Votexx

Extreme Coercion Resistance

David Chaum[1], Richard T. Carback[1], Mario Yaksetig[1], Jeremy Clark[2],
Mahdi Nejadgholi[2], Bart Preneel[3], Alan T. Sherman[4], Filip Zagorski[5],
Bingsheng Zhang[6] and Zeyuan Yin[6]

[1] xx network, USA
[2] Concordia University, Canada
[3] KU Leuven, Belgium
[4] Cyber Defense Lab, University of Maryland, Baltimore County (UMBC), USA
[5] University of Wroclaw, Poland
[6] Zhejiang University, China

Abstract. We provide a novel perspective on a long-standing challenge to the integrity of votes cast without the supervision of a voting booth: "*improper influence*," which we define as any combination of vote buying and voter coercion. In comparison with previous proposals, our system is the first in the literature to protect against a strong adversary who learns all of the voter's keys—we call this property *"extreme coercion resistance."* When keys are stolen, each voter, or their trusted agents (which we call "*hedgehogs*"), may *"nullify"* (effectively cancel) their vote in a way that is unstoppable and irrevocable, and such that the nullification action is forever unattributable to that voter or their hedgehog(s). We demonstrate the security of our VoteXX system in the universal composability model.

As in many other coercion-resistant systems, voters are authorized to vote with public-private keys. Each voter registers their public keys with the *Election Authority (EA)* in a way that convinces the EA that the voter has memorized a passphrase that corresponds to their private keys. As a consequence, if an adversary obtains a voter's keys, the voter also retains a copy. Voters concerned about adversaries stealing their private keys can themselves, or by delegating to one or more untrusted hedgehog(s), monitor the bulletin board for malicious ballots cast with their keys, and can act to nullify these ballots in a privacy-preserving manner with zero-knowledge proofs.

In comparison with previous proposals, our system offers some protection against even the strongest adversary who learns all keys. Other coercion-resistant protocols either do not address these attacks, place strong limitations on adversarial abilities, or rely on fully trusted parties to assist voters with their keys.

Keywords: coercion resistance · high-integrity voting · voter-verifiable elections

1 Introduction

For over 150 years, the voting booth helped prevent voters from being bribed and coerced. For example, a controlling family member might coerce a voter by observing them vote, if votes are cast online from home or by mail. The booth, however, is becoming untenable as information technology provides the means for people to vote more frequently and conveniently without booths, including using combinations of mailed paper forms and online interactions. Moreover, growing use of technology facilitates vote buying and voter

Table 1: Properties of related work for resisting improper influence in online *end-to-end (E2E)* verifiable elections. Properties are fully present (●), partially present (◐), or not present (○). It is best to receive ● for each property. Decoy ballots act indirectly against influence (◉).

Influencer: System resists coercion when the influencer: (0) acts before/during registration; (1) colludes with the EA; (2) colludes with hardware manufactures; (3) acts at any time; (4) learns all information stored by the voter, including all keys required by the protocol (*i.e.,* mitigates extreme coercion); (5) learns every action taken by the voter. **Other:** (6) voter can undo coercion undetectably; (7) system is inexpensive; (8) system has low cognitive burden; (9) system has security proof (none/game-based/UC).

		0	1	2	3	4	5	6	7	8	9
Type	**Example**		Influencer						Other		
Baseline (coercible)	Helios (2008) [Adi08]	○	○	○	○	○	○	○	●	●	◐
Fake credentials	JCJ (2005) [JCJ05]	○	◐	●	○	○	○	●	●	●	◐
Masked ballots	WeBu09 (2009) [WB09]	○	●	●	○	○	○	○	●	○	●
Panic passwords	Selections (2011) [CH11]	○	●	●	○	○	○	●	●	◐	◐
Decoy ballots	RS-Voting (2012) [Cha12]	●	◐	●	◉	◉	○	○	●	◐	●
Secure hardware	AOZZ (2015) [AOZZ15]	●	◐	○	○	◐	○	●	○	●	●
Re-voting (E2E)	VoteAgain (2020) [LQAT20]	●	○	●	○	○	○	●	●	●	◐
Hedgehogs	VoteXX (2022)	●	●	●	●	●	●	○	○	●	●

coercion with electronic payments, live video streaming from voter phones, and various types of online threats.

