A Scalable Coercion-resistant Blockchain Decision-making Scheme

Zeyuan Yin*, Bingsheng Zhang*, Andrii Nastenko†, Roman Oliynykov†‡, Kui Ren*
*The State Key Laboratory of Blockchain and Data Security, Zhejiang University
Email: {zeyuanyin, bingsheng, kuiren}@zju.edu.cn
†IOG Singapore Pte Ltd
Email: {andrii.nastenko, roman.oliynykov}@iohk.io
‡V.N.Karazin Kharkiv National University, Ukraine

***Abstract*—Typically, a decentralized collaborative blockchain decision-making mechanism is realized by remote voting. To date, a number of blockchain voting schemes have been proposed; however, to the best of our knowledge, none of these schemes achieve coercion-resistance. In particular, for most blockchain voting schemes, the randomness used by the voting client can be viewed as a witness/proof of the actual vote, which enables improper behaviors such as coercion and vote-buying. Unfortunately, the existing coercion-resistant voting schemes cannot be directly adopted in the blockchain context. In this work, we design the first scalable coercion-resistant blockchain decision-making scheme that supports private differential voting power and 1-layer liquid democracy as introduced by Zhang *et al.* (NDSS'19). Its overall complexity is $O(n)$, where n is the number of voters. Moreover, the ballot size is reduced from Zhang *et al.*'s $\Theta(m)$ to $\Theta(1)$, where m is the number of experts and/or candidates. Its incoercibility is formally proven under the UC incoercibility framework by Alwen *et al.* (Crypto'15). We implement a prototype of the scheme and the evaluation result shows that our scheme's tally procedure is more than 6x faster than VoteAgain (USENIX'20) in an election with over 10,000 voters and over 50% extra ballot rate.**

1. Introduction

Blockchain technology enjoys its popularity since the invention of Bitcoin in 2008, and it continues to reshape our digital society with its property for decentralization, transparency, and security. Democracy on the blockchain has the potential to transform the way political systems function, creating a more equitable and inclusive future.

Democracy on the blockchain. In a democratic blockchain system, stakeholders have the right to participate in the decision-making process through verifiable remote voting where everyone can express his opinion. Usually, in a blockchain voting, each participant's voting power is different and is proportional to his stake, which is called differential voting power in this work. This is one of the main differences between blockchain voting and conventional elections, where typically one participant has one vote. On the other hand, direct democracy might not always be the best choice for a blockchain decision-making process. In practice, to make a wise decision, a stakeholder needs to rationally put substantial effort into informing himself, and also expert knowledge throughout the process. Therefore, letting elites lead the decision-making might be an optimization in most cases. Some systems such as ZCash [1] use a small committee (consisting of several experts) to make the decisions; however, this has the risk of centralization, i.e., if the committee behaves maliciously, there is no mechanism for stakeholders to alter their decisions whatsoever.

The concept of liquid democracy has been proposed to achieve better collaborative intelligence. Liquid democracy (also known as delegative democracy [2]) is a hybrid of direct democracy and representative democracy. It provides the benefits of both systems (whilst avoiding their drawbacks) by enabling organizations to take advantage of experts in a blockchain voting process, as well as giving the stakeholders the opportunity to vote. For each proposal, a voter can either vote directly or delegate his voting power to an expert who is knowledgeable and renowned in the corresponding area.

Zhang *et al.* [3] proposed a treasury system that supports liquid democracy. However, their scheme has the following two drawbacks: (i) the ballot size is linear in the number of candidates and/or experts; (ii) it is not coercion-resistant.

The coercion problem in remote voting. In real-world voting, a voting booth gives a voter privacy and protects him from being coerced. However, in remote voting, the voting procedure can be viewed as a probabilistic algorithm that takes as input a random coin and the voter's choice. If the output is published on the bulletin board, then we have a problem: the input and randomness used in the voting procedure can be viewed as a proof of casting a certain ballot, and anyone can run the probabilistic algorithm again to verify it, which makes coercion and vote-buying possible. Many well-known e-voting systems are not coercion-resistant, such as snapshot [4], Helios [5], and prêt à voter [6]. What's worse, in the blockchain context, vote-buying becomes easier with the help of smart contracts.

To address the coercion/vote-buying problem, several schemes are proposed. Generally, coercion-resistant voting can be divided into three categories: fake credentials [7], [8], [9], re-voting [10], [11], [12], [13], and secure hard-

ware [14], [15]. In a coercion-resistant voting scheme using *fake credentials*, a voter holds both real and fake credentials. If coerced, a voter will cast a ballot using a fake credential, which is indistinguishable from the real one in the coercer's view, and it will be silently uncounted in the tally phase. In a *re-voting* scheme, a voter can cast his ballot multiple times and only the last one will be tallied. Coercion-resistance relies on that the voter can cast the ballot again after the coercer leaves. In the schemes using *secure hardware*, the secure hardware has its internal randomness source and can do probabilistic encryption for the voter so that the voter can lie about what has been encrypted.

Coercion-resistance v.s. deniability. All types of coercion-resistant voting schemes must give a voter deniability through some technique. Concretely, in "*fake credentials*" schemes, the election authority will provide randomness in the credential-related elements, and generate a designated verifier proof of correctness. To deceive the coercer, a voter can generate a fake credential and claim it as the real one by simulating the designated verifier proof. Namely, the registration procedure is deniable. In *re-voting* schemes, the re-vote operation must be deniable, i.e., the tally procedure will not reveal if a voter has re-voted. In the schemes using *secure hardware*, the secure hardware hides the randomness in the ciphertext so that a voter can claim that it is encryption of another candidate, i.e., the encryption operation is deniable.

Challenges. Could we apply the aforementioned techniques to realize a coercion-resistant voting scheme in the blockchain context? It turns out to be a non-trivial task. First of all, secure hardware based solutions might not be suitable for the blockchain setting, because an open blockchain allows anyone to join and leave freely and not all devices are equipped with a secure hardware, such as a trusted execution environment (TEE).

How about "fake credentials" and "re-voting" schemes? Can we adapt those schemes with differential voting power? There are still some challenges. For instance, JCJ [7] is a well-known coercion-resistance voting scheme, and it can be modified to support differential voting power; however, the scheme has $O(n^2)$ complexity due to the pair-wise plaintext equivalence tests (PETs), where n is the total number of votes, limiting its scalability. On the other hand, although the recently proposed scheme, VoteAgain [13], offers quasi-linear complexity, its verifiability relies on a trusted third party (TTP), which is undesirable in the blockchain setting. The best-known candidate is Araújo *et al.*'s "fake credentials" scheme [9], which achieves $O(n)$ complexity without relying on TTP for verifiability. Unfortunately, the credential of Araújo *et al.*'s scheme is in the form of two group elements satisfying a linear relationship, and this makes it unable to be modified trivially to support differential voting power. To the best of our knowledge, no proper coercion-resistant voting scheme in the literature can support differential voting power and achieve $O(n)$ complexity at the same time. Hereby, we are asking the question:

Can we design a scalable (linear complexity) coercion-resistant delegated voting scheme for blockchain decision-making?

1.1. Our Approach

In this work, we answer the above question affirmatively by proposing a new coercion-resistant voting scheme. Our scheme belongs to the "fake credentials" category. We start with the well-known JCJ scheme [7]. In the JCJ scheme, each encrypted credential is put on the bulletin board in the registration phase, and each ballot generally consists of an encrypted candidate and an encrypted credential. In the tally phase, by a shuffle and pair-wise PETs on the credentials, the ballots with fake credentials will be silently eliminated.

It is intuitive that one can associate credentials with voting power in the JCJ scheme to support differential voting power, i.e., each encrypted credential is tied with an encrypted voting power and we still perform PETs on the encrypted credentials in the tally phase. However, the scheme will have $O(n^2)$ complexity, so it does not scale well when the number of voters is large. To improve scalability, we propose a novel *"dummy voting power"* technique. The key idea is that we allow voters to publish (encrypted) fake credentials associated with (encrypted) zero voting power on the bulletin board. Then, in the tally phase, after shuffle re-encrypting the real and fake credentials, all the credentials can be decrypted. In this way, we transform the pair-wise PETs into "decrypt and match", achieving $O(n)$ complexity (counting cryptographic operations only).

To achieve delegation, we design a *"two-layer homomorphic tally"* procedure consisting of "delegation calculation" and "final tally calculation". In layer one, delegation is calculated by decrypting voters' choices and adding the delegated voting power to the corresponding experts. In layer two, the final tally result is calculated by decrypting experts' choices and adding experts' voting power together with voters' direct votes. Thanks to the additive homomorphism of the encryption scheme, voters' ballots and voting power are hidden throughout the tally.

Combining the "dummy voting power" technique and "two-layer tally" procedure together, we build the first coercion-resistant voting scheme that has linear complexity and supports private differential voting power and liquid democracy. We perform the security analysis under the UC (universal composable) framework, and we prove that our scheme is UC coercion-resistant [15]. We implement the scheme and evaluate its performance. Results show that our scheme's tally execution time is more than 6x faster than VoteAgain [13] in elections with over 10,000 voters and over 50% extra ballot rate[1].

1.2. Related Work

Coercion-resistant voting can be roughly split into three classes: fake credentials [7], [16], [17], [8], [9], [18], re-voting [10], [11], [12], [13], and secure hardware [14],

1. In VoteAgain, it means that more than 50% voters re-voted once; in our scheme, it means that more than 50% voters cast a fake ballot.

TABLE 1. COMPARISON OF VOTING SCHEMES. HERE, n IS THE NUMBER OF VOTERS AND m IS THE NUMBER OF ELECTION CANDIDATES. DEL. MEANS DELEGATION. CRYPTO STATE MEANS THAT THE VOTER NEEDS TO KEEP A CRYPTOGRAPHIC SECRET FROM THE COERCER. EA STANDS FOR ELECTION AUTHORITY. HARDWARE MEANS THAT THE PROPERTY IS GUARANTEED BY SECURE HARDWARE. AN ITEM IN BOLD TEXT MEANS THAT OUR SCHEME IS THE BEST IN THIS ASPECT.

Schemes	Diff. voting power	Del.	Ballot size	Complexity	Crypto state	Ballot privacy	Verifiability	Coercion-resistant
JCJ [7] (fake credentials)	No	No	$O(1)$	$O(n^2)$	Yes	t-out-of-k	trust no one	trust EA
ABBT [9] (fake credentials)	No	No	$O(1)$	$O(n)$	Yes	t-out-of-k	trust no one	trust EA
AKL+ [11], LHK [12] (re-voting)	No	No	$O(1)$	$O(n^2)$	No	t-out-of-k	trust no one	secret credential
VoteAgain [13] (re-voting)	No	No	$O(1)$	$O(n \log n)$	No	t-out-of-k	trust EA	trust EA
MBC [14], AOZZ [15] (secure hardware)	No	No	$O(1)$	$O(n)$	No	hardware	hardware	hardware
Snapshot [4]	Yes	No	$O(1)$	$O(n)$	-	t-out-of-k	trust no one	-
ZOB [3]	Yes	Yes	$O(m)$	$O(mn)$	-	t-out-of-k	trust no one	-
Our scheme (fake credentials)	**Yes**	**Yes**	$\boldsymbol{O(1)}$	$\boldsymbol{O(n)}$	Yes	$\boldsymbol{t}$**-out-of-**$\boldsymbol{k}$	**trust no one**	trust EA

[15]. JCJ [7] is the first paper that introduces the *"fake credentials"* type of coercion-resistant voting. In JCJ, each ballot contains an encrypted credential and there is a list of encrypted valid credentials in the bulletin board. In the tally phase, by pair-wise plaintext equivalence tests (PETs), the ballots with invalid credentials will be eliminated, but the pair-wise PETs obviously have $O(n^2)$ complexity. Other "fake credentials" schemes such as [8], [9] improve the time complexity and achieve better properties such as everlasting privacy. The *re-voting* type of coercion-resistant voting allows a voter to cast multiple ballots and the tally procedure will only count the last one. Achenbach *et al.* [11] and Locher *et al.* [12] utilize a deniable vote update mechanism to realize re-voting with quadratic complexity. The Norwegian Internet voting protocol [10] and VoteAgain [13] achieve (quasi-)linear complexity, but they both need a trusted third party for verifiability. Schemes based on *secure hardware* [14], [15] can achieve coercion-resistance easily, but secure hardware is a strong assumption.

On the other hand, blockchain voting [19], [4], [3], [20] is becoming more and more popular nowadays. Snapshot [4] is a popular DAO (Decentralized Autonomous Organization) voting platform that frees voters from gas fees. It uses IPFS [21] to store the proposals and votes, making the voting process off-chain and gas-free. Zhang *et al.* [3] proposes a treasury system for blockchain governance. It supports liquid democracy and is provably secure, but its ballot size is linear to the candidate number. Besides, to the best of our knowledge, none of the existing blockchain voting schemes are coercion-resistant.

Finally, Table 1 gives a comparison between our scheme and previous work.

2. Preliminaries

Notations. Let $\lambda \in \mathbb{N}$ be the security parameter. Let $\mathbb{G}$ be a cyclic group of prime order p with group generator g. We abbreviate *probabilistic polynomial time* as PPT.

2.1. (Lifted) ElGamal Encryption

ElGamal encryption scheme consists of three PPT algorithms: the key generation algorithm $\mathsf{EC.Keygen}(\mathbb{G}, g, p)$ takes as input the group parameters and outputs a public-private key pair $(\mathsf{pk} := g^{\mathsf{sk}}, \mathsf{sk})$; the encryption algorithm $\mathsf{EC.Enc}_{\mathsf{pk}}(m)$ takes as input the public key pk and the message $m \in \mathbb{G}$ and outputs the ciphertext $c := (c_1, c_2) := (g^r, m \cdot \mathsf{pk}^r)$; the decryption algorithm $\mathsf{EC.Dec}_{\mathsf{sk}}(c)$ takes as input the secret key sk and the ciphertext c and outputs the message $m := c_2/c_1^{\mathsf{sk}}$.

ElGamal encryption is a re-randomizable encryption scheme. The re-encryption algorithm $\mathsf{EC.Rand}_{\mathsf{pk}}(c)$ takes as input a ciphertext $c := (c_1, c_2)$ and outputs the re-randomized ciphertext $c' := (c_1', c_2') := (g^r \cdot c_1, h^r \cdot c_2)$.

Lifted ElGamal encryption is a variant of ElGamal encryption. The encryption algorithm $\mathsf{LE.Enc}_{\mathsf{pk}}(m)$ takes as input the public key pk and the message m and outputs the ciphertext $c := (c_1, c_2) := (g^r, g^m \cdot \mathsf{pk}^r)$; the decryption algorithm $\mathsf{LE.Dec}_{\mathsf{sk}}(c)$ takes as input the secret key sk and the ciphertext c and outputs the message $m := \mathsf{Dlog}(c_2/c_1^{\mathsf{sk}})$, where $\mathsf{Dlog}(x)$ outputs the discrete logarithm of x (note that computing the discrete logarithm is inefficient, thus the message space should be small in practice).

Clearly, the (lifted) ElGamal encryption scheme is IND-CPA secure under the DDH assumption (see Appendix B for formal definition). Lifted ElGamal encryption is additively homomorphic, i.e., $\mathsf{LE.Enc}_{\mathsf{pk}}(m_1) \cdot \mathsf{LE.Enc}_{\mathsf{pk}}(m_2) = \mathsf{LE.Enc}_{\mathsf{pk}}(m_1 + m_2)$. Besides, (lifted) ElGamal encryption can be distributed as a threshold encryption scheme [22].

