, proofs of knowledge are guaranteed to be generated by a prover with explicit knowledge of a quantity. In our protocol, the Fiat-Shamir heuristic is employed to construct non-interactive proofs [20]. Consequently, our security proofs are in the Random Oracle Model [6].

3 Our Proposed Solution: DRE-ip

DRE-ip requires a secure and publicly-accessible bulletin board (BB) and incorporates voter-initiated auditing to achieve end-to-end verifiability. We assume the DRE has append-only write access to the BB over an authenticated channel. We assume voting is conducted in supervised polling stations and there are procedures in place to ensure the "one person, one vote" principle, including secure voter registration and authentication. At the time of voting, a voter is authenticated first and issued a token, unlinked to her identity. She then enters a private voting booth and authenticates herself to the DRE using the token. Up to here, the assumptions and mechanisms are similar to those of DRE-i.

We describe DRE-ip for the case where there are only two candidates, i.e., for v_i representing the vote of the i-th ballot, we have $v_i \in \{0, 1\}$. In DRE-ip the setup establishes two generators g and $\tilde{g}$, whose logarithmic relationship is unknown. The DRE keeps track of the running tally $t = \sum v_i$ for the cast votes v_i, and the sum $s = \sum x_i y_i$ for random x_i and y_i generated on the fly.

To achieve individual verifiability, DRE-ip incorporates Benaloh-style voter-initiated auditing [8], i.e., the voter gets the option to audit the ballot composed by the DRE to gain confidence in that the DRE is preparing the ballots according to her choice. If a ballot is audited, it cannot be used to cast a vote. Therefore, the set of all ballots $\mathbb{B}$ at the closing of the voting phase will be comprised of the audited ballots $\mathbb{A}$ and the cast ballots $\mathbb{C}$, i.e., $\mathbb{B} = \mathbb{A} \cup \mathbb{C}$.

Voting Phase. This phase involves the voter, the DRE, and the BB:

1. The voter enters the booth, initiates voting, and keys in her vote $v_i \in \{0, 1\}$.
2. The DRE generates random $x_i, y_i \in \mathbb{Z}_q^\star$, calculates

$$\begin{array}{llll} X_i = g^{x_i}, & Y_i = g^{y_i}, & \tilde{X}_i = \tilde{g}^{x_i}, & \mathrm{P}_{\text{WF}}\{\tilde{X}_i : g, X_i, \tilde{g}\}, \\ Z_i = g^{x_i y_i} g^{v_i}, & \mathrm{P}_{\text{WF}}\{Z_i : g, X_i, Y_i\}, & \tilde{Z}_i = \tilde{g}^{x_i y_i}, & \mathrm{P}_{\text{WF}}\{\tilde{Z}_i : g, Y_i, \tilde{X}_i\}, \end{array}$$

and provides a signed receipt including the unique ballot index i and the ballot content X_i, Y_i $\tilde{X}_i$, $\mathrm{P}_{\mathrm{WF}}\{\tilde{X}_i\}$, Z_i, $\mathrm{P}_{\mathrm{WF}}\{Z_i\}$, $\tilde{Z}_i$, and $\mathrm{P}_{\mathrm{WF}}\{\tilde{Z}_i\}$ to the voter.

3. The voter observes that the first part of the receipt is provided, and chooses to either audit the ballot or confirm her vote.

In case of audit:

4. The DRE adds i to $\mathbb{A}$, provides a signed receipt of audit, clearly marked `audited`, including x_i, y_i, and v_i to the voter.
5. The voter takes and keeps the receipt, and verifies that v_i reflects her choice. If the verification succeeds, voting continues to Step 1; otherwise, the voter should raise a dispute immediately.

In case of confirmation:

4. The DRE adds i to $\mathbb{C}$, updates the tally and the sum

$$t = \sum_{j \in \mathbb{C}} v_j, \quad s = \sum_{j \in \mathbb{C}} x_j y_j,$$

and provides a signed receipt of confirmation, clearly marked `confirmed`, to the voter, and securely deletes x_i, y_i, and v_i.

5. The voter leaves the booth with her receipts.

6. The DRE posts on the BB all the receipts provided to the voter.
7. The voter verifies that her receipts match those on the BB.

Tallying Phase. This phase involves the DRE, the BB, and the public:

1. The DRE calculates

$$S = g^s, \quad \mathrm{P}_{\mathrm{WF}}\{S : g, \tilde{g}, \prod_{j \in \mathbb{C}} \tilde{Z}_j\},$$

and posts on the BB the final tally t, as well as S and $\mathrm{P}_{\mathrm{WF}}\{S\}$.

2. The public:
   - verify all the well-formedness proofs on the BB (*well-formedness verification*),
   - verify that for all the audited ballots on the BB: X_i, Y_i, $\tilde{X}_i$, Z_i, and $\tilde{Z}_i$ included in the first part of the receipt are consistent with x_i, y_i, and v_i included in the second part (and with the system parameters g and $\tilde{g}$) (*audit consistency verification*), and
   - verify that the following equation holds (*tally verification*):

$$\prod_{j \in \mathbb{C}} Z_j = S g^t. \tag{1}$$

If at any point during the voting or tallying phases, any of the verifications carried out by the voter or the public does not succeed, the election staff should be notified and we assume that there are procedures in place dealing with such verification failures. These include voter verifications in Steps 5 (in case of audit) and 7 of the voting phase and public verifications in Step 2 of the tallying phase.

Figure 1 shows the DRE-ip bulletin board. An audited receipt (with index i) and a confirmed receipt (with index j) are shown. Each receipt has two parts: the first part is provided to the voter before she decides to either audit or confirm her ballot and includes the same information for all receipts; the second part is provided after the voter makes her decision and includes different information based on her choice. Both parts of the receipt are signed by the DRE.