Three daunting challenges make Internet voting difficult: (1) The lack of a secure physical voting precinct facilitates improper influence, including vote selling and coercion. (2) Malware on a voter's device (e.g., phone) might undetectably modify votes and spy on the voter. (3) Determined adversaries might try to launch an online attack, including causing outages. Of these challenges, the most elusive has been mitigating improper influence.

We contribute a solution to the problem of *improper influence* in voting without booths that enables any voter to *"nullify"* (effectively cancel) their vote in a way that is unstoppable, irrevocable, and forever unattributable to that voter. Importantly, our protocol can be accomplished even when an adversary gains access to all private keys of voters. Our approach allows each voter to recruit one or more trusted agents, which we call "*hedgehogs.*" The voter, or their hedgehog(s), can nullify the vote by proving knowledge of the voter's private key using a *zero-knowledge proof (ZKP)* without revealing the private key. This paper provides details for these ideas, which we introduced in a four-page abstract in 2022 [RED22].

Hedgehogs can be recruited before or during the election, from the voter's acquaintances or using a service selected on reputation. Hedgehogs can prove to the voter that they perform their services correctly. We call a "coercer" any entity who obtains a voter's key by coercion or bribery, whereas a "hedgehog" is an entity the voter trusts and to whom the voter voluntarily provides the key to protect the voter against coercers.

We accept that certain types of coercion are impossible to prevent in practice: a coercer can generally block a voter from registering for an election, and if a coercer posses all knowledge and attributes of the voter, they cannot be distinguished from the voter. Our approach differs from previous approaches with end-to-end verifiability (see Section 3 and Table 1)—e.g., revoting, fake credentials, and decoy ballots—by protecting against what we believe to be the strongest possible adversarial model that can be realistically protected against. Specifically we assume adversaries can learn all voter secrets and observe all voter

interactions with the system, excluding interactions with the hedgehogs which distinguish the voter from the coercer. We call this protection "*extreme coercion resistance.*"

Our mechanisms work beyond coercion: they protect voters against any form of loss or theft of their secret keys.

Some may yearn for an ideal world in which every voter can vote and vote their desires, but, unfortunately, the imperfect reality is that, in any system, a powerful coercer can always prevent a voter from voting. Nullification is a useful form of coercion resistance, and VoteXX achieves this form of coercion resistance, which is the best possible given the constraints of reality. There is no widely agreed upon well-defined notion of coercion resistance, and some authors make the weak adversarial assumption that coercers cannot prevent voters from voting. Our work explores what can be done under very strong adversarial assumptions.

This paper presents an architecture, design, implementation, and *universal composability (UC)* [Can01] security proof (see Appendix C) of our voting system, called *VoteXX*. The main feature of VoteXX is that it protects against extreme coercion, which we formally define in terms of UC ideal functionality (see Section 4.2). Our VoteXX protocols include comprehensive mechanisms to handle all of the security requirements, including, for example, *inalienable authentication*, which many other voting systems simply assume without providing constructions. We describe the user experience for several settings, which experiences are intuitive and require few steps. We have implemented the entire VoteXX system and made all of our sourcecode publicly available as an artifact (see Section 7). Performance analysis and benchmarking show that the system is highly practical (see also Section 8.3).

Our primary contributions are: (1) We introduce the new notions of nullification and hedgehogs, and present a new solution to improper influence based on them. (2) We give cryptographic protocols realizing nullification, and show how nullification can be applied to several voting settings, including vote-by-mail and online. (3) We present a new fully-decentralized scalable voting system, *VoteXX*, including registration, voting, nullification, and tallying. (4) We describe our implementation of VoteXX, which uses an *anonymous communication system (ACS)* for registration, vote casting, and other communication. (5) We provide a formal statement and UC proof of VoteXX's ballot secrecy, coercion resistance, and tally integrity. In addition, while other systems complicate registration and vote casting, our approach allows simple registration and vote casting by keeping nullification separate. Consequently, our system can be used as an overlay in conjunction with other coercion-resistant approaches—such as re-voting and decoy ballots—or in conjunction with other voting systems, such as ElectionGuard [BN23].

In the rest of this paper, we explain the notion of extreme coercion resistance, compare our approach with those of previous work, detail our adversarial model, give our problem specification, show the VoteXX architecture, define the VoteXX cryptographic protocols, describe voter interfaces for several settings including vote-by-mail and online, mention possible extensions to VoteXX, sketch the VoteXX implementation and discuss its performance, and explain the significance of our work. Appendix C gives our UC proofs. Throughout, we use the terms "coercion" and "improper influence" synonymously.