2.2. Signature

A signature scheme Sig is defined by three PPT algorithms: A key generation algorithm $\mathsf{Sig.Keygen}(1^\lambda)$ that generates a public-private key pair $(\mathsf{pk}, \mathsf{sk})$; a signing algorithm $\sigma \leftarrow \mathsf{Sig.Sign}_{\mathsf{sk}}(m)$ that generates a signature on message m; and a verification algorithm $\mathsf{Sig.Verify}_{\mathsf{pk}}(\sigma, m)$ that outputs 1 if and only if σ is a valid signature on m.

A secure signature scheme is existentially unforgeable under chosen message attack (see Appendix B for formal definition).

2.3. Non-interactive Zero-knowledge Proof (NIZK)

A non-interactive zero-knowledge proof (NIZK) consists of four PPT algorithms: $\{\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify}, \mathsf{Sim}\}$

and is complete, sound, and zero-knowledge (see Appendix B for formal definition). Our scheme utilizes six zero-knowledge proofs for proving: (i) voting power correctness ($\mathsf{NIZK}_{\mathsf{power}}$); (ii) ElGamal encryption plaintext knowledge ($\mathsf{NIZK}_{\mathsf{knowledge}}$); (iii) re-encryption correctness ($\mathsf{NIZK}_{\mathsf{DVF\text{-}reenc}}$); (iv) knowledge of secret key ($\mathsf{NIZK}_{\mathsf{sk}}$); (v) shuffle correctness ($\mathsf{NIZK}_{\mathsf{shuffle}}$); and (vi) decryption correctness ($\mathsf{NIZK}_{\mathsf{Dec}}$). We will give the details of these NIZKs in Appendix A.

2.4. Universal Composibility

We perform security analysis under the Universally Composable (UC) framework [23], [24]. In the UC framework, a protocol is represented by a set of interactive Turing machines (ITMs). Each ITM contains the program to be run by a party. Security is based on the indistinguishability between the real world execution $\mathsf{EXEC}_{\Pi,\mathcal{A},\mathcal{Z}}$ and the ideal world execution $\mathsf{EXEC}_{\mathcal{F},\mathcal{S},\mathcal{Z}}$. In $\mathsf{EXEC}_{\Pi,\mathcal{A},\mathcal{Z}}$, the parties run protocol Π with the adversary $\mathcal{A}$. In $\mathsf{EXEC}_{\mathcal{F},\mathcal{S},\mathcal{Z}}$, the parties interact with the ideal functionality $\mathcal{F}$ with the ideal adversary (simulator) $\mathcal{S}$. If for all PPT adversary $\mathcal{A}$ there exists a PPT simulator $\mathcal{S}$ such that no PPT environment $\mathcal{Z}$ can distinguish the real world and the ideal world, then we say that the protocol Π UC-realizes the ideal functionality $\mathcal{F}$.

2.5. Distributed Key Generation

Our voting scheme utilizes a distributed key generation protocol for threshold key generation. We use an ideal functionality $\mathcal{F}_{\mathsf{DKG}}^{t,k}$ [25] to abstract the DKG procedure. The functionality $\mathcal{F}_{\mathsf{DKG}}^{t,k}$ is depicted in Fig. 1. It interacts with key generators $\mathcal{P} := \{P_1, \ldots, P_k\}$ to generate a public key pk and deal the secret key shares sk_j to P_j. Meanwhile, it publishes each party P_i's partial public key ppk_i. To realize $\mathcal{F}_{\mathsf{DKG}}^{t,k}$, we can use the threshold distributed key generation protocol proposed by Gennaro *et al.* [22].

2.6. Public Bulletin Board

For a voting scheme, a public bulletin board is needed for broadcasting the ballots and other auxiliary information such as zk proofs. Formally, we use a shared functionality [24] $\mathcal{G}_{\mathsf{PBB}}$ to model the public bulletin board, as depicted in Fig. 2. $\mathcal{G}_{\mathsf{PBB}}$ has two interfaces: READ and WRITE, and it guarantees the anonymity of the sender. In practice, the blockchain serves as the public bulletin board.

2.7. Secure Channel

Coercion-resistant voting requires a secure channel between a voter and the authority. We model it as a UC secure channel functionality $\mathcal{F}_{\mathsf{sc}}$, as depicted in Fig. 3. Note that $\mathcal{F}_{\mathsf{sc}}$ only deals with transmission of a single message. Secure channel for multiple messages is obtained by invoking multiple sessions of $\mathcal{F}_{\mathsf{sc}}$.

Functionality $\mathcal{F}_{\mathsf{DKG}}^{t,k}[\mathbb{G}]$

The ideal functionality $\mathcal{F}_{\mathsf{DKG}}^{t,n}[\mathbb{G}]$ interacts with key generators $\mathcal{P} := \{P_1, \ldots, P_k\}$, an ideal adversary $\mathcal{S}$. It's parameterized with threshold t. Denote $\mathcal{P}_c$ as the set of corrupted generators, $\mathcal{P}_h := \mathcal{P} \backslash \mathcal{P}_c$ as the set of honest generators, and $|\mathcal{P}_c| \leq t-1$. $\mathcal{F}_{\mathsf{DKG}}$ maintains a set $\mathcal{N}$(initially set to $\emptyset$).

Upon receiving $(\textsc{CorruptShares}, \mathsf{sid}, (\{j, \mathsf{sk}_j\}_{P_j \in P_c}))$ from $\mathcal{S}$:

- Pick $\mathsf{sk} \leftarrow \mathbb{Z}_q$, and compute $\mathsf{pk} := g^{\mathsf{sk}}$;
- Construct random polynomial $F(x) := \sum_{b=0}^{t-1} a_b \cdot x^b$ under the restriction $F(j) = \mathsf{sk}_j$ for $P_j \in \mathcal{P}_c$, and $F(0) = \mathsf{sk}$;
- Compute $\mathsf{sk}_i := F(i)$ and $\mathsf{ppk}_i = g^{\mathsf{sk}_i}$ for $i \in [n]$;
- Send $(\textsc{KeyGen}, \mathsf{sid}, \{\mathsf{pk}_j\}_{j\in[n]})$ to $\mathcal{S}$ and send $(\textsc{KeyGen}, \mathsf{sid}, \mathsf{sk}_i)$ to $P_i, i \in [n]$.

Upon receiving $(\textsc{ReadPk}, \mathsf{sid})$ from any party, return $(\textsc{ReadPk}, \mathsf{sid}, \mathsf{pk}, \{\mathsf{ppk}_i\}_{i\in[n]})$ to the requestor.

Figure 1. DKG ideal functionality $\mathcal{F}_{\mathsf{DKG}}^{t,k}[\mathbb{G}]$

Functionality $\mathcal{G}_{\mathsf{PBB}}$

The ideal functionality $\mathcal{G}_{\mathsf{PBB}}$ is globally available to all participants. It is parameterized with a predicate $\mathsf{Validate}$.

Upon initialization, set $\mathsf{Storage} := \emptyset$.
Upon receiving $(\textsc{Read}, \mathsf{sid})$ from P:

- let $\mathsf{val} := \mathsf{Storage}[\mathsf{sid}]$;
- return $(\mathrm{READ}, \mathsf{sid}, \mathsf{val})$ to the requestor.

Upon receiving $(\textsc{Write}, \mathsf{sid}, \mathsf{inp})$ from P, do the following:

- let $\mathsf{val} := \mathsf{Storage}[\mathsf{sid}]$;
- if $\mathsf{Validate}(\mathsf{val}, \mathsf{inp}) = 1$, then set $\mathsf{Storage}[\mathsf{sid}] := \mathsf{val}||\mathsf{inp}$, return $(\textsc{Receipt}, \mathsf{sid})$ to the requestor;
- Otherwise, return $(\textsc{Reject}, \mathsf{sid})$ to the requestor.

Figure 2. Functionality $\mathcal{G}_{\mathsf{PBB}}$

Functionality $\mathcal{F}_{\mathsf{sc}}$

The ideal functionality $\mathcal{F}_{\mathsf{sc}}$ proceeds as follows.

- Upon receiving $(\textsc{Send}, \mathsf{sid}, R, m)$ from a party S, send $(\textsc{Sent}, \mathsf{sid}, S, R, |m|)$ to the adversary $\mathcal{S}$ and send $(\textsc{Sent}, \mathsf{sid}, S, m)$ to R.
- Ignore other requests.

Figure 3. Functionality $\mathcal{F}_{\mathsf{sc}}$

3. System Overview

In this section, we start by modifying the JCJ protocol [7] to build a coercion-resistant voting scheme that supports differential voting power with $O(n^2)$ complexity. Then, we give the intuition and details of our novel "dummy voting power" technique and "two-layer tally" procedure. We also optimize JCJ's ballot structure to achieve smaller ballot size. Finally, we provide an overview of our scheme.

3.1. The JCJ Protocol

We first recall the well-known JCJ protocol [7] and show how to modify it to support differential voting power.

The original protocol. As mentioned above, JCJ is the first protocol that introduces the concept of "fake credentials". Generally, it works as follows. For simplicity, we use $[\![x]\!]$ to denote encryption of x in sec. 3.1 and 3.2. In the registration phase, a voter authenticates to the election authority (EA) and the EA generates a credential $\sigma \leftarrow \mathbb{G}$. Then, the EA publishes $S = [\![\sigma]\!]$ on the PBB and sends σ to the voter along with a designated verifier proof that S is encryption of σ. In the voting phase, a voter cast a ballot $B = ([\![v]\!], [\![\sigma]\!], Pf)$ where $[\![v]\!]$ is the encryption of a candidate v, $[\![\sigma]\!]$ is the encryption of a credential σ, Pf includes the NIZK proofs of knowledge of v and σ, and a NIZK proof that $[\![v]\!]$ encrypts a valid candidate. If a voter is coerced, he generates a random $\sigma' \leftarrow \mathbb{G}$ and claims it as the real credential by simulating the designated verifier proof. In the tally phase, the trustees shuffle the ballots and perform pair-wise PETs on encrypted credentials to eliminate the ballots with fake credentials. Finally, the trustees decrypt the candidates and tally the votes. The overview of JCJ protocol is shown in Fig. 4.

Supporting differential voting power. We can see that if we associate each credential with voting power, then the system can easily support differential voting power. More specifically, in the registration phase, the EA publishes $S = ([\![\sigma]\!], [\![\alpha]\!])$ on the PBB, where $[\![\sigma]\!]$ is the encrypted credential and $[\![\alpha]\!]$ is the encrypted voting power under an additively homomorphic encryption scheme. After the pair-wise PETs, we get tuples of $\langle [\![v]\!], [\![\alpha]\!] \rangle$, where $[\![v]\!]$ is the encrypted candidate and $[\![\alpha]\!]$ is the encrypted voting power. Then, the trustees decrypt the candidates and add the voting power by additive homomorphism. The overview of the modified JCJ protocol with differential voting power is shown in Fig. 5.

3.2. Our Technique

In this part, we will illustrate our novel techniques that can achieve $O(n)$ complexity and delegated voting.

Dummy voting power. We can see that in the JCJ protocol, the pair-wise PETs cause $O(n^2)$ complexity. The idea is that: if we allow voters to publish the fake credentials on the PBB, but associate them with dummy (zero) voting power, then we can directly decrypt the credentials instead of performing PETs in the tally phase. In other words, in the

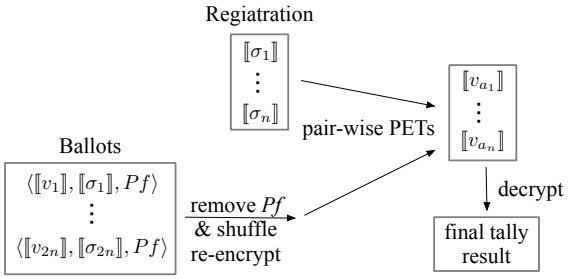

Figure 4. JCJ original protocol

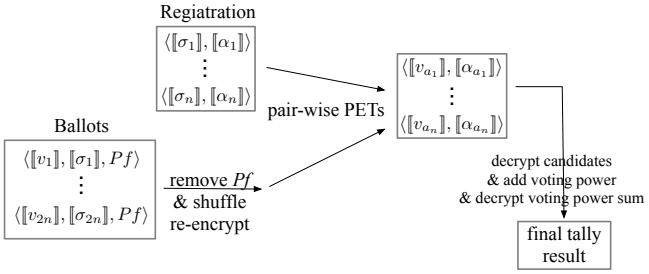

Figure 5. JCJ protocol with differential voting power

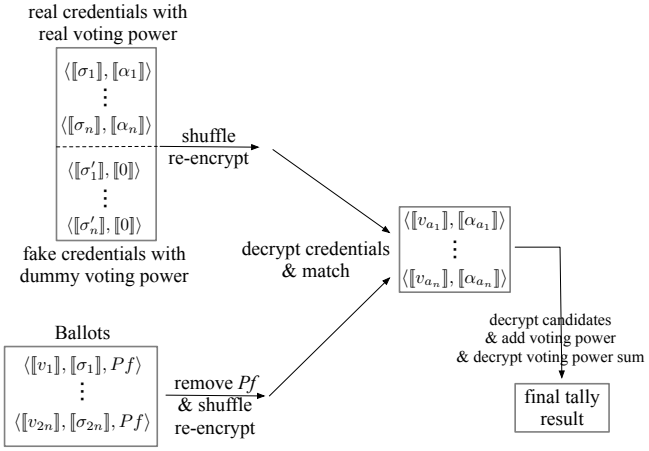

Figure 6. Dummy voting power technique

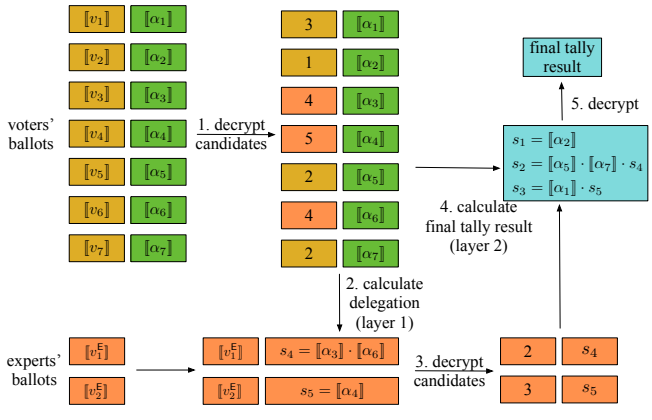

Figure 7. Two-layer tally

registration phase, the EA publishes (encrypted) real credentials and (encrypted) real voting power; at any convenient time, voters can also publish (encrypted) fake credentials and (encrypted) dummy voting power. In this way, we switch pair-wise PETs into shuffle-decrypt and matching, achieving linear complexity. Furthermore, "dummy voting power" also hides the number of votes obtained by each candidate. The idea of "dummy voting power" technique is shown in Fig. 6.

Two-layer tally. To support delegation, the trustees perform a "two-layer tally" procedure in the tally phase. Generally speaking, in layer one, voters' choices are decrypted and the delegated voting power will be added to the corresponding experts; in layer two, experts' choices are decrypted and the final tally result is calculated by adding experts' voting power and voters' direct votes. Note that, experts have input independence instead of ballot privacy so their ballots can be decrypted directly. Fig. 7 shows the process of "two-layer tally" where there are 3 candidates and 2 experts.

3.3. Overview of Our Scheme

In this section, we define the roles in our system and provide an overview of our scheme in the blockchain context.

Roles. There are five roles in the protocol: voters, experts, registration authority (RA), shuffler, and trustees.

- *A voter* has a certain amount of voting power and can either vote on the proposal directly or delegate his voting power to an expert.
- *An expert* does not have voting power himself, but he can be delegated to vote on others' behalf.
- *The RA* is responsible for the registration procedure.
- *The shuffler* performs verifiable shuffle procedures on the ciphertexts.
- *The trustees* are responsible for decrypting the ballot and revealing the final tally result.