The proofs of well-formedness are realized as follows. $\mathrm{P}_{\mathrm{WF}}\{\tilde{X}_i : g, X_i, \tilde{g}\}$, $\mathrm{P}_{\mathrm{WF}}\{\tilde{Z}_i : g, Y_i, \tilde{X}_i\}$, and $\mathrm{P}_{\mathrm{WF}}\{S : g, \tilde{g}, \prod_{j\in\mathbb{C}} \tilde{Z}_j\}$ are all realized as proofs of knowledge and equality of two discrete logarithms as follows:

$$\begin{aligned}
\mathrm{P}_{\mathrm{WF}}\{\tilde{X}_i\} &= \mathrm{P}_{\mathrm{K}}\{\ x_i : X_i = g^{x_i} \wedge \tilde{X}_i = \tilde{g}^{x_i}\ \},\\
\mathrm{P}_{\mathrm{WF}}\{\tilde{Z}_i\} &= \mathrm{P}_{\mathrm{K}}\{\ y_i : Y_i = g^{y_i} \wedge \tilde{Z}_i = \tilde{X}_i^{y_i}\ \},\\
\mathrm{P}_{\mathrm{WF}}\{S\} &= \mathrm{P}_{\mathrm{K}}\{\ s : S = g^{s} \wedge \prod_{j\in\mathbb{C}} \tilde{Z}_j = \tilde{g}^{s}\ \}.
\end{aligned}$$

$\mathrm{P}_{\mathrm{WF}}\{Z_i : g, X_i, Y_i\}$ is realized as a proof of knowledge

$$\mathrm{P}_{\mathrm{WF}}\{Z_i\} = \mathrm{P}_{\mathrm{K}}\{\ x_i :\ \ (X_i = g^{x_i} \wedge Z_i = Y_i^{x_i})\ \vee\ (X_i = g^{x_i} \wedge Z_i/g = Y_i^{x_i})\ \}.$$

This proof guarantees that $Z_i \in \{g^{x_i y_i}, g^{x_i y_i} g\}$, or equivalently $v_i \in \{0, 1\}$.

The well-formedness proofs are based on Schnorr proofs of knowledge of discrete logarithm [34]. Starting with a Schnorr proof, one can apply techniques proposed by Cramer, Damgård, and Schoenmakers [17] to construct proofs of disjunctive knowledge, conjunctive knowledge, and combinations of both. The Fiat-Shamir heuristic [20] is then applied to make the constructed proofs non-interactive. The index i of the ballot is embedded in the proof (as an input to the hash function) to bind the proof to the ballot.

In practice, truncated hash functions may be used to calculate a short digest, e.g., 4 alphanumeric characters long, of each part of the receipt, so that the voter can easily compare the digest on their receipts with those on the bulletin board. In this case, voters are expected to verify the receipts before leaving the polling station and we assume facilities are provided for them to do so in the station.

Although we described the system for only two candidates, there are straightforward methods to extend it to support multiple candidates (see e.g., [24, 3]).

4 Security of DRE-ip

In this section we provide proofs to show that DRE-ip is end-to-end verifiable and ensures ballot secrecy under both non-intrusive and intrusive attacks.

Initial:	$g,\ \tilde{g}$	
Receipts:	$\vdots$	
	$i: X_i,\ Y_i,\ \tilde{X}_i,\ \mathrm{P}_{\text{WF}}\{\tilde{X}_i\}, Z_i,\ \mathrm{P}_{\text{WF}}\{Z_i\},\ \tilde{Z}_i,\ \mathrm{P}_{\text{WF}}\{\tilde{Z}_i\}$	`audited`, $x_i,\ y_i,\ v_i$
	$\vdots$	
	$j: X_j,\ Y_j,\ \tilde{X}_j,\ \mathrm{P}_{\text{WF}}\{\tilde{X}_j\}, Z_j,\ \mathrm{P}_{\text{WF}}\{Z_j\},\ \tilde{Z}_j,\ \mathrm{P}_{\text{WF}}\{\tilde{Z}_j\}$	`confirmed`
	$\vdots$	
Final:	$t,\ S,\ \mathrm{P}_{\text{WF}}\{S\}$	

Fig. 1. DRE-ip bulletin board

4.1 End-to-End Verifiability

We discuss the integrity (i.e., correctness) of the election tally in DRE-ip and show how DRE-ip achieves end-to-end verifiability: we prove that votes are tallied as recorded under the assumption that all proofs of well-formedness are proofs of knowledge; furthermore, we demonstrate how voter-initiated auditing guarantees that votes are recorded as cast, and cast as intended.

We assume the bulletin board is secure, in particular it is append-only and publicly accessible. Besides, there should be a mechanism to establish an authenticated channel between authorized DRE(s) and the bulletin board, to ensure that only an authorized DRE can append new values to the BB, and also that such values are not modified in transit. This can be achieved using standard techniques such as digital signatures. Furthermore, we assume that the number of voters is less than the size of the group q.

Recall that public verification in DRE-ip, i.e., Step 2 of the tallying phase, includes three types of verification: well-formedness verification, audit consistency verification, and tally verification. The following theorem shows that if well-formedness and tally verifications succeed, DRE-ip achieves the tallied-as-recorded property, that is, DRE-ip guarantees that the tally on the bulletin board is the correct tally of all the confirmed ballots on the bulletin board.

Theorem 1. *In DRE-ip, assuming that all proofs of well-formedness are proofs of knowledge, if the public well-formedness and tally verifications succeed, then the reported tally t is the correct tally of all the confirmed votes on the BB.*

The proof is rather straightforward and hence omitted here due to lack of space. In short, one can demonstrate how the proofs of well-formedness collectively guarantee that the tally verification equation (i.e., Equation 1 on page 6) holds if and only if $t = \sum_{i \in \mathbb{C}} v_i$, where $\mathbb{C}$ denotes the set of confirmed votes. Hence, if well-formedness and tally verifications are carried out successfully, the reported tally t is guaranteed to be the correct tally of all the confirmed votes on the BB.

The well-formedness dependency graph for $\tilde{X}_i$, $\tilde{Z}_i$, and S enforced by the corresponding proofs of well-formedness is given in Figure 2. As the graph shows,

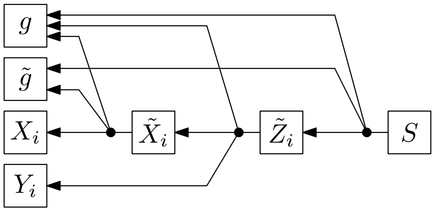

Fig. 2. Well-formedness dependency graph for $\tilde{X}_i$, $\tilde{Z}_i$, and S. The item on the right of each set of edges is well-formed with respect to the items on the left.

$\tilde{X}_i$, $\tilde{Z}_i$, and S are all eventually well-formed with respect to four values: g, $\tilde{g}$, X_i, and Y_i. Therefore, given fixed g, $\tilde{g}$, X_i, and Y_i, the well-formedness proofs guarantee that $\tilde{X}_i$, $\tilde{Z}_i$, and S are fixed.

Voter initiated auditing includes the following checks: first, by observing the first part of the receipt is provided before deciding to either audit or confirm a ballot, the voter makes sure that the DRE commits to the first part of the ballot; second, by checking that the receipts match what is published on the BB, the voter makes sure that her interaction with the machine is captured faithfully on the bulletin board. The public verification of the consistency of the audited ballots, i.e., the audit consistency verification, guarantees that DRE has been successful in responding to the challenges made by voter initiated auditing. Hence, the individual verification and the public audit consistency verification collectively ensure that the votes are cast as intended and recorded as cast. Theorem 1 ensures that votes are tallied as recorded.