2 Extreme Coercion Resistance

As pointed out by Hirt and Sako [Hir01], a basic axiom of coercion resistance is that, if the adversary has/knows/is everything that the voter has/knows/is, the voting system cannot distinguish the voter from the adversary, and therefore cannot provide a mechanism to register the true voter's vote and not the adversary's. The most common path toward providing coercion resistance has been to assume a voter can maintain some secret value for authentication that cannot be stolen by the adversary.

Table 2: VoteXX countermeasures against various forms of key loss. Traditional coercion-resistant voting strategies consider the first row ("normal conditions"), whereas we argue all rows of key loss (see Eyal [Eya21]) should be addressed. We define *extreme coercion resistance* as addressing all rows, and introduce VoteXX as the first voting system that comprehensively mitigates all of these forms of key loss, leak, or theft.

Type of key loss	User has key?	Adversary has key?	Countermeasure in VoteXX
Normal conditions	Yes	No	Coercion resistance
Lost secret	No	No	Inalienable authentication
Leaked secret	Yes	Yes	Nullification
Stolen secret	No	Yes	Inalienable authentication

For fully online voting, maintaining secrets is particularly challenging because digital forms of secrets are generally "something you have/know" rather than "something you are," and secrets can be stolen through a variety of covert means, including malware and physical access to the voter's device. Directly stealing a secret side-steps known coercion-resistant mechanisms—each assumes somewhere that the voter knows of, and possibly plays a role in, resisting coercion (such as providing a fake secret or by voting again with their real secret).

We have used the term "stolen" loosely. In fact, there are three specific failures a user can experience in maintaining a secret value [Eya21]: secrets can be *lost*, where neither the user or adversary has the credential; *leaked*, where the user and adversary both have the credential; or *stolen*, where the user loses the credential and the adversary gains the credential (so only the adversary has the credential). VoteXX deals with each case as summarized in Table 2.

To provide coercion resistance in normal conditions (Row 1), VoteXX can operate as an overlay (that provides Rows 2-4) for many coercion-resistant mechanisms. For pedagogical reasons, we present nullification and inalienable authentication in the context of a full voting system; however, alternatively, VoteXX mechanisms (nullification and inalienable authentication) can be added to other coercion-resistant mechanisms to cover lost/leaked/stolen keys.

For lost and stolen secrets (Rows 2 and 4), the voter loses their secret. Our countermeasure is to ensure voters cannot lose their secret through *inalienable authentication* [CM22], which is a protocol where the voter proves that they have committed their secret to their own human memory without revealing the secret.

The threat of a leaked credential (Row 3) is underappreciated in the voting literature. It is easy to think that coercion resistance already covers this attack; however, coercion resistance addresses only *overt* attempts by a coercer to learn the key, whereas we also need to consider *covert* attacks for which voters may be unaware.

At first sight extreme coercion resistance seems an unsolvable problem: if the adversary knows every secret the voter knows, the remote system cannot distinguish the voter and adversary. The truth of this unpleasant fact implies that a dimension of the problem indeed is unsolvable: the voter cannot always cast their true intent and prevent the adversary from doing the same. Nevertheless, the voter still has a course of action: the voter (or either party) can sabotage (*nullify*) the vote, preventing themselves and the adversary from casting a ballot. Further, because we cannot assume that a voter will know their key was compromised, the election bulletin board can be monitored (by the voter or a recruited hedgehog) for adversarial action involving the voter's key.

It would miss a crucial point to dismiss VoteXX simply on the grounds that it "supports a forced abstention attack." In our strong model, for any voting system, a coercer can prevent any voter from registering or voting. VoteXX offers strategies summarized in Table 2 for mitigating all forms of key loss, which strategies in comparison with those of previous

coercion-resistant mechanisms, offer additional benefits, provide a more comprehensive solution, and work against a stronger adversary.

3   Comparison to Previous Work

When comparing coercion-resistance systems, it is revealing to consider how they deal with the various ways credential can be lost, leaked, or stolen. Most papers do not consider Rows 2, 3, and 4 at all from Table 2, and instead concentrate on only the Row 1 ("normal conditions") of the table, which is challenging enough.