We use n, m, ℓ to denote voter number, expert number, and candidate number, respectively. There are k trustees with a threshold t.

An optimization of JCJ's ballot structure. In our scheme, we optimize JCJ's ballot structure in the following two aspects. Firstly, we observe that it is not necessary to prove that $[\![v]\!]$ encrypts a valid candidate because all of them will be decrypted in the tally phase. If it encrypts an invalid value, we can simply treat it as "abstain" and drop it. Secondly, instead of defining σ as the secret credential, we can define the discrete logarithm of σ as the secret credential (voting secret key) and define σ as the voting public key, i.e., $\sigma := \mathsf{pk} := g^{\mathsf{sk}}$. Then, the ballot can be modified as $\langle \mathsf{pk}, \boldsymbol{u}, \pi, Sig \rangle$, where $\boldsymbol{u} := [\![v]\!]$ is the encrypted choice, π is a NIZK proof of plaintext knowledge of $\boldsymbol{u}$, and $Sig \leftarrow \mathsf{Sign}(\mathsf{sk}, \boldsymbol{u})$. By doing so, we change an Elgamal encryption $[\![\sigma]\!]$ to a group element pk, achieving smaller ballot size.

Overview. Our voting scheme has four phases: preparation phase, registration phase, voting/delegation phase, and tally phase.

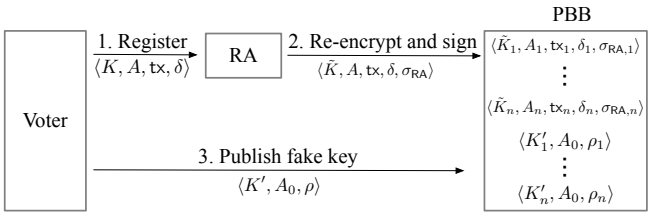

Figure 8. Registration phase

In the preparation phase, the RA generates a public-private signing key pair $\langle \mathsf{pk}_{\mathsf{RA}}, \mathsf{sk}_{\mathsf{RA}} \rangle$. The trustees perform a distributed key generation protocol to generate pk_{T} and they share sk_{T}.

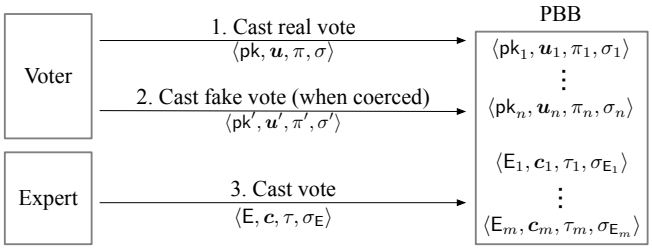

Figure 9. Voting/delegation phase

In the registration phase (see Fig. 8), each voter first freezes some stake by transaction tx. Then, he generates a pair of real voting keys $\langle \mathsf{pk}, \mathsf{sk} \rangle$ and sends a registration message $\langle K, A, \mathsf{tx}, \delta \rangle$ to the RA (step 1), where $K \leftarrow \mathsf{EC.Enc}_{\mathsf{pk}_{\mathsf{T}}}(\mathsf{pk})$ is encryption of real public key, $A \leftarrow \mathsf{LE.Enc}_{\mathsf{pk}_{\mathsf{T}}}(\alpha)$ is encryption of voting power, tx is the transaction that freezes some stake, and δ is a NIZK proof that the encrypted voting power equals the frozen stake (Cf. Appendix A). After authenticating the voter (i.e., checking that the voter knows the sk corresponding to tx's sender's pk by an interactive zero-knowledge protocol), the RA re-encrypts K as $\tilde{K}$ and signs the registration message. Then, the RA sends a designated verifier proof of re-encryption correctness to the voter, and sends $\langle \tilde{K}, A, \mathsf{tx}, \delta, \sigma_{\mathsf{RA}} \rangle$ to the PBB (step 2), where $\sigma_{\mathsf{RA}} \leftarrow \mathsf{Sig.sign}_{\mathsf{sk}_{\mathsf{RA}}}(\tilde{K}||A||\mathsf{tx}||\delta)$. At any convenient time, the voter can generate a pair of fake voting keys $\langle \mathsf{pk}', \mathsf{sk}' \rangle$ and publish a fake key item $\langle K', A_0, \rho \rangle$ on the PBB (step 3), where $K' \leftarrow \mathsf{EC.Enc}_{\mathsf{pk}_{\mathsf{T}}}(\mathsf{pk}')$ is encryption of fake public key, A_0 is a deterministic encryption of 0, and ρ is a NIZK proof of knowledge of sk'. The voter can repeat step 3 multiple times to generate multiple fake keys.

In the voting phase (see Fig. 9), each voter encrypts his choice with the trustees' public key pk_{T}, signs it with the voting secret key and casts it on PBB (step 1). Specifically, a voter's ballot is of the form $\langle \mathsf{pk}, \boldsymbol{u}, \pi, \sigma \rangle$, where $\boldsymbol{u} \leftarrow \mathsf{EC.Enc}_{\mathsf{pk}_{\mathsf{T}}}(v)$ is the encrypted choice, π is a NIZK proof of plaintext knowledge of $\boldsymbol{u}$, and $\sigma \leftarrow \mathsf{Sign}(\mathsf{sk}, \boldsymbol{u})$ is the signature. If a voter is coerced, he will use the fake key pair $\langle \mathsf{pk}', \mathsf{sk}' \rangle$ to perform the voting process (step 2). Thanks to the re-encryption by RA and the designated verifier proof, the voter can claim that $\tilde{K}$ is re-

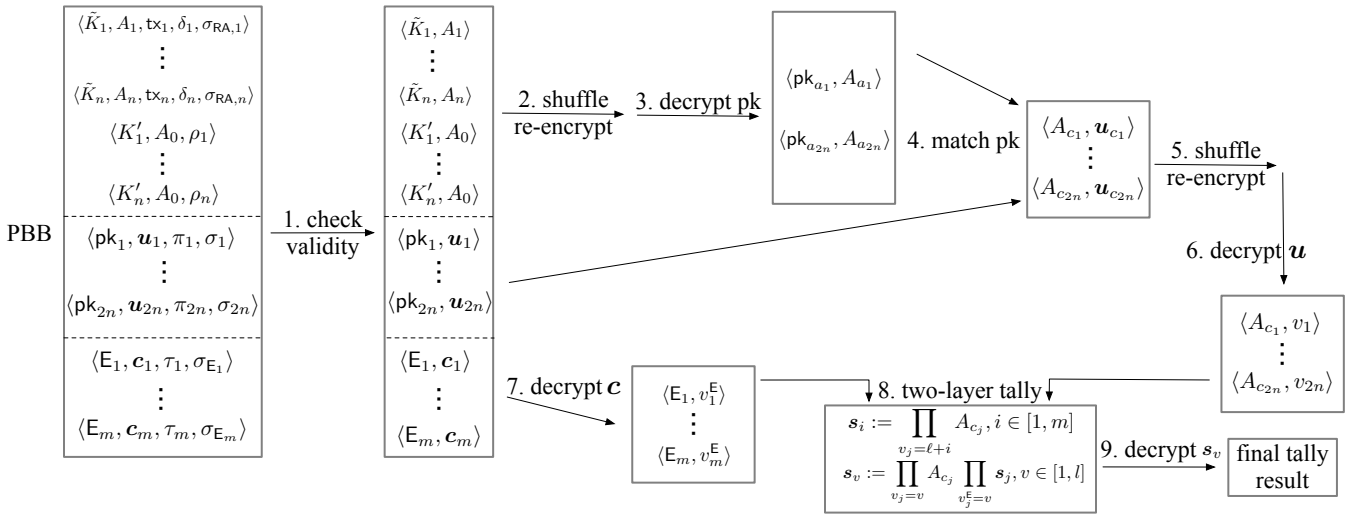

Figure 10. Tally phase

encryption of $\mathsf{EC.Enc}_{\mathsf{pk_T}}(\mathsf{pk}')$ by simulating the designated verifier proof. In this phase, each expert also casts his ballot by simply encrypting the choice and signing it with his blockchain secret key (step 3), i.e., an expert's ballot is of the form $\langle \mathsf{E}, \boldsymbol{c}, \tau, \sigma_{\mathsf{E}} \rangle$, where E is the expert's identity, $\boldsymbol{c} \leftarrow \mathsf{EC.Enc}_{\mathsf{pk_T}}(v^{\mathsf{E}})$ is the encrypted candidate, τ is the NIZK proof of plaintext knowledge, and σ_{E} is the signature.

In the tally phase (see Fig. 10), the PBB contains "encrypted public key items", voters' ballots and experts' ballots at the beginning. For simplicity, in Fig. 10, we assume there are n real ballots and n fake ballots; in reality, a voter can cast any number of fake ballots. Firstly, the shuffler checks the validity of all the "encrypted public key items" and ballots, and removes the NIZK proofs and signatures (step 1). Then, the shuffler shuffle re-encrypts the "encrypted public key items" (step 2). Next, the trustees jointly decrypt the public keys in "encrypted public key items" and voters' ballots (step 3). If the same public key appears more than once, then drop them. The encrypted voting power and encrypted choices with the same public key will be matched (step 4). To ensure ballot privacy, the matched items need to be shuffle re-encrypted (step 5). Next, the trustees jointly decrypt voters' choices (step 6) and experts' choices (step 7). After the decryption, the trustees will add the voting power to the corresponding candidates by a two-layer tally (step 8). In layer 1, each expert's obtained voting power is calculated, i.e., expert E_i's obtained voting power is $\boldsymbol{s}_i := \prod_{v_j=\ell+i} A_{c_j}, i \in [1,m]$, where ℓ is the candidate number and m is the expert number; in layer 2, the votes for each candidate are tallied by adding direct votes and expert votes together, i.e., candidate v's obtained voting power is $\boldsymbol{s}_v := \prod_{v_j=v} A_{c_j} \prod_{v_j^{\mathsf{E}}=v} \boldsymbol{s}_j, v \in [1,l]$. Ballots that encrypt an invalid choice will be dropped. Finally, the trustees jointly decrypt $\{\boldsymbol{s}_v\}_{v\in[1,\ell]}$ to publish the final tally result (step 9).

Blockchain deployment. To deploy the scheme on a blockchain, we need to select the RA, shuffler, and trustees properly. Also, there should be a validator that checks all the NIZK proofs.

The RA. Note that the communication between the voter and the RA must be secret to the coercer. Also, the RA is trusted for coercion-resistance and cannot be distributed (see sec. 4.1 for details). Therefore, it may be instantiated with trusted execution environment (TEE) like Intel SGX.

The shuffler. The shuffler can be implemented by a mixnet [26], and the mixnet nodes can be selected by cryptographic sortition [27]. Ballot privacy is preserved as long as one mixnet node is honest.

The trustees. The trustees can also be selected by cryptographic sortition [27]. In the blockchain context, the majority of trustees are honest with a high probability when the majority of the stake is honest.

The validator. Every participant can be the validator to check all the NIZK proofs if he wants. Since there are shuffle proofs in our scheme, whose verification cost is relatively heavy, it is not recommended to deploy a smart contract to play the role of the validator.

Roles of the blockchain. In our scheme, the blockchain serves as the PBB. It also plays the role of PKI in the registration phase to authenticate a voter.

4. Assumptions and Security Modeling

In this section, we first informally define ballot privacy, verifiability, and coercion-resistant and we analyze under which assumptions these properties are achieved. Then, we formally introduce the UC incoercibility framework.

4.1. Security Properties and Assumptions

Definition 1. (Ballot privacy) The adversary cannot learn the votes of honest voters.

Definition 2. (Verifiability) Honest voters' ballots must be tallied and the adversary cannot cast more votes than the number of voters that he controls.

Definition 3. (Coercion-resistance) A coercer cannot determine if the coerced party is trying to deceive him.

Assumptions. Our scheme only relies on basic assumptions that any "fake credentials" type of coercion-resistant voting scheme needs. We list these assumptions, explain the necessity of them, and discuss how they are achieved in the blockchain context.

Assumption 1. The public bulletin board is honest and the communication with the PBB is anonymous.

PBB is a basic assumption for any electronic voting scheme. A malicious PBB can break verifiability by creating different views for different voters [28]. Then, since ballots can be dropped undetectably, ballot privacy will be undermined [29], and it is not possible to have coercion-resistance without ballot privacy. Thus, PBB is trusted for all three properties. Moreover, communication with PBB must be anonymous; otherwise, the coercer will catch the deceiving voter when he tries to cast the real ballot.

In our system, the blockchain is a public ledger and serves as the honest PBB. A voter can use anonymous channels (e.g., TOR) to broadcast on the blockchain.

Assumption 2. There is a secure (untappable) channel between the voter and the RA.

In all "fake credentials" schemes, a voter needs to establish a secret in the registration phase and keep the secret from the coercer. If the coercer taps all the communication between the voter and the authorities, then the voter's private information is a receipt/witness of what he cast [30].

As mentioned above, the RA can be instantiated with TEE such as Intel SGX. A voter can use TOR to communicate with the RA.

Assumption 3. The authentication is inalienable [11], i.e., the coercer cannot impersonate the voter or stop the voter from authenticating.

Inalienable authentication is a must for all voting schemes. Otherwise, the adversary can vote on the voter's behalf or launch a forced abstention attack.

In the blockchain context, the blockchain plays the role of PKI and each voter authenticates to the RA by proving knowledge of his blockchain secret key. It is assumed that a voter will not reveal his blockchain secret key to the coercer, which will put the voter's stake at risk.

In the following, we analyze how ballot privacy, verifiability, and coercion-resistance are achieved in our scheme.

Ballot privacy. In the registration phase, the voter himself generates the voting public key, and he encrypts it with the trustees' public key before sending it to the RA. Thus, nobody knows the link between the voting public key and the voter as long as the majority of trustees are honest. In the voting/delegation phase, voters' ballots are also encrypted with the trustees' public key. In the tally phase, ballots and "encrypted public key items" will be shuffled before decryption so that the link between identity and ballot/voting public key is broken by the shuffle. Therefore, ballot privacy is achieved if the shuffler and the majority of trustees are honest.

TABLE 2. TRUST ASSUMPTIONS ON THE ENTITIES.

	Ballot privacy	Verifiability	Coercion-resistance
PBB	Trusted	Trusted	Trusted
RA	Untrusted	Untrusted	Trusted
Shuffler	Trusted	Untrusted	Trusted
Trustees	t-out-of-k	Untrusted	t-out-of-k

Note: In delegated voting, usually the experts have input independence rather than ballot privacy, i.e., when casting the ballot, it should be independent of the others; later in the tally phase, it will be decrypted directly without shuffle. This is an important requirement for delegated voting because we want to detect if an expert's behavior deviates from what he claimed.

Verifiability. A process composed of several subroutines is verifiable if each subroutine is verifiable itself. In the preparation phase, the trustees perform a verifiable distributed key generation protocol [22]. In the registration phase, we use two NIZKs to ensure that (i) the encrypted voting power is equal to the frozen stake; (ii) the RA does the re-encryption correctly. In the voting/delegation phase, EUF-CMA property of the signature scheme prevents anyone who does not know the secret key from casting a valid ballot. In the tally phase, the shuffle correctness is guaranteed by the shuffle NIZK [31], and decryption correctness is guaranteed by the decryption NIZK [22]. In conclusion, all subroutines in our scheme are publicly verifiable so no one needs to be trusted for verifiability.

Coercion-resistance. A coercer may ask the voter to reveal his real voting key pair, but the voter can claim a fake key pair as real by simulating the designated verifier proof, as long as the RA is not colluding with the coercer. In the tally phase, after shuffling all the "encrypted public key items", real keys and fake keys become indistinguishable from the coercer's perspective. Besides, the majority of trustees must be honest to ensure that the coercer cannot decrypt the ciphertexts.