4.2 Ballot Secrecy

Ballot secrecy corresponds to the natural expectation from a voting system to protect the secrecy of cast ballots. We consider a definition of ballot secrecy which requires that an adversary controlling the voting behaviour of a group of dishonest voters should not be able to distinguish between any two elections, regardless of how honest voters vote, as long as the two elections have the same sub-tally of honest votes. This definition originates from Benaloh [10, p. 74].

We assume a secure setup phase; that is, we assume that the discrete logarithm of $\tilde{g}$ in base g is either not known to any party or securely deleted after the two generators are computed. We also assume secure deletion of values x_i, y_i, and v_i after each vote is cast[1].

Ballot Secrecy under Non-Intrusive Attacks. Let us consider an adversary that does not get access to the voting machine (DRE). The adversary is able to read the publicly available information on the bulletin board, which includes

[1] See, for instance, [23] and the references within for an overview of available solutions to secure data deletion.

the total tally. Besides, we assume that the adversary can control an arbitrary number of voters and in effect cast an arbitrary number of votes. Let us call the votes cast by the adversary (or more generally known by the adversary) the *adversarial* votes. Knowledge of the adversarial votes along with the total tally enables the adversary to find out the tally of the non-adversarial votes. We prove that under the DDH assumption, this is the only information the adversary gains about the non-adversarial votes. In particular, we show that any two elections with the same non-adversarial tally are indistinguishable to the adversary.

We first consider two elections in which all votes are the same except for two votes that are swapped. We show that the bulletin boards of these two elections remain indistinguishable to the adversary even if the adversary controls all the votes other than the two that are swapped. More formally, we have:

Lemma 1. *In DRE-ip, assuming that all proofs of well-formedness are zero knowledge, if the DDH assumption holds, then an adversary that determines an arbitrary number of votes cannot distinguish between two bulletin boards in which two votes are swapped.*

The proof comes in Appendix A. We consider an adversary that can determine an arbitrary number of voter except two votes v_i and v_j. Assuming that such an adversary is able to distinguish the bulletin boards in which v_i and v_j are swapped, we show how the adversary can be used to break the DDH assumption.

Given Lemma 1, we expand it to prove that any two elections with the same tally remain indistinguishable to an adversary who controls an arbitrary number of votes. This shows that the only knowledge the adversary can gain about the non-adversarial votes is that disclosed by the election tally.

Theorem 2. *In DRE-ip, assuming that all proofs of well-formedness are zero knowledge, if the DDH assumption holds, then an adversary that determines an arbitrary number of votes cannot gain any knowledge about the non-adversarial votes other than their tally.*

Proof. To prove this theorem, we show that under the DDH assumption, given any two sets of non-adversarial votes with the same tally, one can simulate two corresponding bulletin boards that are indistinguishable to an adversary that chooses an arbitrary number of adversarial votes.

First, note that any two given sets of non-adversarial votes with the same tally differ on an even number of votes, say $2d$. This means that with d "swaps" one set of these votes can be converted to the other, where in each swap, for some i and j, the i-th vote is replaced with the j-th one, and vice versa. In Lemma 1 we proved that the bulletin boards before and after each swap remain indistinguishable to the adversary under DDH. Consequently, the bulletin boards corresponding to the two given sets of non-adversarial votes remain indistinguishable to the adversary and the proof is complete. □

In comparison with DRE-i, DRE-ip provides essentially the same level of security against such non-intrusive attacks as both systems guarantee ballot secrecy under the DDH assumption.

Ballot Secrecy under Intrusive Attacks. Now let us consider a stronger adversary that apart from the ability to determine an arbitrary number of votes, also gets read access to the voting machine (DRE) storage for a period during the voting phase. Obviously, such an adversary would be able to observe the votes cast during the access period and hence be able to at least work out the tally of the non-adversarial votes cast outside the access period. We prove that under the Square DDH assumption, this is the only information the adversary gains about the non-adversarial votes. In particular, we show that any two elections in which the non-adversarial votes cast outside the adversarial access period have the same tally are indistinguishable to the adversary. Note that in DRE-i, in case of an adversarial access to the voting machine storage, the privacy of the ballots cast outside the adversarial access period is also lost. Therefore, while DRE-i falls victim to such intrusive attacks, DRE-ip guarantees vote privacy under the Square DDH assumption.

We first prove the following lemma:

Lemma 2. *In DRE-ip, assuming that all proofs of well-formedness are zero knowledge, if the Square DDH assumption holds, then an adversary that determines an arbitrary number of votes and gets temporary read access to the voting machine (DRE) storage cannot distinguish between two bulletin boards in which two votes cast outside the access period are swapped.*

The proof of the lemma comes in Appendix B. The proof considers an adversary that not only can determine an arbitrary number of votes except two votes v_i and v_j, but gets access to DRE storage for an arbitrary period. Assuming that such an adversary is able to distinguish the bulletin boards in which v_i and v_j are swapped, we show how it can be used to break the Square DDH assumption. Basically, the proof shows that even if the value of the sum s is leaked to the adversary, ballot secrecy is still guaranteed, albeit under a stronger assumption.

Lemma 2 can then be similarly expanded to prove our main theorem for ballot secrecy under intrusive attacks as follows. We omit the proof of this theorem as it is similar to that of Theorem 2.

Theorem 3. *In DRE-ip, assuming that all proofs of well-formedness are zero knowledge, if the Square DDH assumption holds, then an adversary that determines an arbitrary number of votes and gets temporary read access to the voting machine (DRE) storage cannot gain any knowledge about the non-adversarial votes other than their tally.*

The main difference between Theorems 2 and 3 is that the latter depends on a stronger assumption (i.e., Square DDH) because of the additional data that becomes accessible to the adversary under an intrusive attack.

5 Comparison

In this section we look at how DRE-ip compares with other DRE-based verifiable e-voting systems. In particular, we consider Chaum's Votegrity [16], Neff's MarkPledge [30], VoteBox [33], STAR-Vote [4], and DRE-i [24].

Table 1. Selected security assumptions for DRE-based verifiable e-voting systems. TA: tallying authority, VIA: voter-initiated auditing, BB: bulletin board, RNG: random number generation, ■: assumption is required, □: assumption is not required.