Coercion resistance guarantees that each voter may vote freely. Informally, a voting system is *coercion resistant* if and only if no ballot is "*counted as coerced*," that is, no voter can prove to any coercer that the voter cast a counted ballot according to the coercer's instructions. As explained in Sections 4.2 and C.2, we uniquely adopt a very strong form of coercion resistance, which we call *extreme coercion*, in which the coercer learns all of the voter's keys. By contrast, other researchers assume only weaker forms of coercion, such as *semi-honest coercion* (receipt-freeness) in which the the voter must follow the voting protocol (see Table 1), or *active coercion* in which the voter can interact with the coercer during the voting protocol (see Table 1). Some researchers aim only to detect coercion rather than to mitigate it (e.g., Caveat Coercitor [GRBR13]).

In an unpublished manuscript, Smyth [Smy19] surveys four definitions of coercion resistance and finds that "coercion resistance has not been adequately formalized." According to Smyth, three of the definitions are too weak, and the general definition by Küsters [KTV10] is complex and too strong. His observations are controversial but demonstrate that settling definitions is still an elusive goal. Similarly, there remains some debate on the definition of receipt freeness [DKR06].

Previous work often makes strong assumptions: the voter knows an honest *Election Authority (EA)* official [CCM08]; the voter needs a special device to evade coercion [AFT10, ARR⁺10, CCM08, JCJ05]; the voter needs to perform mental arithmetic to evade coercion [WB09]; the voter needs to generate a fake password to evade coercion [CH11, ECH12]; the voter must complete registration before being coerced [JCJ05]; the election will not close before the voter can cast a ballot after coercion [LQAT20, SHD10, VG06]; a trusted server stores the voters' private keys [LQAT20]; and the probability of successful coercion is lowered by flooding voters with decoy ballots [Cha12]. VoteXX makes none of these assumptions.

Giustolisi et al. [GGS23] propose a coercion-resistance strategy based on flexible vote updating, which can evade last-minute voter coercion. This approach, however, is limited to over the shoulder coercion and assumes a voting server trusted for coercion resistance. VoteXX does not require this assumption.

We do assume the voter can use an *untappable channel*, as all coercion-resistant systems must—if an adversary can always influence the voter, they are indistinguishable from the voter [HS00]. Some systems establish windows for this channel, such as during registration, or after coercion occurs. VoteXX is as flexible as it could be. The channel is used once or twice between the voter and each hedgehog (who can be any person in the world): first to induct the hedgehog (any time before the end of the election), and possibly second to signal the hedgehog (after coercion and before the end of the election).

VoteXX guarantees that the voter is able to nullify their coerced vote. Unlike some systems, in VoteXX, the voter cannot change their coerced ballot selection. VoteXX can be used as an overlay, providing an additional coercion-resistant mechanism to others already in place. Thus, VoteXX can support re-voting (as outlined in our protocol description): if a voter were unable to re-vote (due to coercion at the end of the election), nullification would be a failsafe. Similarly, VoteXX can be used together with decoy ballots.

Table 1 compares our solution to previous proposed mechanisms, as well as to Helios [Adi08], which we take as a baseline coercible system. We do so by scoring each mechanism with regard to six properties of the influencer and four other properties. We state each assumption and property positively, meaning it is better to receive ● than ○. Appendix B explains our assumptions and the basis for our scoring.

In any system, an adversary could always prevent a voter from voting. In this sense, VoteXX achieves an optimal solution. Furthermore, we conjecture, that for Column 4 and Column 6 in Table 1, no system that resists extreme coercion can also undo coercion undetectably.

4 System Overview

In VoteXX, each voter has a public-private key pair for "YES" votes, and another such pair for "NO" votes. Without revealing their private keys, each voter registers their public keys with the Election Authority. This EA comprises a set of independent and non-colluding (up to a threshold) entities called trustees. Each voter may share their keys with one or more hedgehogs. During nullification, the voter, or one or more of their hedgehog(s), can interact with the ACS to nullify a vote by proving knowledge of one of the voter's private keys via a ZKP. We describe a fully decentralizable implementation of VoteXX, including its public *bulletin board (BB)*, which could be implemented on a blockchain.

4.1 Adversarial Model

The adversary could be anyone—including a voter or an EA trustee, located close to or far away from their target. The adversary might be covert or overt. The adversary's goal might include any or all of the following: tamper with the tally, influence a voter's ballot choice through coercion, learn how a voter voted, or disrupt or discredit an election. The adversary can engage in coercion at any time, including before or during voter registration.