Finally, Table 2 summarizes the trust assumptions on the entities for achieving each property.

Notes on distributing the shuffler and RA: To distribute the shuffler, a mixnet [26] can be utilized and we can perform a cryptographic sortition [27] on the blockchain to select the mixnet nodes. The assumption becomes that at least one of the mixnet nodes is honest instead of trusting a single shuffler.

However, simply distributing the RA does not lead to a weaker assumption because the coercer can ask the voter to provide the entire view of the registration phase. Even if only one of the RA parties is colluding with the coercer, the voter who does not know which RA party is colluding cannot simulate the registration view with negligible fail probability. Concretely, if the voter fakes a message sent by a RA member, then he will have at least $1/n$ probability of being caught (in the case that the RA member is malicious),

where n is the number of RA parties. Therefore, it is better not to distribute the RA for "fake credentials" schemes. We suggest using TEE to instantiate the RA on the blockchain.

4.2. Security Modeling

Framework. We formally perform security analysis under the Universal Composable (UC) framework [23]. As mentioned in sec. 2.4, security is based on the indistinguishability between the ideal world and the real/hybrid world.

Modeling coercion-resistance. We adopt the UC incoercibility definition proposed by Alwen, Ostrovsky, Zhou, and Zikas (AOZZ) [15]. In their definition, the ideal deception strategy DI in the ideal world is controlled by the environment $\mathcal{Z}$ and the ideal adversary $\mathcal{S}$ plays the role of coercer. When DI receives an input x from the ideal adversary (coercer) $\mathcal{S}$, the environment $\mathcal{Z}$ maps x to x' and instructs DI to forward x' to the ideal functionality $\mathcal{F}$ (x' can be equal to x). In the real world, the real deception strategy DR is a twist on the protocol Π and DR internally runs DI. DR will interact with the real adversary (coercer) $\mathcal{A}$ and do actions according to DI's output.

We can see that, in the ideal world, the ideal deceiving strategy can map the coercer's input x to any other input x', which represents "cast-as-intend". Also, the ideal adversary $\mathcal{S}$ determines if a party is deceiving no better than someone who only sees the outputs of the computation. Thus, if for every DI there exists DR that makes the ideal world and the real world indistinguishable, we say that the protocol Π is IUC (incoercible UC) secure.

The ideal world execution. In the ideal world, the voters $\mathcal{V}$, experts $\mathcal{E}$, trustees $\mathcal{T}$, shuffler, registration authority RA, and ideal adversary $\mathcal{S}$ communicate with the ideal functionality $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ (Cf. Fig. 11). The ideal functionality $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ has four phases: preparation phase, registration phase, voting/delegation phase, and tally phase.

Preparation phase. In the preparation phase, the ideal functionality $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ receives the initialization commands from the trustees and the RA, and it sends initialization notification to the simulator $\mathcal{S}$. At the end of the preparation phase, $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ sets $\mathsf{state} := 1$.

Registration phase. In the registration phase, the ideal functionality $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ receives registration requests from voters and records their voting power. If there are t or more corrupted trustees, voting power will be leaked to the ideal world adversary $\mathcal{S}$. At the end of the registration phase, $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ sets $\mathsf{state} := 2$.

Voting/Delegation phase. In the voting/delegation phase, the ideal functionality $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ receives vote commands from voters and experts, and records their votes. Similarly, if there are t or more corrupted trustees, the votes will be leaked to $\mathcal{S}$. At the end of the voting/delegation phase, $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ sets $\mathsf{state} := 3$.

Tally phase. In the tally phase, the ideal functionality $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ receives tally commands from the trustees and performs the tally algorithm. The tally algorithm $\mathsf{TallyAlg}$ (Cf. Fig 12) takes as input all the ballots and voting power and outputs the final tally result. During the tally, experts' ballots and voters' ballot count are leaked (Cf. Fig. 13), but this does not affect ballot privacy of voters.

Ideal deception. The ideal deceiving strategy DI is controlled by the environment $\mathcal{Z}$. When a voter V_i is coerced by the ideal adversary $\mathcal{S}$ to vote for x, $\mathcal{Z}$ maps x to x' and instrusts DI_i to send $(\textsc{Vote}, \mathsf{sid}, x')$ to $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$. Note that, in the modeling, DI_i can either obey (i.e., $x = x'$) or deceive (i.e., $x \neq x'$).

Connection with the properties. It is easy to see that a protocol Π IUC-realizing $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ has ballot privacy, verifiability and coercion-resistance. Firstly, voters' ballots are never leaked by $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ if the shuffler and majority of trustees are honest. Secondly, nobody can falsify the final tally result computed by $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$. Thirdly, the definition of the ideal deception DI ensures that a voter can cast-as-intend even if he is being coerced.

The real world execution. In the real world, we invoke the distributed key generation functionality $\mathcal{F}_{\mathsf{DKG}}^{t,k}[\mathbb{G}]$ (Cf. Fig. 1) and the secure channel functionality $\mathcal{F}_{\mathsf{sc}}$ (Cf. Fig. 3). The parties perform the voting protocol. When a voter V_i is coerced, he switches to the real deceiving strategy DR_i to resist coercion.

5. The Protocol

In this section, we give a detailed protocol description of our voting scheme and formally prove security of our scheme in the UC framework.

5.1. Protocol Description

Coercion-resistant voting $\Pi_{\mathsf{vote}}^{n,m,\ell,t,k}$

Denote the voters as $\mathcal{V} := \{\mathsf{V}_1, \ldots, \mathsf{V}_n\}$, the experts as $\mathcal{E} := \{\mathsf{E}_1, \ldots, \mathsf{E}_m\}$, the candidates as $\mathcal{C} := \{\mathsf{C}_1, \ldots, \mathsf{C}_\ell\}$, the registration authority as RA, the trustees as $\mathcal{T} := \{\mathsf{T}_1, \ldots, \mathsf{T}_k\}$. The protocol is parameterized with threshold t.

Preparation Phase:

Upon receiving $(\textsc{Init}, \mathsf{sid})$ from the environment $\mathcal{Z}$, the trustee T_i does the following:

- Send $(\textsc{KeyGen}, \mathsf{sid})$ to $\mathcal{F}_{\mathsf{DKG}}^{t,k}$ and receive $(\textsc{KeyGen}, \mathsf{sid}, (\mathsf{pk}_{\mathsf{T}}, \mathsf{sk}_{\mathsf{T},i}))$.
- Send $(\textsc{Write}, \mathsf{sid}, \mathsf{pk}_{\mathsf{T}})$ to $\mathcal{G}_{\mathsf{PBB}}$.

Upon receiving $(\textsc{Init}, \mathsf{sid})$ from the environment $\mathcal{Z}$, the RA does the following:

- Generate a public-private signing key-pair $(\mathsf{pk}_{\mathsf{RA}}, \mathsf{sk}_{\mathsf{RA}}) \leftarrow \mathsf{Sig.Keygen}(1^\lambda)$.
- Send $(\textsc{Write}, \mathsf{sid}, \mathsf{pk}_{\mathsf{RA}})$ to $\mathcal{G}_{\mathsf{PBB}}$.

Registration Phase:

Upon receiving $(\textsc{Reg}, \mathsf{sid}, \alpha_j)$ from the environment $\mathcal{Z}$, the voter V_j does the following:

- Generate a public-private signing key-pair $(\mathsf{pk}_j, \mathsf{sk}_j) \leftarrow \mathsf{Sig.Keygen}(1^\lambda)$.

Voting ideal functionality $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$

The functionality $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ interacts with a set of voters $\mathcal{V} := \{\mathsf{V}_1, \ldots, \mathsf{V}_n\}$, a set of experts $\mathcal{E} := \{\mathsf{E}_1, \ldots, \mathsf{E}_m\}$, a set of trustees $\mathcal{T} := \{\mathsf{T}_1, \ldots, \mathsf{T}_k\}$, the shuffler, the registration authority RA and the adversary $\mathcal{S}$. It is parameterized with candidate number ℓ and threshold t and it internally keeps variables $\mathcal{J}, \eta, \mathsf{state}, \mathsf{ballots}$. Denote $\mathcal{T}_{\mathsf{cor}}$ and $\mathcal{T}_{\mathsf{honest}}$ as the set of corrupted and honest trustees, respectively. Denote the candidates as $\mathcal{C}$.
Initially, $\mathcal{J} := \eta := \mathsf{ballots} := \emptyset$, $\mathsf{state} := 0$.

Preparation phase.
Upon receiving $(\textsc{Init}, \mathsf{sid})$ from the trustee T_i, send a notification message $(\textsc{InitNotify}, \mathsf{sid}, \mathsf{T}_i)$ to $\mathcal{S}$.
Upon receiving $(\textsc{Init}, \mathsf{sid})$ from the RA, send a notification message $(\textsc{InitNotify}, \mathsf{sid}, \mathsf{RA})$ to $\mathcal{S}$.

Registration phase.
When $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ receives $(\textsc{Init}, \mathsf{sid})$ from all trustees and RA, set $\mathsf{state} := 1$.
Upon receiving $(\textsc{Reg}, \mathsf{sid}, \alpha_j)$ from the voter V_j, $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ does the following: (ignore the request if $\mathsf{state} \neq 1$)

- Set $\mathsf{power}[\mathsf{V}_j] := \alpha_j$.
- If $|\mathcal{T} \cap \mathcal{T}_{\mathsf{cor}}| \geq t$, send $(\textsc{LeakPower}, \mathsf{sid}, \langle \alpha_i, \mathsf{V}_i \rangle)$ to $\mathcal{S}$.
- Otherwise, send a notification message $(\textsc{RegNotify}, \mathsf{sid}, \mathsf{V}_j)$ to the adversary $\mathcal{S}$.

Voting/Delegation phase:
When $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ receives $(\textsc{Reg}, \mathsf{sid})$ from all voters, set $\mathsf{state} := 2$.
Upon receiving $(\textsc{Vote}, \mathsf{sid}, v_i)$ from the voter V_i or the expert E_i, $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ does the following: (ignore the request if $\mathsf{state} \neq 2$)

- Set $\mathsf{ballots}[\mathsf{V}_i(\text{or } \mathsf{E}_i)] := v_i$.
- If the request is from a voter, send $(\textsc{VoteNotify}, \mathsf{sid}, \mathsf{VOTER})$ to $\mathcal{S}$. Otherwise, send $(\textsc{VoteNotify}, \mathsf{sid}, \mathsf{EXPERT})$ to $\mathcal{S}$.
- If $|\mathcal{T} \cap \mathcal{T}_{\mathsf{cor}}| \geq t$, send $(\textsc{LeakVote}, \mathsf{sid}, \langle v_i, \mathsf{V}_i(\text{or } \mathsf{E}_i) \rangle)$ to $\mathcal{S}$.

Tally phase.
Upon receiving $(\textsc{VoteEnd}, \mathsf{sid})$ from the shuffler, send $(\textsc{VoteEnd}, \mathsf{sid})$ to the adversary $\mathcal{S}$, and set $\mathsf{state} := 3$.
Upon receiving $(\textsc{Tally}, \mathsf{sid})$ from the trustee T_i, $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ does the following: (ignore the request if $\mathsf{state} \neq 3$)

- Set $\mathcal{J} := \mathcal{J} \cup \{\mathsf{T}_i\}$, if $|\mathcal{J} \cap \mathcal{T}_{\mathsf{honest}}| + |\mathcal{T}_{\mathsf{cor}}| \geq t$, compute $\eta \leftarrow \mathsf{TallyAlg}(\mathcal{V}, \mathcal{E}, \mathcal{C}, \mathsf{ballots}, \mathsf{power})$ (Cf. Fig. 12) and $\mathsf{bc}_{\mathsf{V}} \leftarrow \mathsf{CountAlg}(\mathcal{V}, \mathcal{E}, \mathcal{C}, \mathsf{ballots})$ (Cf. Fig. 13). Denote experts' ballots as $\mathsf{ballots}_{\mathsf{E}}$. Send $(\textsc{LeakTally}, \mathsf{sid}, \eta, \mathsf{ballots}_{\mathsf{E}}, \mathsf{bc}_{\mathsf{V}}))$ to the adversary $\mathcal{S}$.
- If $|\mathcal{J} \cap \mathcal{T}_{\mathsf{honest}}| + |\mathcal{T}_{\mathsf{cor}}| \geq t$ and the shuffler is corrupted, send $(\textsc{LeakBallots}, \mathsf{sid}, \mathsf{ballots})$ to the adversary $\mathcal{S}$.
- Send $(\textsc{TallyNotify}, \mathsf{sid}, \mathsf{T}_i)$ to adversary $\mathcal{S}$.
- If $|\mathcal{J}| \geq t$, set $\eta \leftarrow \mathsf{TallyAlg}(\mathcal{V}, \mathcal{E}, \mathcal{C}, \mathsf{ballots}, \mathsf{power})$ (Cf. Fig. 12).

Upon receiving $(\textsc{ReadTally}, \mathsf{sid})$ from any party, return $(\textsc{ReadTallyReturn}, \mathsf{sid}, \eta)$ to the requestor.

Figure 11. The voting ideal functionality $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$

The tally algorithm TallyAlg

Input: a set of the voters $\mathcal{V}$, a set of the experts $\mathcal{E}$, a set of candidates $\mathcal{C}$, a table of ballots ballots and a table of voting power power.
Output: the tally result η.
Use n, m, ℓ to denote $|\mathcal{V}|, |\mathcal{E}|, |\mathcal{C}|$, respectively.

Init:

- For each candidate C_i, set initial score as $s_i := 0$. For each expert E_j, set initial score as $s_{\ell+j} := 0$.

Tally Computation:

- For each $\mathsf{V}_j \in \mathcal{V}$, let $v := \mathsf{ballots}[\mathsf{V}_j]$, set $s_v := s_v + \mathsf{power}[\mathsf{V}_i]$.
- For each $\mathsf{E}_i \in \mathcal{E}$, let $v := \mathsf{ballots}[\mathsf{E}_i]$, set $s_v := s_v + s_{\ell+i}$.

Output:

- Return $\eta := \{s_i\}_{i \in [1,\ell]}$.

Figure 12. The tally algorithm

The ballot-counting algorithm CountAlg

Input: a set of the voters $\mathcal{V}$, a set of the experts $\mathcal{E}$, a set of candidates $\mathcal{C}$, a table of ballots ballots.
Output: voters' ballot count bc_{V}.
Use n, m, ℓ to denote $|\mathcal{V}|, |\mathcal{E}|, |\mathcal{C}|$, respectively.

Init:

- For each candidate C_i, set initial ballot count as $c_i := 0$. For each expert E_j, set initial ballot count as $c_{\ell+j} := 0$.

Computation:

- For each $\mathsf{V}_j \in \mathcal{V}$, let $v := \mathsf{ballots}[\mathsf{V}_j]$, set $c_v := c_v + 1$.

Output:

- Return $\mathsf{bc}_{\mathsf{V}} := \{c_i\}_{i \in [1,\ell+m]}$.

Figure 13. The ballot-counting algorithm

- Send $(\text{SEND}, \mathsf{sid}, \mathsf{RA}, \langle K_j, A_j, \mathsf{tx}_j, \delta_j \rangle)$ to $\mathcal{F}_{\mathsf{sc}}$, where $K_j \leftarrow \mathsf{EC.Enc}_{\mathsf{pk_T}}(\mathsf{pk}_j)$, $A_j \leftarrow \mathsf{LE.Enc}_{\mathsf{pk_T}}(\alpha_j)$, tx_j is the transaction that freezes the voter's stake, and $\delta_j \leftarrow \mathsf{NIZK}_{\mathsf{power}}.\mathsf{Prove}(A_j, \mathsf{tx}_j)$ is a NIZK proof that the frozen stake equals α_j (Cf. Appendix A).