	Availability	Integrity		Privacy				
System	Reliable TA(s)	Sufficient VIA	Secure BB	Secure setup	Secure RNG	Secure deletion	Secure ballot storage	Trust-worthy TA(s)
Votegrity	■	■	■	■	■	■	□	■
MarkPledge	■	■	■	■	■	■	□	■
VoteBox	■	■	■	■	■	■	□	■
STAR-Vote	■	■	■	■	■	■	□	■
DRE-i	□	■	■	■	■	■	■	□
DRE-ip	□	■	■	■	■	■	□	□

Votegrity is based on visual cryptography and uses onion encryption. MarkPledge employs a purpose-designed encryption scheme that allows challenge-response-style individual verifiability. VoteBox and STAR-Vote are both based on exponential ElGamal encryption which allows homomorphic tallying. DRE-i and DRE-ip on the other hand use encryption that does not admit to a fixed decryption key. All these systems consider voter registration and voter authentication outside their scope and assume they are carried out correctly and securely.

In general, systems that require tallying authorities, i.e. Votegrity, MarkPledge, VoteBox, and STAR-Vote, assume a minimum number of them are available at the tallying phase to compute the election tally. DRE-i and DRE-ip do not require such an assumption to guarantee availability.

To guarantee integrity, all these systems rely on a secure bulletin board and on a sufficient number of voters carrying out individual verification. Systems that require tallying authorities, i.e. Votegrity, MarkPledge, VoteBox, and STAR-Vote, also require that the tallying authorities perform the decryption of the tally correctly. In a verifiable system, this is enforced by requiring the tallying authorities to produce universally verifiable proofs of correct decryption. Hence, we consider assumptions underlying all the systems to guarantee integrity to be comparable, whether the system requires tallying authorities or not.

To guarantee privacy, all these systems assume a secure setup phase to generate and distribute system parameters and keys, as well as secure random number generators to produce the randomness required for probabilistic encryption. Furthermore, all systems assume that the captured votes and any ephemeral secrets generated for the cryptographic operations during the voting phase are securely erased. Votegrity is based on decryption mix-nets and requires that the tallying authorities do not collude to compromise voter privacy. MarkPledge employs a re-encryption mix-net to shuffle encrypted ballots before decryption, and assumes that the tallying authorities do not decrypt ballots before mixing although they are available on the bulletin board. VoteBox and STAR-Vote require that the tallying authorities do not collude to decrypt individual ballots. DRE-i does not require this assumption, but instead relies on a secure ballot storage mech-

Table 2. Computation complexity of selected DRE-based verifiable e-voting systems. $\mathbb{B}$, $\mathbb{A}$, $\mathbb{C}$: all, audited, confirmed ballots, exp: exponentiation, mul: multiplication.

System	Ballot calculation	Well-formedness and consistency verification	Tally calculation and verification
VoteBox	$6\lvert\mathbb{B}\rvert$ exp	$(7\lvert\mathbb{A}\rvert + 5\lvert\mathbb{C}\rvert)$ exp	$\lvert\mathbb{C}\rvert$ mul + 5 exp
STAR-Vote	$6\lvert\mathbb{B}\rvert$ exp	$(7\lvert\mathbb{A}\rvert + 5\lvert\mathbb{C}\rvert)$ exp	$\lvert\mathbb{C}\rvert$ mul + 5 exp
DRE-i	$11\lvert\mathbb{B}\rvert$ exp	$(10\lvert\mathbb{A}\rvert + 5\lvert\mathbb{C}\rvert)$ exp	$\lvert\mathbb{B}\rvert$ mul + 1 exp
DRE-ip	$13\lvert\mathbb{B}\rvert$ exp	$(15\lvert\mathbb{A}\rvert + 10\lvert\mathbb{C}\rvert)$ exp	$\lvert\mathbb{C}\rvert$ mul + 1 exp

anism to keep the pre-computed ballots safe after the setup phase. DRE-ip does not require trust assumptions on tallying authorities or ballot storage.

Table 1 summarizes the main similarities and differences in terms of their underlying security assumptions between the voting systems we consider.

The computation complexity of DRE-ip is compared with other DRE-based verifiable e-voting systems in Table 2. We do not consider Votegrity and Mark-Pledge since they use mix-nets and their computation complexity depend on how verifiable mix-nets are implemented. All calculations are based on a two-candidate election, one TA if present, and encryption implemented based on exponential ElGamal. Note that having multiple TAs increases the complexity of tally calculation and verification for all the schemes requiring tallying authorities. We assume in all systems that the TA, if present, provides proofs of correct decryption as required by end-to-end verifiability. We assume that the simultaneous multiple exponentiation technique [29] is used to optimize computations. Using this technique, computing $\mathrm{P}_{\mathrm{WF}}\{\tilde{X}_i\}$, $\mathrm{P}_{\mathrm{WF}}\{\tilde{Z}_i\}$, and $\mathrm{P}_{\mathrm{WF}}\{Z_i\}$ require 2, 2, and (equivalent to) around 4.4 exponentiations, respectively, and verification of these proofs take (equivalent to) around 2.4, 2.4, and 4.8 exponentiations, respectively. All systems make use of well-formedness proofs similar to these and we assume the same complexities as above for equivalent proofs in all systems. All numerical values in the table are rounded to the nearest integer.

6 Concluding Remarks

In this paper we revisited the design of the DRE-i voting system and proposed a new system: DRE-ip. On the theoretical level, we have shown that it is possible to have end-to-end verifiable DRE-base voting systems in which the privacy of the ballots does not rely on trustworthy tallying authorities or trusted hardware. On the practical level, we have shown that DRE-ip provides an efficient and practical DRE-based voting solution which is end-to-end verifiable and is able to preserve the privacy of the ballots even if the adversary gets temporary read access to the voting machine during the voting phase. Our system is described for two candidates, and can be extended to support multiple candidates via standard methods. Designing a system without tallying authorities that can efficiently support more complex electoral systems such as single transferable vote (STV) or write-in candidates remains an open problem.

Acknowledgement. The authors wish to thank Changyu Dong and the anonymous reviewers of ESORICS 2016 for their valuable comments. The authors are supported by the ERC Starting Grant No. 306994.

References

[1] B. Adida. Helios: Web-based open-audit voting. In *USENIX Security Symp.*, volume 17, pages 335–348, 2008.

[2] B. Adida, O. de Marneffe, O. Pereira, and J.-J. Quisquater. Electing a university president using open-audit voting: Analysis of real-world use of Helios. In *EVT/WOTE'09*, page 10. USENIX, 2009.

[3] O. Baudron, P.-A. Fouque, D. Pointcheval, J. Stern, and G. Poupard. Practical multi-candidate election system. In *ACM Symp. on Principles of Distributed Computing*, PODC '01, pages 274–283. ACM, 2001.