We assume a secure ACS that protects against traffic analysis. Examples include TOR with hidden services [TOR02], I2P [I2P03], xx network [xx.21], and Oxen [Oxe20]. We further assume that the adversary cannot defeat standard cryptographic functions and protocols, including encryption, digital signatures, cryptographic hashing, pseudorandom number generation, and ZKPs. We assume an untappable channel between the voter and their hedgehog(s), as explained in Section 3.

4.2 Ideal Functionality

A foundational component of our UC proof (Appendix C) is the *voting ideal functionality* $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$, which we now introduce and define in Fig. 1.

The ideal functionality. The voting ideal functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ has four phases: preparation, registration, voting, and tally. In the voting phase, $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ receives ballots from the voters and records them. In particular, $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ accepts a special type of request: "nullify." Upon receiving a nullify request, the former choice of the voter will not be counted in the final tally.

Extreme coercion. In our UC model, the adversary has the power of extreme coercion. When the adversary $\mathcal{A}$ sends an "extreme coercion" request to a voter, V_i, V_i will hand his state to $\mathcal{A}$ and follow $\mathcal{A}$'s instructions, but V_i can still communicate with his hedgehog(s) H_i secretly.

Connection with the properties. It is easy to see that our UC definition implies the basic properties of a secure voting scheme. First, $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ does not leak the ballot of a voter to anyone else, so it implies ballot privacy. Second, as mentioned above, the ideal deception is able to nullify the ballot and the coercer cannot know if the coercion was

Functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$

The functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ interacts with a set of voters $\mathcal{V} := \{\mathsf{V}_1, \dots, \mathsf{V}_n\}$, a set of hedgehogs $\mathcal{H} := \{\mathsf{H}_1, \dots, \mathsf{H}_n\}$, a set of trustees $\mathcal{T} := \{\mathsf{T}_1, \dots, \mathsf{T}_k\}$, the *Election Authority (EA)*, and the adversary $\mathcal{S}$. Internally it keeps variables $\mathsf{status}, \mathsf{ballots}, \tau$, and $\mathcal{J}$. Let $\mathcal{P}_{\mathsf{cor}}$ be the set of corrupted parties.
Initially, set $\mathsf{status} := 0, \mathsf{ballots} := \tau := \mathcal{J} := \emptyset$.

Preparation:
- Upon receiving (START, sid) from the trustee $\mathsf{T}_j \in \mathcal{T}$, set $\mathcal{J} := \mathcal{J} \cup \{\mathsf{T}_j\}$, and send a notification (START, sid, T_j) to the adversary $\mathcal{S}$. (If $\mathsf{status} \neq 0$, then ignore the request.)
- Upon receiving (BEGIN, sid) from the EA, if $|\mathcal{J}| < k$ ignore the request. Otherwise, send a notification (BEGIN, sid) to the adversary $\mathcal{S}$, and set $\mathsf{status} := 1$. (If $\mathsf{status} \neq 0$, then ignore the request.)

Registration:
- Upon receiving (REGISTER, sid) from the voter V_i, send (REGISTER, sid, V_i) to the adversary $\mathcal{S}$. (If $\mathsf{status} \neq 1$, then ignore the request.)
- Upon receiving (ENDREG, sid) from EA, send (ENDREG, sid) to the adversary $\mathcal{S}$ and set $\mathsf{status} := 2$. (If $\mathsf{status} \neq 1$, then ignore the request.)

Voting:
- Upon receiving (VOTE, sid, x) from a voter $\mathsf{V}_i \in \mathcal{V}$, set $\mathsf{ballots}[i] := x$ (x=YES/NO), and send (VOTENOTIFY, sid, V_i) to the adversary $\mathcal{S}$. (If $\mathsf{status} \neq 2$, then ignore the request.)
- Upon receiving (ENDVOTE, sid) from EA, compute $\delta \leftarrow \mathsf{TallyAlg}(\mathsf{ballots})$ (Cf Fig. 6). Send (PRETALLY, sid, δ) to the adversary $\mathcal{S}$. Set $\mathsf{status} := 3$. (If $\mathsf{status} \neq 2$, then ignore the request.)
- Upon receiving (NULLIFY, sid) from a voter $\mathsf{V}_i \in \mathcal{V}$ or V_i's hedgehog H_i, set $\mathsf{ballots}[i] := \mathsf{nullify}$. Send (NULLIFYNOTIFY, sid) to the adversary $\mathcal{S}$. (If $\mathsf{status} \neq 3$, then ignore the request.)