Upon receiving $(\text{SENT}, \mathsf{sid}, \mathsf{V}_j, \langle K_j, A_j, \mathsf{tx}_j, \delta_j \rangle)$ from $\mathcal{F}_{\mathsf{sc}}$, the registration authority RA does the following:

- Compute $b := \mathsf{NIZK}_{\mathsf{power}}.\mathsf{Verify}(\delta_j, A_j || \mathsf{tx}_j)$. If $b = 0$, send $(\text{SEND}, \mathsf{sid}, \mathsf{V}_j, \text{REJECT})$ to $\mathcal{F}_{\mathsf{sc}}$.
- Else, compute $\tilde{K}_j \leftarrow \mathsf{EC.Rand}_{\mathsf{pk_T}}(K_j)$. Send $(\text{WRITE}, \mathsf{sid}, \langle \tilde{K}_j, A_j, \mathsf{tx}_j, \delta_j, \sigma_{\mathsf{RA},j} \rangle)$ to $\mathcal{G}_{\mathsf{PBB}}$, where $\sigma_{\mathsf{RA},j} \leftarrow \mathsf{Sig.sign}_{\mathsf{sk_{RA}}}(\tilde{K}_j || A_j || \mathsf{tx}_j || \delta_j)$.
- Generate a designated verifier proof π_{DVP} of re-encryption correctness (Cf. Appendix A) and send $(\text{SEND}, \mathsf{sid}, \mathsf{V}_j, \pi_{\mathsf{DVP}})$ to $\mathcal{F}_{\mathsf{sc}}$.

Voting/Delegation Phase:

Upon receiving $(\text{VOTE}, \mathsf{sid}, v_i)$ from the environment $\mathcal{Z}$, the expert E_i does the following:

- Compute $\boldsymbol{c}_i \leftarrow \mathsf{EC.Enc}_{\mathsf{pk_T}}(v_i)$ and the corresponding NIZK $\tau_i \leftarrow \mathsf{NIZK}_{\mathsf{knowledge}}.\mathsf{Prove}(\boldsymbol{c}_i)$, which is a proof of plaintext knowledge of $\boldsymbol{c}_i$ (Cf. Fig. 18).
- Compute $\sigma_{\mathsf{E}_i} \leftarrow \mathsf{Sig.sign}_{\mathsf{sk}_{\mathsf{E}_i}}(\boldsymbol{c}_i)$.
- Denote the ballot as $B_i^{\mathsf{E}} := (\mathsf{E}_i, \boldsymbol{c}_i, \tau_i, \sigma_{\mathsf{E}_i})$.
- Send $(\text{WRITE}, \mathsf{sid}, B_i^{\mathsf{E}})$ to $\mathcal{G}_{\mathsf{PBB}}$.

Upon receiving $(\text{VOTE}, \mathsf{sid}, v_j)$ from the environment $\mathcal{Z}$, the voter V_j does the following:

- Compute $\boldsymbol{u}_j \leftarrow \mathsf{EC.Enc}_{\mathsf{pk_T}}(v_j)$ and the corresponding NIZK $\pi_i \leftarrow \mathsf{NIZK}_{\mathsf{knowledge}}.\mathsf{Prove}(\boldsymbol{u}_i)$, which is a proof of plaintext knowledge of $\boldsymbol{u}_i$ (Cf. Fig. 18).
- Compute $\sigma_j \leftarrow \mathsf{Sig.sign}_{\mathsf{sk}_j}(\boldsymbol{u}_j)$.
- Denote the ballot as $B_j = (\mathsf{pk}_j, \boldsymbol{u}_j, \pi_j, \sigma_j)$.
- Send $(\text{WRITE}, \mathsf{sid}, B_j)$ to $\mathcal{G}_{\mathsf{PBB}}$.

Upon coerced, the voter V_j switches to the real deception strategy as described in Fig. 14.

Real Deceiving Strategy DR

The real deception strategy DR_i internally runs DI_i. Upon coerced, it does the following:

- Generate a fake signing key pair: $\mathsf{sk}'_j, \mathsf{pk}'_j \leftarrow \mathsf{Sig.Keygen}(1^\lambda)$.
- Compute $K'_j \leftarrow \mathsf{EC.Enc}_{\mathsf{pk_T}}(\mathsf{pk}'_j)$.
- Send $(\text{WRITE}, \mathsf{sid}, \langle K'_j, A_0, \rho_j \rangle)$ to $\mathcal{G}_{\mathsf{PBB}}$, where $A_0 = \mathsf{LE.Enc}_{\mathsf{pk_T}}(0; 0)$ and $\rho_j \leftarrow \mathsf{NIZK}_{\mathsf{sk}}.\mathsf{Prove}(K'_j)$ is a NIZK proof of knowledge of sk'_j (Cf. Fig. 20).
- Use $(\mathsf{sk}'_j, \mathsf{pk}'_j)$ to perform the voting procedure, following the coercer's instruction.
- Use $(\mathsf{sk}_j, \mathsf{pk}_j)$ to perform the voting procedure again, where the candidate v_i is the same as the one submitted by DI_i.

Figure 14. Real Deceiving Strategy DR

Tally Phase:

Upon receiving $(\text{VOTEEND}, \mathsf{sid})$ from the environment $\mathcal{Z}$, the shuffler does the following:

- Send $(\text{READ}, \mathsf{sid})$ to $\mathcal{G}_{\mathsf{PBB}}$ and get all the ballots and "encrypted public key items".
- For each voter's ballot $B_j = (\mathsf{pk}_j, \boldsymbol{u}_j, \pi_j, \sigma_j)$, check (i) $\mathsf{NIZK}_{\mathsf{knowledge}}.\mathsf{Verify}(\pi_j, \boldsymbol{u}_j) \stackrel{?}{=} 1$; (ii) $\mathsf{Sig.Verify}_{\mathsf{pk}_j}(\sigma_j, \boldsymbol{u}_j) \stackrel{?}{=} 1$. Remove the invalid ballots and remove π_j, σ_j. Now a voter's ballot β is of the form $(K, \boldsymbol{u})$.
- For each "encrypted public key item" $\langle \tilde{K}_j, A_j, \mathsf{tx}_j, \delta_j, \sigma_{\mathsf{RA},j} \rangle$ published by RA, check (i) $\mathsf{NIZK}_{\mathsf{power}}.\mathsf{Verify}(\delta_j, A_j || \mathsf{tx}_j) \stackrel{?}{=} 1$; (ii) $\mathsf{Sig.Verify}_{\mathsf{pk_{RA}}}(\sigma_{\mathsf{RA},j}, \tilde{K}_j || A_j || \mathsf{tx}_j || \delta_j) \stackrel{?}{=} 1$. Remove the invalid ones and remove $\mathsf{tx}_j, \delta_j, \sigma_{\mathsf{RA},j}$.
- For each "encrypted public key item" $\langle K'_j, A_0, \rho_j \rangle$ published by a voter, check $\mathsf{NIZK}_{\mathsf{knowledge}}.\mathsf{Verify}(\rho_j, K'_j) \stackrel{?}{=} 1$. Remove the invalid ones and remove ρ_j.
- Put all the valid "encrypted public key items" together. At this point, each item W is of the form (K, A), where K is the encrypted public key, A is the encrypted voting power.
- Verifiably shuffle re-encrypt the encrypted public key items (Cf. [31]).
- Send $(\text{WRITE}, \mathsf{sid}, \langle \{\beta\}, \{W\}, \pi \rangle)$ to $\mathcal{G}_{\mathsf{PBB}}$, where π is the proof of shuffle correctness (Cf. [31]).

Upon receiving $(\text{TALLY}, \mathsf{sid})$ from the environment $\mathcal{Z}$, the trustee $\mathsf{T}_t, t \in [k]$ does the following:

- Send $(\text{READ}, \mathsf{sid})$ to $\mathcal{G}_{\mathsf{PBB}}$ to get $\{\beta\}, \{W\}$.
- Jointly decrypt the public keys (i.e. K) in $\{W\}$ (Cf. [22]).
- For each ballot β, if the public key does not match any public key in $\{W\}$, drop it.
- Put the corresponding A together with $\boldsymbol{u}$, i.e., assume $\beta = (K_B, \boldsymbol{u})$ and $W = (K_W, A)$, and K_B, K_A are decrypted to the same public key, then we put A and $\boldsymbol{u}$ together to form a new item $I := (A, \boldsymbol{u})$.
- Send $(\text{WRITE}, \mathsf{sid}, \{I\})$ to $\mathcal{G}_{\mathsf{PBB}}$.
- Send $(\text{SHUFFLE}, \mathsf{sid})$ to the shuffler.

Upon receiving $(\text{SHUFFLE}, \mathsf{sid})$ from the trustees, the shuffler does the following:

- Send $(\text{READ}, \mathsf{sid})$ to $\mathcal{G}_{\mathsf{PBB}}$ and get $\{I\}$.
- Shuffle $\{\text{I}\}$ and send $(\text{WRITE}, \mathsf{sid}, \langle \{I'\}, \pi \rangle)$ to $\mathcal{G}_{\mathsf{PBB}}$, where π is the proof of shuffle correctness (Cf. [31]).
- Send $(\text{SHUFFLEEND}, \mathsf{sid})$ to all the trustees.

Upon receiving $(\text{SHUFFLEEND}, \mathsf{sid})$ from the shuffler, the trustee $\mathsf{T}_t, t \in [k]$ does the following:

- For each candidate C_i, set initial score as $\boldsymbol{s}_i := \mathsf{LE.Enc}_{\mathsf{pk_T}}(0)$. For each expert E_j, set initial score as $\boldsymbol{s}_{\ell+j} := \mathsf{LE.Enc}_{\mathsf{pk_T}}(0)$.

- For each item $I := (A, \boldsymbol{u})$, jointly decrypt $\boldsymbol{u}$ to v (Cf. [22]) and update candidate (or expert) v's score $\boldsymbol{s}_v := \boldsymbol{s}_v \cdot A$.
- After tallying all the voters' ballots, for each expert's ballot $B_i^{\mathsf{E}} := (\mathsf{E}_i, \boldsymbol{c}_i, \tau_i, \sigma_{\mathsf{E}_i})$, check (i) $\mathsf{NIZK}_{\mathsf{knowledge}}.\mathsf{Verify}(\tau_i, \boldsymbol{c}_i) \stackrel{?}{=} 1$; (ii) $\mathsf{Sig}.\mathsf{Verify}_{\mathsf{pk}_{\mathsf{E}_i}}(\sigma_{\mathsf{E}_i}, \boldsymbol{c}_i) \stackrel{?}{=} 1$. Remove the invalid ballots and $\tau_i, \sigma_{\mathsf{E}_i}$. Now an expert's ballot β^{E} is of the form $(\mathsf{E}_i, \boldsymbol{c}_i)$.
- Form a list of each expert's encrypted score and encrypted candidate, i.e., expert E_i's entry is of the form $(\boldsymbol{s}_{\ell+i}, \boldsymbol{c}_i)$.
- For each expert's entry $(\boldsymbol{s}, \boldsymbol{c})$, jointly decrypt $\boldsymbol{c}$ to v (Cf. [22]) and update candidate v's score $\boldsymbol{s}_v := \boldsymbol{s}_v \cdot \boldsymbol{s}$.
- For each candidate C_i, jointly decrypt the score $\boldsymbol{s}_i$ to s_i (Cf. [22]).
- Send $(\textsc{Write}, \mathsf{sid}, \{s_i\}_{i\in[\ell]})$ to $\mathcal{G}_{\mathsf{PBB}}$.

Upon receiving $(\textsc{ReadTally}, \mathsf{sid})$ from the environment $\mathcal{Z}$, the party P does the following:

- Send $(\textsc{Read}, \mathsf{sid})$ to $\mathcal{G}_{\mathsf{PBB}}$ and get $\{s_i\}_{i\in[\ell]}$.
- Return $(\textsc{ReadTallyReturn}, \mathsf{sid}, \{s_i\}_{i\in[\ell]})$ to the requestor.

5.2. Security

We show the security of our construction via the following theorem.

Theorem 1. *Assume that the NIZKs* $\mathsf{NIZK}_i, i \in \{\mathsf{power}, \mathsf{knowledge}, \mathsf{DVF\text{-}reenc}, \mathsf{sk}, \mathsf{shuffle}, \mathsf{Dec}\}$ *are complete, sound, and zero-knowledge. Assume that the ElGamal encryption scheme* EC *is IND-CPA secure. Assume that* Sig *is a signature scheme satisfying EUF-CMA. Then the protocol* $\Pi_{\mathsf{vote}}^{n,m,\ell,t,k}$ *IUC-realizes* $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ *against static corruption and adaptive coercion in the* $\{\mathcal{F}_{\mathsf{DKG}}^{t,k}[\mathbb{G}], \mathcal{F}_{\mathsf{sc}}, \mathcal{G}_{\mathsf{PBB}}\}$*-hybrid world.*

The proof of the theorem is deferred to Appendix C.

5.3. Discussion

Vote-buying via stake-buying. In blockchain voting, typically a voter's voting power is proportional to his stake. Since blockchain coins are publicly traded in open exchanges, one may argue that it is always possible to realize vote-buying via stake-buying. However, acquiring sufficient voting power through stake-buying is impractical. For an adversary to be successful, he needs to purchase a substantial amount of stake. Here's the catch: when buying from an exchange, there is often limited availability of stakes. Furthermore, rapidly purchasing large volumes of stake will inevitably drive up the price due to the basic principles of supply and demand. This surge in price could put the adversary's capital at risk. In contrast, vote-buying is a much simpler method and is detrimental to the decision-making process. Our coercion-resistant voting scheme effectively prevents vote-buying on the blockchain.

Stake renting/smart contract vote-buying. Another attack on blockchain voting is to use a smart contract to "rent" stakes. The smart contract collects stakes, uses the stakes for voting, and returns them back with an extra payment after the election. We defend this attack by prohibiting contract accounts from participating in the voting.

Inalienable authentication. Coercion-resistant voting requires inalienable authentication [11], i.e., the coercer can neither impersonate the voter nor prevent the voter from authenticating. In the blockchain context, in-person authentication is inappropriate. Instead, a voter authenticates to the RA by proving knowledge of his blockchain secret key and it is assumed that the voter will not give his secret key to the coercer. However, there is still an attack if the coercer can use TEEs. Specifically, the coercer can set up a TEE running a "cryptocurrency wallet" and use remote attestation to prove that the wallet will only do authentication and will not steal money. Then, a voter can input his secret key to this TEE in exchange for payments. This attack enables vote-selling without the coercer knowing the secret key. As pointed out in [32], this problem is inherent in any remote voting scheme where the secret key is generated by the voter. Kelkar *et al.* [33] proposed two schemes to defend such kind of attacks using TEEs and ASICs (Application-Specific Integrated Circuit), respectively. But neither of them is suitable in the blockchain context. In this work, we assume that the authentication is inalienable. How to defend this type of attack is out of the scope of this paper. We leave this as an interesting open problem.

Complexity. In the *preparation phase*, the RA takes $O(1)$ time to generate the signing key pair, and the trustees take $O(k)$ time to perform the DKG protocol, where k is the number of trustees. In the *registration phase*, a voter generates a NIZK proof of voting power correctness and verifies a designated verifier proof, which has $O(1)$ complexity. In the *voting/delegation phase*, a voter encrypts his choice and generates a NIZK proof of plaintext knowledge, which has $O(1)$ complexity, too. In the *tally phase*, the shuffler shuffles the ballots and the encrypted public key items, which has $O(n)$ complexity (counting cryptographic operations only), where n is the number of voters. Then, the trustees decrypt the public keys, do the matching between the voting power and candidates, and decrypt the candidates. Thus, the time complexity of a trustee is also $O(n)$. Adding them together, the whole scheme has $O(n)$ time complexity.