[4] S. Bell, J. Benaloh, M. D. Byrne, D. DeBeauvoir, B. Eakin, G. Fisher, P. Kortum, N. McBurnett, J. Montoya, M. Parker, O. Pereira, P. B. Stark, D. S. Wallach, and M. Winn. STAR-Vote: A secure, transparent, auditable, and reliable voting system. *USENIX Journal of Election Technology & Systems*, 1(1):18–37, 2013.

[5] M. Bellare and O. Goldreich. On defining proofs of knowledge. In E. F. Brickell, editor, *Crypto'92*, volume 740 of *LNCS*, pages 390–420. Springer, 1993.

[6] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM CCS'93*, pages 62–73. ACM, 1993.

[7] J. Ben-Nun, M. Llewellyn, B. Riva, A. Rosen, A. Ta-Shma, and D. Wikström. A new implementation of a dual (paper and cryptographic) voting system. In *EVOTE2012: 5th Int'l Conf. on Electronic Voting*, pages 315–329, 2012.

[8] J. Benaloh. Ballot casting assurance via voter-initiated poll station auditing. In *USENIX Workshop on Accurate E-Voting Technology (EVT)*, page 14, 2007.

[9] J. Benaloh and D. Tuinstra. Receipt-free secret-ballot elections. In *ACM Symp. on Theory of Computing*, STOC '94, pages 544–553. ACM, 1994.

[10] J. D. C. Benaloh. *Verifiable Secret-Ballot Elections.* PhD thesis, Department of Computer Science, Yale University, 1987.

[11] J.-M. Bohli, J. Müller-Quade, and S. Röhrich. Bingo voting: Secure and coercion-free voting using a trusted random number generator. In *E-Voting and Identity*, pages 111–124. Springer, 2007.

[12] J. Camenisch and M. Stadler. Efficient group signature schemes for large groups. In *Crypto'97*, volume 1294 of *LNCS*, pages 410–424. Springer, 1997.

[13] R. Carback, D. Chaum, J. Clark, J. Conway, A. Essex, P. Herrnson, T. Mayberry, S. Popoveniuc, R. Rivest, E. Shen, A. Sherman, and P. Vora. Scantegrity II municipal election at Takoma Park: The first E2E binding governmental election with ballot privacy. In *USENIX Security Symp.*, pages 291–306, 2010.

[14] D. Chaum, R. Carback, J. Clark, A. Essex, S. Popoveniuc, R. Rivest, P. Ryan, E. Shen, A. Sherman, and P. Vora. Scantegrity II: End-to-end verifiability by voters of optical scan elections through confirmation codes. *Information Forensics and Security, IEEE Transactions on*, 4(4):611–627, Dec 2009.

[15] D. L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–90, 1981.

[16] D. L. Chaum. Secret-ballot receipts: True voter-vrifiable elections. *IEEE security & privacy*, 2(1):38–47, 2004.

[17] R. Cramer, I. Damgård, and B. Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In Y. Desmedt, editor, *Crypto'94*, volume 839 of *LNCS*, pages 174–187. Springer, 1994.

[18] C. Culnane, P. Y. A. Ryan, S. Schneider, and V. Teague. vVote: A verifiable voting system. *ACM Trans. Inf. Syst. Secur.*, 18(1):3:1–3:30, June 2015.

[19] W. Diffie and M. E. Hellman. New directions in cryptography. *Information Theory, IEEE Transactions on*, 22(6):644–654, Nov 1976.

[20] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In A. M. Odlyzko, editor, *Crypto'86*, volume 263 of *LNCS*, pages 186–194. Springer, 1987.

[21] K. Fisher, R. Carback, and A. T. Sherman. Punchscan: Introduction and system definition of a high-integrity election system. In *Workshop on Trustworthy Elections (WOTE)*, 2006.

[22] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.

[23] F. Hao, D. Clarke, and A. Zorzo. Deleting secret data with public verifiability. *IEEE Transactions on Dependable and Secure Computing*, PP(99):1, 2015.

[24] F. Hao, M. N. Kreeger, B. Randell, D. Clarke, S. F. Shahandashti, and P. H.-J. Lee. Every vote counts: Ensuring integrity in large-scale electronic voting. *USENIX Journal of Election Technology & Systems*, 2(3):1–25, 2014.

[25] A. Juels, D. Catalano, and M. Jakobsson. Coercion-resistant electronic elections. In *Privacy in Electronic Society*, WPES'05, pages 61–70. ACM, 2005.

[26] E. Kiltz. A tool box of cryptographic functions related to the Diffie-Hellman function. In *IndoCrypt'01*, volume 2247 of *LNCS*, pages 339–349. Springer, 2001.

[27] A. Lim. Travis County, TX developing electronic voting system with a paper trail. *Government Technology*, July 2014. `www.govtech.com` (accessed Oct. 2015).

[28] U. M. Maurer and S. Wolf. Diffie-Hellman oracles. In N. Koblitz, editor, *Crypto'96*, volume 1109 of *LNCS*, pages 268–282. Springer, 1996.

[29] A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone. *Handbook of Applied Cryptography.* CRC press, 1996.

[30] C. A. Neff. Practical high certainty intent verification for encrypted votes, 2004. Avalable from `http://citeseer.ist.psu.edu`.

[31] P. Ryan, D. Bismark, J. Heather, S. Schneider, and Z. Xia. Prêt à Voter: a voter-verifiable voting system. *IEEE T. Inf. Foren. Sec.*, 4(4):662–673, Dec 2009.

[32] K. Sako and J. Kilian. Receipt-free mix-type voting scheme. In *EuroCrypt'95*, volume 921 of *LNCS*, pages 393–403. Springer, 1995.

[33] D. Sandler, K. Derr, and D. S. Wallach. VoteBox: A tamper-evident, verifiable electronic voting system. In *USENIX Security Symp.*, volume 4, page 87, 2008.

[34] C.-P. Schnorr. Efficient signature generation by smart cards. *Journal of cryptology*, 4(3):161–174, 1991.

[35] S. Wolf. *Information-Theoretically and Computationally Secure Key Agreement in Cryptography.* PhD thesis, ETH Zurich, 1999.

A Proof of Lemma 1

Proof. First, we consider the following assumption:

Assumption 3. For randomly chosen $a, b, c \in \mathbb{Z}_q^\star$, given $(g, g^a, g^b, g^c, g^{ac}, g^{bc}, g^{abc}, \Omega)$ where $\Omega \in \{g^{ab}, g^{ab+1}\}$, it is hard to decide whether $\Omega = g^{ab}$ or $\Omega = g^{ab+1}$.