Tally:
- Upon receiving (TALLY, sid) from EA, compute $\tau \leftarrow \mathsf{TallyAlg}(\mathsf{ballots})$ (Cf Fig. 6). Send (TALLY, sid, τ) to the adversary $\mathcal{S}$. (If $\mathsf{status} \neq 3$, then ignore the request.)
- Upon receiving (RESULT, sid) from any party P, if $\tau := \emptyset$, then ignore the request, otherwise return (RESULT, sid, τ) to the requester.

Figure 1: Functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$.

successful, so our definition implies coercion resistance. Third, $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ ensures that the tally procedure is performed correctly, so it implies verifiability.

4.3   Problem Specification

Our main requirement is a coercion-resistant remote voting system that achieves a level of security at least as strong as that for a precinct-based in-person voter-verifiable paper secret-ballot system. The system must maximize the ability to prevent or remediate serious failures by eliminating undetectable attacks, preventing scalable "wholesale" attacks, and making "retail" attacks as difficult as possible. The key requirements, specific to our context, are *coercion resistance*, *malware resistance*, and *availability*.

Coercion resistance. An adversary cannot be convinced that the voter's ballot is "counted as coerced," that is, counted the way the coercer instructed the voter to vote. This property is related to *ballot secrecy* but we assume that the adversary can watch the voter vote or vote for them. The adversary, however, cannot be sure how that vote is counted, so they have no incentive to threaten or pay the voter to vote a certain way. While rarely a significant issue in polling place elections, this problem is much more important in uncontrolled environments such as absentee voting or Internet voting.

Malware resistance. Any modification of the hardware or software that changes the result must be detectable. This property is similar to *software independence* but with the caveat that a version of the software exists without the undetected change before or after the election. In other words, the adversary does not, for all time, control everything read or written to all devices used by the voter for voting.

Availability. The system must not have single points of failure. It should resist denial of service attacks, and no single entity should be able to prevent completion of the election.

4.4   System Architecture

As shown in Fig. 2, we describe VoteXX in terms of the following entities and elements. There are *n voters* $v_1, v_2, \ldots, v_n$ who interact with a publicly readable *BB*, which is a distributed ledger such as a blockchain. The writing interactions take place via an ACS. The ACS disassociates the device, physical location, and other associated metadata by all clients posting to the BB, protecting the metadata of voters and hedgehogs, as well as sensitive election authority equipment. Read operations can take place through the ACS or via a direct interaction with the BB. Each voter may have one or more trusted *hedgehog(s)*. Each hedgehog interacts with the BB via the ACS. The EA comprises a set of independent and non-colluding (up to a threshold) entities called *trustees*. The trustees of the EA are authoritative over registration, voting, and tallying. The EA can read and write to the BB directly or via the ACS. The EA is a multiparty computation; it cannot compute any tally (preliminary or otherwise) without cooperation of the specified threshold of its members. The system includes a set of *auditors* who can read from the BB and verify the correctness of operations performed by the EA.
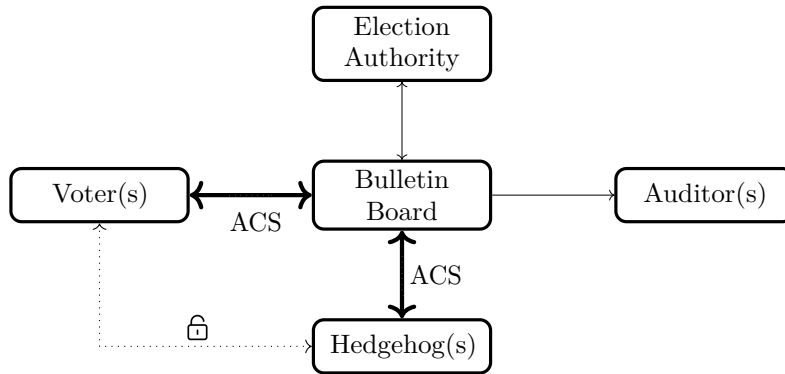
Figure 2: VoteXX Architecture. Arrows represent information flows (i.e., read/write) between system entities. Thick lines represent communications that take place over the ACS. The EA can write to the BB directly or via the ACS. Voters submit (encrypted) ballots over the ACS, read from the BB, and share their secret key over an untappable channel (dotted) to their hedgehog(s). The hedgehogs submit a ZKP to nullify the corresponding vote(s) over the ACS. Auditors read from the BB.