6. Implementation and Evaluation

We implement a prototype of our voting scheme in Rust. The implementation uses OpenSSL 1.1.1t to provide the basic elliptic curve math and it uses Schnorr signature as the signature scheme. We evaluated all the cryptographic building blocks and the time consumption in each phase. The experiments are performed on a workstation with Intel Core i7-1165G7 @2.80GHz and 32GB RAM running Ubuntu 20.04.4 LTS x64, using the elliptic curve secp256r1.

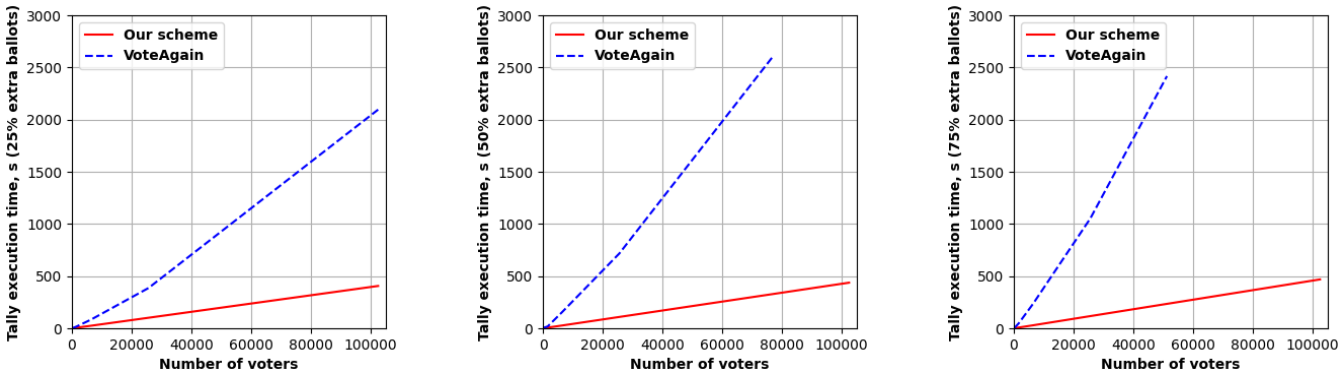

Figure 15. Comparison of tally execution time between our scheme and VoteAgain [13] (with extra ballot rate as 25%, 50%, 75% from left to right). In VoteAgain, $x\%$ extra ballot rate means that $x\%$ voters re-vote once; in our scheme, $x\%$ extra ballot rate means that $x\%$ voters cast one fake ballot.

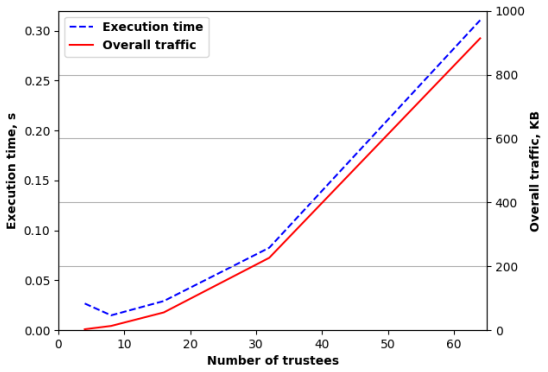

Figure 16. DKG execution time and overall traffic with respect to different numbers of trustees

Preparation phase. We evaluate the DKG execution time and traffic with respect to different numbers of trustees: from 4 to 64. Results are given in Fig. 16.

Registration phase. In the registration phase, the voter uses $\mathsf{NIZK}_{\mathsf{power}}$ in the registration request to prove that the frozen stake is equal to the encrypted voting power, and the RA proves re-encryption correctness by a designated verifier proof $\mathsf{NIZK}_{\mathsf{DVF\text{-}reenc}}$. Generating a registration request costs 430.95 µs and its size is 475 bytes. The designated verifier proof costs 102.91 µs to generate and 181.69 µs to verify, and its size is 141 bytes. Generating a fake "encrypted public key item" costs 336.05 µs and the size is 267 bytes.

Voting/Delegation phase. In the voting/delegation phase, voters and experts encrypt the choice, sign it, and use $\mathsf{NIZK}_{\mathsf{knowledge}}$ to prove plaintext knowledge of the ballot. It takes 283.27 µs to generate a ballot. The size of a voter's ballot is 358 bytes and the size of an expert's ballot is 332 bytes.

Tally phase. We evaluate the tally execution time with respect to different numbers of voters and different extra ballot rates and compare the results with VoteAgain [13]. Here, extra ballot rate represents how many voters cast extra ballots, i.e., in VoteAgain, 50% extra ballot rate means that 50% voters re-vote once; in our scheme, 50% extra ballot rate means that 50% voters cast one fake ballot. Note that, in our scheme, a voter's ballot takes more time to tally than an expert's ballot (because voters' ballots are shuffled and experts' ballots are not shuffled), so we set expert number as zero in the experiments.

Fig. 15 shows the tally execution time compared with VoteAgain [13] when the extra ballot rates are 25%, 50%, and 75% from left to right. VoteAgain's benchmark fails in 102400 voters with 50% and 75% extra ballot rates (probably because of too large ciphertext input). We can see that our scheme's execution time grows linearly and VoteAgain's execution time grows quasi-linearly ($O(n \log n)$). A higher rate of extra ballots confers a greater advantage, as it necessitates VoteAgain to introduce a substantial quantity of dummy ballots in this scenario. In large-scale voting with more than 10000 voters and over 50% extra ballot rate, our scheme's tally execution time is over 6x faster than VoteAgain.

7. Conclusion

In this work, we propose the first scalable coercion-resistant blockchain decision-making scheme that supports differential voting power and liquid democracy. It is scalable in the sense that it has constant ballot size and linear complexity. We formally prove the scheme secure under the UC incoercibility framework without any extra strong assumptions. Compared with existing voting schemes, our scheme has an advantage over all of them, so it is suitable to be applied to large-scale coercion-resistant blockchain voting programs.

References

[1] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, "Zerocash: Decentralized anonymous payments from bitcoin," in *2014 IEEE Symposium on Security and Privacy*. Berkeley, CA, USA: IEEE Computer Society Press, May 18–21, 2014, pp. 459–474.

[2] B. A. Ford, "Delegative democracy," Tech. Rep., 2002.

[3] B. Zhang, R. Oliynykov, and H. Balogun, "A treasury system for cryptocurrencies: Enabling better collaborative intelligence," in *ISOC Network and Distributed System Security Symposium – NDSS 2019*. San Diego, CA, USA: The Internet Society, Feb. 24–27, 2019.

[4] Snapshot, "Snapshot," online: https://snapshot.org (Last accessed: 2023-10-16).

[5] B. Adida, "Helios: Web-based open-audit voting," in *USENIX Security 2008: 17th USENIX Security Symposium*, P. C. van Oorschot, Ed. San Jose, CA, USA: USENIX Association, Jul. 28 – Aug. 1, 2008, pp. 335–348.

[6] P. Y. Ryan, D. Bismark, J. Heather, S. Schneider, and Z. Xia, "Prêt à voter: a voter-verifiable voting system," *IEEE transactions on information forensics and security*, vol. 4, no. 4, pp. 662–673, 2009.

[7] A. Juels, D. Catalano, and M. Jakobsson, "Coercion-resistant electronic elections," in *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, 2005, pp. 61–70.

[8] J. Clark and U. Hengartner, "Selections: Internet voting with over-the-shoulder coercion-resistance," in *FC 2011: 15th International Conference on Financial Cryptography and Data Security*, ser. Lecture Notes in Computer Science, G. Danezis, Ed., vol. 7035. Gros Islet, St. Lucia: Springer, Heidelberg, Germany, Feb. 28 – Mar. 4, 2012, pp. 47–61.

[9] R. Araújo, A. Barki, S. Brunet, and J. Traoré, "Remote electronic voting can be efficient, verifiable and coercion-resistant," in *FC 2016 Workshops*, ser. Lecture Notes in Computer Science, J. Clark, S. Meiklejohn, P. Y. A. Ryan, D. S. Wallach, M. Brenner, and K. Rohloff, Eds., vol. 9604. Christ Church, Barbados: Springer, Heidelberg, Germany, Feb. 26, 2016, pp. 224–232.

[10] K. Gjøsteen, "The norwegian internet voting protocol," in *E-Voting and Identity: Third International Conference, VoteID 2011, Tallinn, Estonia, September 28-30, 2011, Revised Selected Papers 3*. Springer, 2012, pp. 1–18.

[11] D. Achenbach, C. Kempka, B. Löwe, and J. Müller-Quade, "Improved coercion-resistant electronic elections through deniable re-voting," {*USENIX*} *Journal of Election Technology and Systems ({JETS})*, vol. 3, pp. 26–45, 2015.

[12] P. Locher, R. Haenni, and R. E. Koenig, "Coercion-resistant internet voting with everlasting privacy," in *FC 2016 Workshops*, ser. Lecture Notes in Computer Science, J. Clark, S. Meiklejohn, P. Y. A. Ryan, D. S. Wallach, M. Brenner, and K. Rohloff, Eds., vol. 9604. Christ Church, Barbados: Springer, Heidelberg, Germany, Feb. 26, 2016, pp. 161–175.

[13] W. Lueks, I. Querejeta-Azurmendi, and C. Troncoso, "VoteAgain: A scalable coercion-resistant voting system," in *USENIX Security 2020: 29th USENIX Security Symposium*, S. Capkun and F. Roesner, Eds. USENIX Association, Aug. 12–14, 2020, pp. 1553–1570.

[14] E. Magkos, M. Burmester, and V. Chrissikopoulos, "Receipt-freeness in large-scale elections without untappable channels," *Towards the E-Society: E-Commerce, E-Business, and E-Government*, pp. 683–693, 2001.

[15] J. Alwen, R. Ostrovsky, H.-S. Zhou, and V. Zikas, "Incoercible multi-party computation and universally composable receipt-free voting," in *Advances in Cryptology – CRYPTO 2015, Part II*, ser. Lecture Notes in Computer Science, R. Gennaro and M. J. B. Robshaw, Eds., vol. 9216. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 16–20, 2015, pp. 763–780.

[16] M. R. Clarkson, S. Chong, and A. C. Myers, "Civitas: Toward a secure voting system," in *2008 IEEE Symposium on Security and Privacy*. Oakland, CA, USA: IEEE Computer Society Press, May 18–21, 2008, pp. 354–368.

[17] S. Bursuc, G. S. Grewal, and M. D. Ryan, "Trivitas: Voters directly verifying votes," in *International Conference on E-Voting and Identity*. Springer, 2011, pp. 190–207.

[18] D. Chaum, "Random-sample voting," *White Paper*, 2016.

[19] Votem, "Votem," online: https://votem.com (Last accessed: 2023-10-16).

[20] Y. Li, W. Susilo, G. Yang, Y. Yu, D. Liu, X. Du, and M. Guizani, "A blockchain-based self-tallying voting protocol in decentralized iot," *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 1, pp. 119–130, 2020.

[21] J. Benet, "Ipfs-content addressed, versioned, p2p file system," *arXiv preprint arXiv:1407.3561*, 2014.

[22] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin, "Secure distributed key generation for discrete-log based cryptosystems," in *Advances in Cryptology – EUROCRYPT'99*, ser. Lecture Notes in Computer Science, J. Stern, Ed., vol. 1592. Prague, Czech Republic: Springer, Heidelberg, Germany, May 2–6, 1999, pp. 295–310.

[23] R. Canetti, "Universally composable security: A new paradigm for cryptographic protocols," in *42nd Annual Symposium on Foundations of Computer Science*. Las Vegas, NV, USA: IEEE Computer Society Press, Oct. 14–17, 2001, pp. 136–145.

[24] R. Canetti, Y. Dodis, R. Pass, and S. Walfish, "Universally composable security with global setup," in *TCC 2007: 4th Theory of Cryptography Conference*, ser. Lecture Notes in Computer Science, S. P. Vadhan, Ed., vol. 4392. Amsterdam, The Netherlands: Springer, Heidelberg, Germany, Feb. 21–24, 2007, pp. 61–85.

[25] D. Wikström, "Universally composable dkg with linear number of exponentiations," in *Security in Communication Networks: 4th International Conference, SCN 2004, Amalfi, Italy, September 8-10, 2004, Revised Selected Papers 4*. Springer, 2005, pp. 263–277.

[26] D. L. Chaum, "Untraceable electronic mail, return addresses, and digital pseudonyms," *Communications of the ACM*, vol. 24, no. 2, pp. 84–90, 1981.

[27] J. Chen and S. Micali, "Algorand: A secure and efficient distributed ledger," *Theoretical Computer Science*, vol. 777, pp. 155–183, 2019.

[28] L. Hirschi, L. Schmid, and D. A. Basin, "Fixing the achilles heel of E-voting: The bulletin board," in *CSF 2021: IEEE 34th Computer Security Foundations Symposium*, R. Küsters and D. Naumann, Eds. Virtual Conference: IEEE Computer Society Press, Jun. 21–24, 2021, pp. 1–17.

[29] V. Cortier and J. Lallemand, "Voting: You can't have privacy without individual verifiability," in *ACM CCS 2018: 25th Conference on Computer and Communications Security*, D. Lie, M. Mannan, M. Backes, and X. Wang, Eds. Toronto, ON, Canada: ACM Press, Oct. 15–19, 2018, pp. 53–66.

[30] M. Hirt and K. Sako, "Efficient receipt-free voting based on homomorphic encryption," in *Advances in Cryptology – EUROCRYPT 2000*, ser. Lecture Notes in Computer Science, B. Preneel, Ed., vol. 1807. Bruges, Belgium: Springer, Heidelberg, Germany, May 14–18, 2000, pp. 539–556.

[31] S. Bayer and J. Groth, "Efficient zero-knowledge argument for correctness of a shuffle," in *Advances in Cryptology – EUROCRYPT 2012*, ser. Lecture Notes in Computer Science, D. Pointcheval and T. Johansson, Eds., vol. 7237. Cambridge, UK: Springer, Heidelberg, Germany, Apr. 15–19, 2012, pp. 263–280.

[32] P. Daian, T. Kell, I. Miers, and A. Juels, "On-chain vote buying and the rise of dark daos," online: https://hackingdistributed.com/2018/07/02/on-chain-vote-buying/ (Last accessed: 2023-11-27).

[33] M. Kelkar, K. Babel, P. Daian, J. Austgen, V. Buterin, and A. Juels, "Complete knowledge: Preventing encumbrance of cryptographic secrets," *Cryptology ePrint Archive*, 2023.

[34] A. Fiat and A. Shamir, "How to prove yourself: Practical solutions to identification and signature problems," in *Advances in Cryptology – CRYPTO'86*, ser. Lecture Notes in Computer Science, A. M. Odlyzko, Ed., vol. 263. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 1987, pp. 186–194.

[35] J. Groth, "On the size of pairing-based non-interactive arguments," in *Advances in Cryptology – EUROCRYPT 2016, Part II*, ser. Lecture Notes in Computer Science, M. Fischlin and J.-S. Coron, Eds., vol. 9666. Vienna, Austria: Springer, Heidelberg, Germany, May 8–12, 2016, pp. 305–326.

[36] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell, "Bulletproofs: Short proofs for confidential transactions and more," in *2018 IEEE Symposium on Security and Privacy*. San Francisco, CA, USA: IEEE Computer Society Press, May 21–23, 2018, pp. 315–334.

[37] R. Cramer, I. Damgård, and B. Schoenmakers, "Proofs of partial knowledge and simplified design of witness hiding protocols," in *Advances in Cryptology – CRYPTO'94*, ser. Lecture Notes in Computer Science, Y. Desmedt, Ed., vol. 839. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 21–25, 1994, pp. 174–187.