The following lemma is proven in [24]:

Lemma 3. *The DDH assumption implies Assumption 3.*

In the rest of the proof we show that Lemma 1 holds under Assumption 3.

Let A be an adversary that after determining a number of votes distinguishes the two bulletin boards. We construct an algorithm D that given g, g^a, g^b, g^c, g^{ac}, g^{bc}, g^{abc}, and a challenge $\Omega \in \{g^{ab}, g^{ab+1}\}$ distinguishes which Ω is given.

Consider an abridged bulletin board resulting from removing the well-formedness proofs. Let us call this the *bare* bulletin board. Let the adversary determine any subset of votes other than the swapped votes v_i and v_j. D simulates the bare bulletin board as follows. We describe how confirmed ballots are constructed. Audited ballots can be easily calculated since x_k, y_k, and v_k are known to D for all $k \notin \{i, j\}$. Note that ballots i and j are confirmed ballots.

D posts g and $\tilde{g} = g^c$ as the initial parameters on the bulletin board. For all $k \notin \{i, j\}$, D simply chooses x_k and y_k randomly and generates the ballot according to the protocol. D generates random values α_i, β_i, α_j, and β_j and calculates the i-th and j-th ballots as follows. First, D sets

$$\begin{aligned} X_i &= g^{\alpha_i} g^a, & Y_i &= g^{\beta_i} g^b, & Z_i &= g^{\alpha_i \beta_i} (g^a)^{\beta_i} (g^b)^{\alpha_i} \Omega, \\ X_j &= g^{\alpha_j} g^a, & Y_j &= g^{\beta_j} / g^b, & Z_j &= g^{\alpha_j \beta_j + 1} (g^a)^{\beta_j} / ((g^b)^{\alpha_j} \Omega), \end{aligned}$$

i.e., we implicitly have $x_i = \alpha_i + a$, $y_i = \beta_i + b$, $x_j = \alpha_j + a$, and $y_j = \beta_j - b$. Since α_i, β_i, α_j, and β_j are random, X_i, Y_i, X_j, and Y_j are randomly distributed. Furthermore, D sets

$$\begin{aligned} \tilde{X}_i &= (g^c)^{\alpha_i} g^{ac}, & \tilde{Z}_i &= (g^c)^{\alpha_i \beta_i} (g^{ac})^{\beta_i} (g^{bc})^{\alpha_i} g^{abc}, \\ \tilde{X}_j &= (g^c)^{\alpha_j} g^{ac}, & \tilde{Z}_j &= (g^c)^{\alpha_j \beta_j} (g^{ac})^{\beta_j} / ((g^{bc})^{\alpha_j} g^{abc}). \end{aligned}$$

$\tilde{X}_i$, $\tilde{Z}_i$, $\tilde{X}_j$, and $\tilde{Z}_j$ are well-formed since we have:

$$\begin{aligned} \tilde{X}_i &= (g^c)^{\alpha_i} g^{ac} = (g^c)^{\alpha_i + a} = \tilde{g}^{x_i}, \\ \tilde{Z}_i &= (g^c)^{\alpha_i \beta_i} (g^{ac})^{\beta_i} (g^{bc})^{\alpha_i} g^{abc} = (g^c)^{(\alpha_i + a)(\beta_i + b)} = \tilde{g}^{x_i y_i}, \\ \tilde{X}_j &= (g^c)^{\alpha_j} g^{ac} = (g^c)^{\alpha_j + a} = \tilde{g}^{x_j}, \\ \tilde{Z}_j &= (g^c)^{\alpha_j \beta_j} (g^{ac})^{\beta_j} / ((g^{bc})^{\alpha_j} g^{abc}) = (g^c)^{(\alpha_j + a)(\beta_j - b)} = \tilde{g}^{x_j y_j}. \end{aligned}$$

Now if $\Omega = g^{ab}$, then we have

$$\begin{aligned} Z_i &= g^{\alpha_i \beta_i} (g^a)^{\beta_i} (g^b)^{\alpha_i} \Omega = g^{\alpha_i \beta_i} g^{a \beta_i} g^{b \alpha_i} g^{ab} = g^{(\alpha_i + a)(\beta_i + b)} = g^{x_i y_i} \quad \text{and} \\ Z_j &= g^{\alpha_j \beta_j + 1} (g^a)^{\beta_j} / ((g^b)^{\alpha_j} \Omega) = g^{\alpha_j \beta_j + 1} g^{a \beta_j} / (g^{b \alpha_j} g^{ab}) = g^{(\alpha_j + a)(\beta_j - b)} g = g^{x_j y_j} g \ . \end{aligned}$$

On the other hand, if $\Omega = g^{ab+1}$, then we have

$$\begin{aligned} Z_i &= g^{\alpha_i \beta_i} (g^a)^{\beta_i} (g^b)^{\alpha_i} \Omega = g^{\alpha_i \beta_i} g^{a \beta_i} g^{b \alpha_i} g^{ab+1} = g^{(\alpha_i + a)(\beta_i + b)} g = g^{x_i y_i} g \quad \text{and} \\ Z_j &= g^{\alpha_j \beta_j + 1} (g^a)^{\beta_j} / ((g^b)^{\alpha_j} \Omega) = g^{\alpha_j \beta_j + 1} g^{a \beta_j} / (g^{b \alpha_j} g^{ab+1}) = g^{(\alpha_j + a)(\beta_j - b)} = g^{x_j y_j} \ . \end{aligned}$$

In other words, $\Omega = g^{ab}$ corresponds to a bulletin board with $v_i = 0$ and $v_j = 1$, and $\Omega = g^{ab+1}$ corresponds to a bulletin board with $v_i = 1$ and $v_j = 0$, with all other votes being identical in the two bulletin boards.