Appendix A. NIZKs

In this section, we show the construction of the NIZKs used in our system. There are six zero-knowledge proofs in our scheme for proving: (i) voting power correctness ($\mathsf{NIZK}_{\mathsf{power}}$); (ii) ElGamal encryption plaintext knowledge ($\mathsf{NIZK}_{\mathsf{knowledge}}$); (iii) re-encryption correctness ($\mathsf{NIZK}_{\mathsf{DVF\text{-}reenc}}$); (iv) knowledge of secret key ($\mathsf{NIZK}_{\mathsf{sk}}$) (v) shuffle correctness ($\mathsf{NIZK}_{\mathsf{shuffle}}$); and (vi) decryption correctness ($\mathsf{NIZK}_{\mathsf{Dec}}$). We adopt Bayer and Groth's scheme [31] for shuffle correctness and Gennaro *et al.*'s scheme [22] for decryption correctness. Here, we will demonstrate how to use Sigma protocols to construct the other four zero-knowledge proofs. In practice, they will be transformed into NIZKs by Fiat-Shamir heuristic [34].

Proof of voting power correctness. In the registration phase, the voter will create a transaction that freezes some stake. Then, he sends the transaction tx, encrypted voting power A, and proves that the encrypted voting power is the same as the value of tx. In a privacy-preserving blockchain cryptocurrency system, tx usually contains an encrypted transaction value v. In this case, the zero-knowledge proof proves that A and v encrypt the same value. If the transaction value is encrypted by Lifted Elgamal, then Fig. 17 shows the Sigma protocol for voting power correctness. If it is encrypted by a hybrid encryption scheme (e.g., ZCash [1]), then we can utilize other zero-knowledge protocols for general circuits (e.g. Groth16 [35], Bulletproofs [36]).

Proof of ElGamal encryption plaintext knowledge. In the voting phase, we use a NIZK for ballot plaintext knowledge to prevent copying the other voter's choice. This can also be proven with Sigma protocol, depicted in Fig. 18.

Proof of re-encryption correctness. In the registration phase, the RA needs to generate a designated verifier proof for re-encryption correctness. To make it a designated verifier proof, the statement is "this is a correct re-encryption OR I know the verifier's blockchain secret key" so that the verifier can simulate the proof. The Sigma protocol for re-encryption correctness is depicted in Fig. 19. By the CDS composition [37], we can compose the Sigma protocol for re-encryption correctness and the standard Schnorr protocol to construct the designated verifier proof of re-encryption correctness.

Sigma protocol for voting power correctness

CRS: g, h, m.
Statement: $A = (A_1, A_2), v = (v_1, v_2)$.
Witness: α, r_1, r_2 such that $A = (g^{r_1}, g^{\alpha} h^{r_1}) \ \wedge \ v = (g^{r_2}, g^{\alpha} m^{r_2})$.

Prover:

- Pick random $\alpha', r_1', r_2' \leftarrow \mathbb{Z}_q$;
- Compute $a_1 := g^{r_1'}, a_2 := g^{\alpha'} h^{r_1'}, a_3 := g^{r_2'}, a_4 := g^{\alpha'} m^{r_2'}$;
- $P \rightarrow V$: a_1, a_2, a_3, a_4.

Verifier:

- $V \rightarrow P$: random $e \leftarrow \mathbb{Z}_q$.

Prover:

- Compute $z_1 := r_1' + e \cdot r_1, z_2 := r_2' + e \cdot r_2, z_3 := \alpha' + e \cdot \alpha$;
- $P \rightarrow V$: z_1, z_2, z_3.

Verifier:

- Output 1 if and only if the following holds:
  - $g^{z_1} = a_1 \cdot A_1^e$;
  - $g^{z_3} h^{z_1} = a_2 \cdot A_2^e$;
  - $g^{z_2} = a_3 \cdot v_1^e$;
  - $g^{z_3} h^{z_2} = a_4 \cdot v_2^e$.

Figure 17. Sigma protocol for voting power correctness

Sigma protocol for plaintext knowledge

CRS: g, h.
Statement: $c = (c_1, c_2)$.
Witness: m, r such that $c_1 = g^r \ \wedge \ c_2 = m \cdot h^r$.

Prover:

- Pick random $r' \leftarrow \mathbb{Z}_q, m' \leftarrow \mathbb{G}$;
- Compute $a_1 := g^{r'}, a_2 := m' \cdot h^{r'}$;
- $P \rightarrow V$: a_1, a_2.

Verifier:

- $V \rightarrow P$: random $e \leftarrow \mathbb{Z}_q$.

Prover:

- Compute $z_1 := r' + e \cdot r, z_2 := m' \cdot m^e$;
- $P \rightarrow V$: z_1, z_2.

Verifier:

- Output 1 if and only if the following holds:
  - $g^{z_1} = a_1 \cdot c_1^e$;
  - $z_2 \cdot h^{z_1} = a_2 \cdot c_2^e$.

Figure 18. Sigma protocol for ElGamal encryption plaintext knowledge

Sigma protocol for re-encryption correctness

CRS: g, h.
Statement: $u = (u_1, u_2), v = (v_1, v_2)$.
Witness: r such that $v_1 = u_1 \cdot g^r \ \wedge \ v_2 = u_2 \cdot h^r$.

Prover:

- Pick random $r' \leftarrow \mathbb{Z}_q$;
- Compute $a_1 = g^{r'}, a_2 := h^{r'}$;
- $P \rightarrow V$: a_1, a_2.

Verifier:

- $V \rightarrow P$: random $e \leftarrow \mathbb{Z}_q$.

Prover:

- Compute $z := r' + e \cdot r$;
- $P \rightarrow V$: z.

Verifier:

- Output 1 if and only if the following holds:
  - $g^z = a_1 \cdot (v_1/u_1)^e$;
  - $h^z = a_2 \cdot (v_2/u_2)^e$.

Figure 19. Sigma protocol for re-encryption correctness

Sigma protocol for knowledge of secret key

CRS: g, h.
Statement: $c = (c_1, c_2)$.
Witness: x, r such that $c_1 = g^r \ \wedge \ c_2 = g^x \cdot h^r$.

Prover:

- Pick random $r' \leftarrow \mathbb{Z}_q, x' \leftarrow \mathbb{G}$;
- Compute $a_1 := g^{r'}, a_2 := g^{x'} \cdot h^{r'}$;
- $P \rightarrow V$: a_1, a_2.

Verifier:

- $V \rightarrow P$: random $e \leftarrow \mathbb{Z}_q$.

Prover:

- Compute $z_1 := r' + e \cdot r, z_2 := x' + e \cdot x$;
- $P \rightarrow V$: z_1, z_2.

Verifier:

- Output 1 if and only if the following holds:
  - $g^{z_1} = a_1 \cdot c_1^e$;
  - $g^{z_2} \cdot h^{z_1} = a_2 \cdot c_2^e$.

Figure 20. Sigma protocol for knowledge of secret key

Proof of knowledge of secret key. To publish an (encrypted) fake voting public key item on the PBB, the voter needs to prove knowledge of the corresponding secret key. This is a variant of the Schnorr protocol, depicted in Fig. 20.

Appendix B.
Security Definitions of NIZK, Encryption, and Signature

Here, we give formal game-based definitions of completeness, soundness, zero-knowledge of a NIZK, IND-CPA property of an encryption scheme, and EUF-CMA property of a signature scheme.

NIZK. A non-interactive zero-knowledge proof (NIZK) for relation $\mathcal{R}$ has four PPT algorithms $(\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify}, \mathsf{Sim})$ such that
$(\sigma, \tau) \leftarrow \mathsf{Setup}(\mathcal{R})$: The setup algorithm outputs a common reference string σ and a simulation trapdoor τ for relation $\mathcal{R}$.
$\pi \leftarrow \mathsf{Prove}(\mathcal{R}, \sigma, \phi, w)$: the prover algorithm takes as input a common reference string σ and $(\phi, w) \in \mathcal{R}$ and outputs a proof π.
$0/1 \leftarrow \mathsf{Verify}(\mathcal{R}, \sigma, \phi, \pi)$: the verification algorithm takes as input a common reference string σ, a statement ϕ and a proof π, and it returns 0 or 1 for rejection or acceptance, respectively.
$\pi \leftarrow \mathsf{Sim}(\mathcal{R}, \tau, \phi)$: the simulation algorithm takes as input a simulation trapdoor τ and a statement ϕ, and it outputs a proof π.
Completeness. A NIZK protocol is complete if an honest prover can always successfully convince an honest verifier. Formally, for all $(\phi, w) \in \mathcal{R}$,

$$\Pr[(\sigma, \tau) \leftarrow \mathsf{Setup}(\mathcal{R}); \pi \leftarrow \mathsf{Prove}(\mathcal{R}, \sigma, \phi, w) : \mathsf{Verify}(\mathcal{R}, \sigma, \phi, \pi) = 1] = 1$$

Zero-knowledge. A proof is zero-knowledge if no other information is leaked except that the statement is true. Consider the following experiment:
Experiment $\mathrm{EXPT}^{\mathsf{zk}}_{\mathcal{A},\mathsf{NIZK}}(\lambda)$:

1) For a relation $\mathcal{R}$, $(\sigma, \tau) \leftarrow \mathsf{Setup}(\mathcal{R})$, $(\phi, w) \in \mathcal{R}$, the challenger computes $\pi_0 \leftarrow \mathsf{Prove}(\mathcal{R}, \sigma, \phi, w)$ and $\pi_1 \leftarrow \mathsf{Sim}(\mathcal{R}, \tau, \phi)$.
2) The challenger picks a random bit $b \in \{0, 1\}$.
3) $\mathcal{A}$ is given (σ, π_b) as input, and it outputs a guess bit $b' \in \{0, 1\}$.
4) If $b = b'$, output 1; otherwise, output 0.

A NIZK is zero-knowledge if the adversary $\mathcal{A}$'s advantage $\mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}}(\mathcal{A}, \lambda) := |2 \cdot \Pr[\mathrm{EXPT}^{\mathsf{zk}}_{\mathcal{A},\mathsf{NIZK}}(\lambda) = 1] - 1|$ is negligible in λ.
Soundness. A proof is sound if it is not possible for a prover to prove a false statement. Consider the following experiment:
Experiment $\mathrm{EXPT}^{\mathsf{sound}}_{\mathcal{A},\mathsf{NIZK}}(\lambda)$:

1) For a relation $\mathcal{R}$, $(\sigma, \tau) \leftarrow \mathsf{Setup}(\mathcal{R})$.
2) Given σ as input, $\mathcal{A}$ outputs (ϕ, π).

3) If $\mathsf{Verify}(\mathcal{R}, \sigma, \phi, \pi) = 1$ and $\phi \notin L_{\mathcal{R}}$, output 1; otherwise, output 0.

A NIZK is sound if the adversary $\mathcal{A}$'s advantage $\mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}}(\mathcal{A}, \lambda) := \Pr[\mathrm{EXPT}^{\mathsf{sound}}_{\mathcal{A},\mathsf{NIZK}}(\lambda) = 1]$ is negligible in λ.

Encryption scheme. An encryption scheme consists of three PPT algorithms $(\mathsf{Keygen}, \mathsf{Enc}, \mathsf{Dec})$. The ElGamal encryption scheme we used is IND-CPA secure. Formally, consider the following IND-CPA experiment:

Experiment $\mathrm{EXPT}^{\mathsf{IND\text{-}CPA}}_{\mathcal{A},\mathsf{EC}}(\lambda)$:

1) The challenger performs the key generation algorithm $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Keygen}(\lambda)$ and sends pk to the adversary $\mathcal{A}$.
2) $\mathcal{A}$ sends m_0, m_1 to the challenger.
3) The challenger picks a random bit $b \in \{0, 1\}$ and sends $c \leftarrow \mathsf{Enc}_{\mathsf{pk}}(m_b)$ to $\mathcal{A}$.
4) $\mathcal{A}$ outputs a guess bit $b' \in \{0, 1\}$. If $b = b'$, output 1; otherwise, output 0.

An encryption scheme is IND-CPA secure if the adversary $\mathcal{A}$'s advantage $\mathsf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathsf{EC}}(\mathcal{A}, \lambda) := |2 \cdot \Pr[\mathrm{EXPT}^{\mathsf{IND\text{-}CPA}}_{\mathcal{A},\mathsf{EC}}(\lambda) = 1] - 1|$ is negligible in λ.

Signature. A signature scheme consists of three PPT algorithms $(\mathsf{Keygen}, \mathsf{Sign}, \mathsf{Verify})$. We require the underlying signature scheme to be existentially unforgeable under chosen message attack (EUF-CMA). The EUF-CMA experiment is as follows:

Experiment $\mathrm{EXPT}^{\mathsf{EUF\text{-}CMA}}_{\mathcal{A},\mathsf{Sig}}(\lambda)$:

1) The challenger performs the key generation algorithm $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Keygen}(\lambda)$ and sends pk to the adversary $\mathcal{A}$.
2) $\mathcal{A}$ can repeatedly request for signatures on chosen messages $(m_0, \ldots, m_q)$, and receives the valid signatures $(\sigma_0, \ldots, \sigma_q)$ in response.
3) $\mathcal{A}$ outputs a message and signature (m^*, σ^*).
4) If m^* is not one of the messages requested in step 2, and $\mathsf{Verify}_{\mathsf{pk}}(m^*, \sigma^*) = 1$, output 1; otherwise, output 0.

A signature scheme is EUF-CMA if the adversary $\mathcal{A}$'s advantage $\mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig}}(\mathcal{A}, \lambda) := \Pr[\mathrm{EXPT}^{\mathsf{EUF\text{-}CMA}}_{\mathcal{A},\mathsf{Sig}}(\lambda) = 1]$ is negligible in λ.

Appendix C. Proof of Theorem 1

Theorem 1 *Assume that the NIZKs* $\mathsf{NIZK}_i, i \in \{\mathsf{power}, \mathsf{knowledge}, \mathsf{DVF\text{-}reenc}, \mathsf{sk}, \mathsf{shuffle}, \mathsf{Dec}\}$ *are complete, sound, and zero-knowledge. Assume that the ElGamal encryption scheme* EC *is IND-CPA secure. Assume that* Sig *is a signature scheme satisfying EUF-CMA. Then the protocol* $\Pi^{n,m,\ell,t,k}_{\mathsf{vote}}$ *IUC-realizes* $\mathcal{F}^{n,m,\ell,t,k}_{\mathsf{vote}}$ *against static corruption and adaptive coercion in the* $\{\mathcal{F}^{t,k}_{\mathsf{DKG}}[\mathbb{G}], \mathcal{F}_{\mathsf{sc}}, \mathcal{G}_{\mathsf{PBB}}\}$*-hybrid world.*

Proof. To prove the theorem, we construct the real deception strategies DR and a simulator $\mathcal{S}$ such that no non-uniform PPT environment $\mathcal{Z}$ can distinguish (i) the real execution $\mathrm{EXEC}^{\mathcal{F}^{t,k}_{\mathsf{DKG}}[\mathbb{G}], \mathcal{F}_{\mathsf{sc}}, \mathcal{G}_{\mathsf{PBB}}}_{\Pi^{n,m,\ell,t,k}_{\mathsf{vote}}, \mathsf{DR}, \mathcal{A}, \mathcal{Z}}$ from the (ii) the ideal execution $\mathrm{EXEC}^{\mathcal{G}_{\mathsf{PBB}}}_{\mathcal{F}^{n,m,\ell,t,k}_{\mathsf{vote}}, \mathsf{DI}, \mathcal{S}, \mathcal{Z}}$.

Real Deception Strategy. The real deception strategy DR_i internally runs DI_i, forwarding messages between DI_i and the environment $\mathcal{Z}$. DR_i performs as described in Fig. 14.

When the coercer asks for the view of the registration phase, DR_i generates $K' \leftarrow \mathsf{EC.Enc}_{\mathsf{pk}_\mathsf{T}}(\mathsf{pk}'_j)$. Then, the voter claims that $(K'_j, A_j, \mathsf{tx}_j, \delta_j)$ is the message sent to the RA by simulating the designated verifier proof of re-encryption.

Simulator. The simulator $\mathcal{S}$ internally runs $\mathcal{A}$, forwarding messages to and from the environment $\mathcal{Z}$. $\mathcal{S}$ simulates the voters $\mathcal{V} := \{\mathsf{V}_1, \ldots, \mathsf{V}_n\}$, the experts $\mathcal{E} := \{\mathsf{E}_1, \ldots, \mathsf{E}_m\}$, the registration authority RA, the shuffler, the trustees $\mathcal{T} := \{\mathsf{T}_1, \ldots, \mathsf{T}_k\}$, and the ideal functionalities $\mathcal{F}^{t,k}_{\mathsf{DKG}}[\mathbb{G}], \mathcal{F}_{\mathsf{sc}}$. It works as follows:

In the preparation phase:

Upon receiving $(\textsc{InitNotify}, \mathsf{sid}, \mathsf{T}_i)$ from the ideal functionality $\mathcal{F}^{n,m,\ell,t,k}_{\mathsf{vote}}$, $\mathcal{S}$ simulates the trustee T_i following the protocol $\Pi^{n,m,\ell,t,k}_{\mathsf{vote}}$ as if he receives $(\textsc{Init}, \mathsf{sid})$ from the environment $\mathcal{Z}$.

Upon receiving $(\textsc{InitNotify}, \mathsf{sid}, \mathsf{RA})$ from the ideal functionality $\mathcal{F}^{n,m,\ell,t,k}_{\mathsf{vote}}$, $\mathcal{S}$ simulates the RA following the protocol $\Pi^{n,m,\ell,t,k}_{\mathsf{vote}}$ as if he receives $(\textsc{Init}, \mathsf{sid})$ from the environment $\mathcal{Z}$.

In the registration phase:

Upon receiving $(\textsc{LeakPower}, \mathsf{sid}, \langle \alpha_i, \mathsf{V}_i \rangle)$ from the ideal functionality $\mathcal{F}^{n,m,\ell,t,k}_{\mathsf{vote}}$, $\mathcal{S}$ simulates the voter V_j following the protocol $\Pi^{n,m,\ell,t,k}_{\mathsf{vote}}$ as if he receives $(\textsc{Reg}, \mathsf{sid}, \alpha_j)$ from the environment $\mathcal{Z}$.

Upon receiving $(\textsc{RegNotify}, \mathsf{sid}, \mathsf{V}_j)$ from the ideal functionality $\mathcal{F}^{n,m,\ell,t,k}_{\mathsf{vote}}$, $\mathcal{S}$ simulates the voter V_j following the protocol $\Pi^{n,m,\ell,t,k}_{\mathsf{vote}}$ as if he receives $(\textsc{Reg}, \mathsf{sid}, 0)$ from the environment $\mathcal{Z}$.

When the the registration authority RA receives $(\textsc{Sent}, \mathsf{sid}, \mathsf{V}_j, \langle K_j, A_j, \mathsf{tx}_j, \delta_j \rangle)$ from the simulated $\mathcal{F}_{\mathsf{sc}}$, $\mathcal{S}$ simulates the RA following the protocol $\Pi^{n,m,\ell,t,k}_{\mathsf{vote}}$.

In the voting/delegation phase:

Upon receiving $(\textsc{VoteNotify}, \mathsf{sid}, \mathsf{VOTER})$ from the ideal functionality $\mathcal{F}^{n,m,\ell,t,k}_{\mathsf{vote}}$, $\mathcal{S}$ simulates an honest voter V_j following the protocol $\Pi^{n,m,\ell,t,k}_{\mathsf{vote}}$ as if he receives $(\textsc{Vote}, \mathsf{sid}, v_0)$ from the environment $\mathcal{Z}$.

Upon receiving $(\textsc{VoteNotify}, \mathsf{sid}, \mathsf{EXPERT})$ from the ideal functionality $\mathcal{F}^{n,m,\ell,t,k}_{\mathsf{vote}}$, $\mathcal{S}$ simulates an honest expert E_i following the protocol $\Pi^{n,m,\ell,t,k}_{\mathsf{vote}}$ as if he receives $(\textsc{Vote}, \mathsf{sid}, v_0)$ from the environment $\mathcal{Z}$.

When a voter V_i is coerced, $\mathcal{S}$ simulates V_i performing the real deceiving strategy DR_i to cast a fake ballot following the coercer's instructions.

Once $\mathcal{G}_{\mathsf{PBB}}$ receives $(\textsc{Write}, \mathsf{sid}, B^{\mathsf{E}}_i)$, $\mathcal{S}$ does the following:

- Parse B^{E}_i as $(\mathsf{E}_i, \boldsymbol{c}_i, \tau_i, \sigma_{\mathsf{E}_i})$.
- Check (i) $\mathsf{NIZK}_{\mathsf{knowledge}}.\mathsf{Verify}(\tau_i, \boldsymbol{c}_i) \stackrel{?}{=} 1$; (ii) $\mathsf{Sig.Verify}_{\mathsf{pk}_{\mathsf{E}_i}}(\sigma_{\mathsf{E}_i}, \boldsymbol{c}_i) \stackrel{?}{=} 1$.

- If it is valid, compute $v_i = \mathsf{EC.Dec}_{\mathsf{sk_T}}(\boldsymbol{c}_i)$;
- If E_i is not corrupted, $\mathcal{S}$ will abort.
- Send $(\textsc{Vote}, \mathsf{sid}, v_i)$ to $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ on behalf of E_i.

Once $\mathcal{G}_{\mathsf{PBB}}$ receives $(\textsc{Write}, \mathsf{sid}, B_j)$, $\mathcal{S}$ does the following:

- Parse B_j as $(\mathsf{pk}_j, \boldsymbol{u}_j, \pi_j, \sigma_j)$.
- Check (i) $\mathsf{NIZK_{knowledge}.Verify}(\pi_j, \boldsymbol{u}_j) \stackrel{?}{=} 1$; (ii) $\mathsf{Sig.Verify}_{\mathsf{pk}_j}(\sigma_j, \boldsymbol{u}_j) \stackrel{?}{=} 1$.
- If it is valid, compute $v_j = \mathsf{EC.Dec}_{\mathsf{sk_T}}(\boldsymbol{u}_i)$;
- Find the owner of pk_j by decrypting all the registration messages. Denote him as V_j.
- If V_j is not corrupted, $\mathcal{S}$ will abort.
- Send $(\textsc{Vote}, \mathsf{sid}, v_j)$ to $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$ on behalf of V_j.

In the tally phase:

Upon receiving $(\textsc{VoteEnd}, \mathsf{sid})$ from the ideal functionality $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$, $\mathcal{S}$ simulates the shuffler following the protocol $\Pi_{\mathsf{vote}}^{n,m,\ell,t,k}$ as if he receives $(\textsc{VoteEnd}, \mathsf{sid})$ from the environment $\mathcal{Z}$.

Upon receiving $(\textsc{LeakTally}, \mathsf{sid}, \eta, \mathsf{ballots_E}, \mathsf{bc_V})$ from the ideal functionality $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$, $\mathcal{S}$ records $\eta, \mathsf{ballots_E}, \mathsf{bc_V}$.

Upon receiving $(\textsc{LeakBallots}, \mathsf{sid}, \mathsf{ballots})$ from the ideal functionality $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$, $\mathcal{S}$ records $\mathsf{ballots}$.

Upon receiving $(\textsc{TallyNotify}, \mathsf{sid}, \mathsf{T}_i)$ from the ideal functionality $\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}$, $\mathcal{S}$ does the following:

- Set $\mathcal{J} := \mathcal{J} \cup \{\mathsf{T}_i\}$.
- If $|\mathcal{J} \cap \mathcal{T}_{\mathsf{honest}}| + |\mathcal{T}_{\mathsf{cor}}| < t$, $\mathcal{S}$ simulates the trustee T_i following the protocol $\Pi_{\mathsf{vote}}^{n,m,\ell,t,k}$ as if he receives $(\textsc{Tally}, \mathsf{sid})$ from the environment $\mathcal{Z}$.
- Otherwise, $\mathcal{S}$ simulates T_i's decryption and the corresponding NIZK based on the tally result η and the known information about ballots.

Indistinguishability.

We prove indistinguishability through a series of hybrid worlds $\mathcal{H}_0, \ldots, \mathcal{H}_6$.

Hybrid $\mathcal{H}_0$: This is the real world execution $\mathsf{EXEC}^{\mathcal{F}_{\mathsf{DKG}}^{t,k}[\mathbb{G}], \mathcal{F}_{\mathsf{sc}}, \mathcal{G}_{\mathsf{PBB}}}_{\Pi_{\mathsf{vote}}^{n,m,\ell,t,k}, \mathsf{DR}, \mathcal{A}, \mathcal{Z}}$.

Hybrid $\mathcal{H}_1$: $\mathcal{H}_1$ is the same as $\mathcal{H}_0$ except that during the tally phase, the honest trustees' decryption NIZKs are generated by the NIZK simulator.

Claim 1: If $\mathsf{NIZK_{Dec}}$ is zero-knowledge with adversary advantage $\mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK_{Dec}}}(\mathcal{A}, \lambda)$, then $\mathcal{H}_1$ and $\mathcal{H}_0$ are indistinguishable with distinguishing advantage at most $(4n + m + \ell) \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK_{Dec}}}(\mathcal{A}, \lambda)$.

Proof 1: With each voting casting at most one fake ballot in our modeling, there are at most $2n$ voters' ballots and $2n$ encrypted public key items. For experts and candidates, each of them has one ciphertext to decrypt. Therefore, the overall advantage is at most $(4n+m+\ell) \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK_{Dec}}}(\mathcal{A}, \lambda)$ by a standard hybrid argument.

Hybrid $\mathcal{H}_2$: $\mathcal{H}_2$ is the same as $\mathcal{H}_1$ except that during the tally phase, the honest trustees' decryption shares are backward calculated from the tally result.

Claim 2: $\mathcal{H}_2$ and $\mathcal{H}_1$ are perfectly indistinguishable.

Proof 2: In out threshold cryptosystem, the backward calculated shares in $\mathcal{H}_2$ and the shares in $\mathcal{H}_1$ have the same distribution.

Hybrid $\mathcal{H}_3$: $\mathcal{H}_3$ is the same as $\mathcal{H}_2$ except that in the voting/delegation phase, the honest voters' ballots are replaced with ballots for candidate v_0.

Claim 3: If the encryption scheme EC is IND-CPA with advantage $\mathsf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathsf{EC}}(\mathcal{A}, \lambda)$ and $\mathsf{NIZK_{knowledge}}$ is zero-knowledge with adversary advantage $\mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK_{knowledge}}}(\mathcal{A}, \lambda)$, then $\mathcal{H}_3$ and $\mathcal{H}_2$ are indistinguishable with distinguishing advantage at most $n \cdot \mathsf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathsf{EC}}(\mathcal{A}, \lambda) + n \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK_{knowledge}}}(\mathcal{A}, \lambda)$.

Proof 3: With at most n honest voters in the system, the overall advantage is at most $n \cdot \mathsf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathsf{EC}}(\mathcal{A}, \lambda) + n \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK_{knowledge}}}(\mathcal{A}, \lambda)$ by a standard hybrid argument.

Hybrid $\mathcal{H}_4$: $\mathcal{H}_4$ is the same as $\mathcal{H}_3$ except that if a corrupted voter generates a valid ballot for an honest voter, the execution will abort.

Claim 4: If the signature scheme Sig is EUF-CMA with advantage $\mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig}}(\mathcal{A}, \lambda)$, then $\mathcal{H}_4$ and $\mathcal{H}_3$ are indistinguishable with distinguishing advantage at most $n \cdot \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig}}(\mathcal{A}, \lambda)$.

Proof 4: There are at most n honest voters, so the probability of abortion is no more than $n \cdot \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig}}(\mathcal{A}, \lambda)$ by a standard hybrid argument.

Hybrid $\mathcal{H}_5$: $\mathcal{H}_5$ is the same as $\mathcal{H}_4$ except that if a corrupted expert generates a valid ballot for an honest expert, the execution will abort.

Claim 5: If the signature scheme Sig is EUF-CMA with advantage $\mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig}}(\mathcal{A}, \lambda)$, then $\mathcal{H}_5$ and $\mathcal{H}_4$ are indistinguishable with distinguishing advantage at most $m \cdot \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig}}(\mathcal{A}, \lambda)$.

Proof 5: Same as the previous proof, the probability of abortion is no more than $m \cdot \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig}}(\mathcal{A}, \lambda)$ by a standard hybrid argument.

Hybrid $\mathcal{H}_6$: This is the ideal execution $\mathsf{EXEC}^{\mathcal{G}_{\mathsf{PBB}}}_{\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}, \mathsf{DI}, \mathcal{S}, \mathcal{Z}}$.

Claim 6: If the decryption NIZK $\mathsf{NIZK_{Dec}}$ is sound with adversary advantage $\mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK_{Dec}}}(\mathcal{A}, \lambda)$ and the shuffle NIZK $\mathsf{NIZK_{shuffle}}$ is sound with adversary advantage $\mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK_{shuffle}}}(\mathcal{A}, \lambda)$, then $\mathcal{H}_6$ and $\mathcal{H}_5$ are indistinguishable with distinguishing advantage at most $(4n + m + \ell) \cdot \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK_{Dec}}}(\mathcal{A}, \lambda) + \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK_{shuffle}}}(\mathcal{A}, \lambda)$.

Proof 6: It suffices to argue that the ideal tally and the real tally output the same result. We can see that, as long as the re-encryption in the registration phase is correct and the shuffle and decryption in the tally phase are sound, the ideal tally and the real tally did the same computation by the additive homomorphism of lifted ElGamal encryption scheme. Thus, the overall advantage is no more than $n \cdot \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK_{DVP\text{-}reenc}}}(\mathcal{A}, \lambda) + (4n + m + \ell) \cdot \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK_{Dec}}}(\mathcal{A}, \lambda) + \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK_{shuffle}}}(\mathcal{A}, \lambda)$.

Combining together, the real execution $\mathsf{EXEC}^{\mathcal{F}_{\mathsf{DKG}}^{t,k}[\mathbb{G}], \mathcal{F}_{\mathsf{sc}}, \mathcal{G}_{\mathsf{PBB}}}_{\Pi_{\mathsf{vote}}^{n,m,\ell,t,k}, \mathsf{DR}, \mathcal{A}, \mathcal{Z}}$ and the ideal execution $\mathsf{EXEC}^{\mathcal{G}_{\mathsf{PBB}}}_{\mathcal{F}_{\mathsf{vote}}^{n,m,\ell,t,k}, \mathsf{DI}, \mathcal{S}, \mathcal{Z}}$ are indistinguishable with distinguishing advantage at most

$$
\begin{aligned}
&(4n + m + \ell) \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{Dec}}}(\mathcal{A}, \lambda) + n \cdot \mathsf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathsf{EC}}(\mathcal{A}, \lambda) \\
&+ n \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{knowledge}}}(\mathcal{A}, \lambda) + (n + m) \cdot \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig}}(\mathcal{A}, \lambda) \\
&+ n \cdot \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{DVP\text{-}reenc}}}(\mathcal{A}, \lambda) + (4n + m + \ell) \cdot \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{Dec}}}(\mathcal{A}, \lambda) \\
&+ \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{shuffle}}}(\mathcal{A}, \lambda)
\end{aligned}
$$

This concludes the proof.

□