Since all the votes other than v_i and v_j, including the ones chosen by the adversary, are known to D, it can calculate the partial tally $t_1 = \sum_{\forall k \notin \{i,j\}} v_k$. Since we have either $v_i = 0$ and $v_j = 1$, or $v_i = 1$ and $v_j = 0$, the total tally can be calculated as $t = t_1 + 1$. Now note that we implicitly have:

$$\begin{aligned} x_i y_i + x_j y_j &= (\alpha_i + a)(\beta_i + b) + (\alpha_j + a)(\beta_j - b) \\ &= (\beta_i + \beta_j)a + (\alpha_i - \alpha_j)b + (\alpha_i \beta_i + \alpha_j \beta_j), \end{aligned}$$

and hence, defining $s_1 = \sum_{\forall k \notin \{i,j\}} x_k y_k$, we have:

$$s = \sum_{\forall k} x_k y_k = (\beta_i + \beta_j)a + (\alpha_i - \alpha_j)b + (\alpha_i \beta_i + \alpha_j \beta_j) + s_1 .$$

Therefore, D can calculate a well-formed $S = g^s$ as follows:

$$S = (g^a)^{\beta_i + \beta_j} (g^b)^{\alpha_i - \alpha_j} g^{\alpha_i \beta_i + \alpha_j \beta_j + s_1} .$$

Thus, D is able to simulate all the elements of a bare bulletin board. Since the proofs of well-formedness are assumed to be zero knowledge, they can be simulated in the Random Oracle Model for ballots i and j. The zero knowledge property ensures that simulated proofs remain indistinguishable from real proofs. Consequently, D is able to simulate a full bulletin board corresponding to one of the two cases, with $\Omega = g^{ab}$ corresponding to the case where $v_i = 0$ and $v_j = 1$, and $\Omega = g^{ab+1}$ corresponding to $v_i = 1$ and $v_j = 0$, with all other votes being identical in the two bulletin boards. If A is able to distinguish the two cases, so will be D, and hence the proof is complete. □

B Proof of Lemma 2

First, we consider the following assumption:

Assumption 4. For randomly chosen $a, b, c, m, n \in \mathbb{Z}_q^\star$, given $(m, n, ma + nb, g, g^a, g^b, g^c, g^{ac}, g^{bc}, g^{abc}, \Omega)$ where $\Omega \in \{g^{ab}, g^{ab+1}\}$, it is hard to decide whether $\Omega = g^{ab}$ or $\Omega = g^{ab+1}$.

In Section B.1 we prove that the Square DDH assumption implies Assumption 4. In Section B.2 we show that the lemma holds under Assumption 4.

B.1 Square DDH Implies Assumption 4

The proof is composed of three parts: Lemma 4 which shows that Square DDH implies Assumption 5, Lemma 5 which shows that Assumption 5 implies Assumption 6, and Lemma 6 which shows that Assumption 6 implies Assumption 4. Figure 3 summarizes the relations between these assumptions. The definitions of the assumptions, the lemmas, and their proofs follow.

$$\begin{tikzcd} \text{Square DDH} \arrow[r, Rightarrow, "\text{Lem 4}"] & \text{Asm 5} \arrow[r, Rightarrow, shift left, "\text{Lem 5}"] & \text{Asm 6} \arrow[l, Rightarrow, shift left, "\text{Note 2}"] \arrow[r, Rightarrow, shift left, "\text{Lem 6}"] & \text{Asm 4} \arrow[l, Rightarrow, shift left, "\text{Note 3}"] \arrow[r, Rightarrow, "\text{Note 1}"] & \text{DDH} \end{tikzcd}$$

Fig. 3. Relations between assumptions. Asm: Assumption, Lem: Lemma, $\Longrightarrow$: implies.

Assumption 5. For randomly chosen $a, c \in \mathbb{Z}_q^\star$, and $R \in \mathbb{G}_q$, given $(g, g^a, g^c, g^{ac}, g^{a^2c}, \varOmega)$ where $\varOmega \in \{g^{a^2}, R\}$, it is hard to decide whether $\varOmega = g^{a^2}$ or $\varOmega = R$.

Lemma 4. *Square DDH (Assumption 2) implies Assumption 5.*

Proof. Considering g^c as the new g and c^{-1} as the new c, it is not hard to see that Assumption 5 is equivalent to the following assumption: For randomly chosen $a, c \in \mathbb{Z}_q^\star$, and $R \in \mathbb{G}_q$, given $(g, g^a, g^c, g^{ac}, g^{a^2}, \varOmega)$ where $\varOmega \in \{g^{a^2c}, R\}$, it is hard to decide whether $\varOmega = g^{a^2c}$ or $\varOmega = R$. To prove this latter assumption based on Square DDH, consider the following distributions for a random $b \in \mathbb{Z}_q^\star$:

$$D_1 = (g, g^a, g^c, g^{ac}, g^{a^2}, g^{a^2c}), \qquad D_2 = (g, g^a, g^c, g^{ac}, g^b, g^{bc}),$$
$$D_3 = (g, g^a, g^c, g^{ac}, g^b, R), \qquad D_4 = (g, g^a, g^c, g^{ac}, g^{a^2}, R).$$

Let $\overset{c}{\approx}$ denote computational indistinguishability. We prove in the following that Square DDH implies $D_1 \overset{c}{\approx} D_2 \overset{c}{\approx} D_3 \overset{c}{\approx} D_4$.

First, we claim that Square DDH implies $D_1 \overset{c}{\approx} D_2$. Otherwise, given an algorithm D that distinguishes between D_1 and D_2, and given a Square DDH challenge $(g, A = g^a, \varOmega)$, we choose a random $c \in \mathbb{Z}_q^\star$ and construct the tuple $\tau = (g, A, g^c, A^c, \varOmega, \varOmega^c)$ and give it as input to D. Note that if $\varOmega = g^{a^2}$, then τ belongs to D_1, and if $\varOmega$ is random, then τ belongs to D_2. Hence, a successful D can be employed to solve a Square DDH challenge, and the claim is proven.

Second, we claim that DDH implies $D_2 \overset{c}{\approx} D_3$. Otherwise, given an algorithm D that distinguishes between D_2 and D_3, and given a DDH challenge $(g, B = g^b, C = g^c, \varOmega)$, we choose a random $a \in \mathbb{Z}_q^\star$ and construct the tuple $\tau = (g, g^a, C, C^a, B, \varOmega)$ and give it as input to D. Note that if $\varOmega = g^{bc}$, then τ belongs to D_2, and if $\varOmega$ is random, then τ belongs to D_3. Hence, a successful D can be employed to solve a DDH challenge, and the claim is proven.

Third, we claim that Square DDH implies $D_3 \overset{c}{\approx} D_4$. Otherwise, given an algorithm D that distinguishes between D_3 and D_4, and given a Square DDH challenge $(g, A = g^a, \varOmega)$, we choose random $c \in \mathbb{Z}_q^\star$ and $R \in \mathbb{G}_q$ and construct the tuple $\tau = (g, A, g^c, A^c, \varOmega, R)$ and give it as input to D. Note that if $\varOmega$ is random, then τ belongs to D_3, and if $\varOmega = g^{a^2}$, then τ belongs to D_4. Hence, a successful D can be employed to solve a Square DDH challenge, and the claim is proven.

The three claims above along with the fact that Square DDH implies DDH together imply that $D_1 \overset{c}{\approx} D_4$ and complete the proof. □

Note 1. One can show that $D_1 \overset{c}{\approx} D_2$, $D_2 \overset{c}{\approx} D_3$, and $D_3 \overset{c}{\approx} D_4$ each imply DDH or Square DDH, and hence prove that Assumption 5 implies DDH.

Assumption 6. For randomly chosen $a, b, c, m, n \in \mathbb{Z}_q^\star$, and $R \in \mathbb{G}_q$, given $(m, n, ma + nb, g, g^a, g^b, g^c, g^{ac}, g^{bc}, g^{abc}, \Omega)$ where $\Omega \in \{g^{ab}, R\}$, it is hard to decide whether $\Omega = g^{ab}$ or $\Omega = R$.

Lemma 5. *Assumption 5 implies Assumption 6.*

Proof. We show that given an algorithm D that breaks Assumption 6 we can break Assumption 5. Given $(g, A = g^a, C = g^c, \bar{B} = g^{ac}, Y = g^{a^2 c}, \Omega_1)$ where Ω_1 is either g^{a^2} or random, we choose random $m, n, \ell \in \mathbb{Z}_q^\star$ and calculate $B = g^{\ell/n}/A^{m/n}$, $\bar{A} = C^{\ell/n}/\bar{B}^{m/n}$, $X = \bar{B}^{\ell/n}/Y^{m/n}$, and $\Omega_2 = A^{\ell/n}/\Omega_1^{m/n}$, and pass $(m, n, \ell, g, A, B, C, \bar{B}, \bar{A}, X, \Omega_2)$ to D. Let us implicitly set $b = (\ell - ma)/n$. We now have $\ell = ma + nb$, and also

$$B = g^{\ell/n}/A^{m/n} = g^{(\ell-ma)/n} = g^b, \qquad \bar{A} = C^{\ell/n}/\bar{B}^{m/n} = g^{c(\ell-ma)/n} = g^{bc},$$
$$X = \bar{B}^{\ell/n}/Y^{m/n} = g^{ac(\ell-ma)/n} = g^{abc},$$

thus B, $\bar{A}$, and X are well-formed. Also note that if $\Omega_1 = g^{a^2}$, then

$$\Omega_2 = A^{\ell/n}/\Omega_1^{m/n} = g^{a(\ell-ma)/n} = g^{ab},$$

and if Ω_1 is random, then Ω_2 is also random. Therefore, a successful D can be used to distinguish between the two cases for Ω_1, and the claim is proven. □

Note 2. One can also show the reverse of the above lemma holds and hence prove that Assumptions 5 and 6 are in fact equivalent.

Lemma 6. *Assumption 6 implies Assumption 4.*

Proof. Let $\ell = ma + nb$. Consider the following distributions for a random $R \in \mathbb{G}_q$:

$$D_1 = (m, n, \ell, g, g^a, g^b, g^c, g^{ac}, g^{bc}, g^{abc}, g^{ab}),$$
$$D_2 = (m, n, \ell, g, g^a, g^b, g^c, g^{ac}, g^{bc}, g^{abc}, R),$$
$$D_3 = (m, n, \ell, g, g^a, g^b, g^c, g^{ac}, g^{bc}, g^{abc}, Rg),$$
$$D_4 = (m, n, \ell, g, g^a, g^b, g^c, g^{ac}, g^{bc}, g^{abc}, g^{ab+1}).$$

We have the following: $D_1 \overset{c}{\approx} D_2$ is equivalent to Assumption 6; $D_2 \overset{c}{\approx} D_3$ always holds; $D_3 \overset{c}{\approx} D_4$ is equivalent to Assumption 6; and finally $D_1 \overset{c}{\approx} D_4$ is equivalent to Assumption 4. Therefore, Assumption 6 implies $D_1 \overset{c}{\approx} D_2$ and $D_3 \overset{c}{\approx} D_4$, which together with $D_2 \overset{c}{\approx} D_3$ imply $D_1 \overset{c}{\approx} D_4$, which in turn implies Assumption 4. □

Note 3. One can also show the reverse of the above lemma holds and hence prove that Assumptions 6 and 4 are in fact equivalent.

In summary, Assumptions 4, 5, and 6 are all equivalent; they are all implied by Square DDH; and they all imply DDH. In other words, they are in between Square DDH and DDH in terms of hardness.

B.2 Assumption 4 Implies Lemma 2

Proof. The proof shares some ideas with Lemma 1. Let A be an adversary that after determining a number of votes and obtaining temporary access to the voting machine distinguishes the two cases. We construct an algorithm D that given m, n, $\ell = ma + nb$, g, g^a, g^b, g^c, g^{ac}, g^{bc}, g^{abc}, and a challenge $\Omega \in \{g^{ab}, g^{ab+1}\}$ distinguishes which Ω is given.

Let the adversary determine any subset of votes other than the swapped votes v_i and v_j. A has access to the bulletin board similar to the adversary in Lemma 1. Furthermore, A has temporary access to the voting machine which enables A to observe some votes and their respective secret values x_i and y_i, and also the value of $s = \sum x_j y_j$ for j belonging to all the votes or a subset thereof. Therefore, apart from simulating the values on the bulletin board, D ought to provide the adversary with the values of x_i and y_i for a subset of the votes and the value of $s = \sum x_j y_j$ for all the votes or a subset thereof.

D simulates the bare bulletin board broadly analogously to the proof of Lemma 1, except that it generates α_i and β_i randomly and sets $\alpha_j = \alpha_i - n$ and $\beta_j = m - \beta_i$, instead of generating all four randomly. Since m and n are random, α_j and β_j will still be distributed randomly. The rest of the simulation of the bulletin board is similar to that in the proof of Lemma 1. Let $s_1 = \sum_{\forall k \notin \{i,j\}} x_k y_k$. Similar to the proof of Lemma 1, at the end of the simulation we have $s = (\beta_i + \beta_j)a + (\alpha_i - \alpha_j)b + (\alpha_i\beta_i + \alpha_j\beta_j) + s_1$. However, since now $m = \beta_i + \beta_j$ and $n = \alpha_i - \alpha_j$, we have:

$$s = ma + nb + (\alpha_i\beta_i + \alpha_j\beta_j) + s_1 = \ell + (\alpha_i\beta_i + \alpha_j\beta_j) + s_1$$

Hence, s can be simulated given ℓ. In addition to the bulletin board, the adversary is then further provided with s and x_k and y_k for k in the period of access to the voting machine. Similar to the proof of Lemma 1, a successful adversary can be used to distinguish the two cases for Ω and hence the proof is complete. □