5 Protocols

Protocol Boxes 1–3 explain the four main stages of the VoteXX protocols: registration, voting, and tallying (including nullification). Section 6 explains nullification in more detail.

The VoteXX protocol assumes a number of cryptographic primitives that are common in the voting literature. All operations are performed in the same elliptic curve group, where the decisional *Diffie-Hellman (DDH)* problem (and by extension, the discrete logarithm problem) is hard. Digital signatures are performed with the Schnorr signature scheme. Encryption is performed with ElGamal [CGS97], which can be augmented with *distributed key generation (DKG)* and threshold decryption (for m out of n key holders [Ped91]).

We use standard $\Sigma-$Protocols to prove knowledge of discrete logarithms (Schnorr [Sch91]), knowledge of representations (Okamoto [Oka92]), and knowledge of Diffie-Hellman tuples (Chaum-Pedersen [CP92]), which also corresponds to ElGamal re-randomizations and decryptions. We also use techniques to allow the trustees to compute jointly, verifiably (*i.e.,* produce $\Sigma-$Protocol proofs), and privately on ElGamal ciphertexts the following: (i) a random shuffle of ciphertexts (Verificatum [Dou10]), and (ii) the evaluation of an *exclusive-or (XOR)* operation based on its logic lookup table (mix and match [JJ00]).

Protocol 1 describes registration. Registration can be re-opened by re-running set-up. Each voter needs to carry out registration only once, and the resulting keys can be reused in subsequent elections.

Protocol 1 is one way of performing registration, but any method that results in a posting of the voter's public key (in encrypted format) is fine. A simple way is for the voter to have the private key (full entropy, not based on a passphrase) on a hardware device and provide the public key. One problem that we tackle in Protocol 1 is providing assurance that the voter actually knows their private key—and it is not, for example, supplied by a coercer. This assurance is one of two properties of so-called *inalienable authentication*. The other property is that the adversary cannot impersonate the voter. Other authors do not provide concrete constructions for inalienable authentication; some simply tacitly assume it in their proofs of coercion resistance. In VoteXX, we provide a concrete instance of the first half of an unalienable authentication protocol, and we present a voting protocol that does not need the second half. That is, we care only that the voter knows their secret key—if the adversary knows it too, we can still achieve coercion resistance. There are

Registration is an in-person ceremony between the voter, using a *voting client* device, and an officer for the EA. At completion, the voter registers two public keys $\langle \mathsf{pk_{yes}}, \mathsf{pk_{no}} \rangle$, which are not learned by the EA officer and will be used to vote YES and NO, respectively. The keys are for a digital signature. They are based on a passphrase that can be regenerated from any voting client. The EA additionally does not learn the passphrase but has high assurance through the protocol that the human voter knows the passphrase.

Registration Set-up.
Registration uses a trapdoor commitment scheme. The commitment aspect allows the voter to present her passphrase in a hidden form to the EA and answer queries about specific characters within it. The trapdoor is revealed after registration closes and allows each voter to convert the format of their commitments into the format of a public key.

1. The generator g_0 is a parameter of the election.

2. The EA computes a generator g_1 as follows: each trustee $T, T', T'', \ldots$ privately chooses one random value a_1, reveals $g_0^{a_1}$, and proves knowledge of a_1 with a Schnorr $\Sigma-$Protocol. Then $g_1 = g_0^{(a_1 + a_1' + a_1'' + \ldots)}$.

3. This process is repeated, with new random a_i values, to complete a set of N generators: $\mathsf{base} \leftarrow \langle g_0, g_1, g_2, \ldots, g_{N-1} \rangle$. The same base is used for all voters in a registration period.

4. Call the set of all a values (split across the trustees): $\mathsf{trapdoor}$.

Registration.

1. Each voter generates two N-character passphrases (for YES and NO). Steps 2–4 describe the process for the first passphrase and are repeated for the second.

2. The voting client parses the passphrase as a sequence of Base64 characters $\langle c_0, c_1, c_2, \ldots, c_N \rangle$ and computes its deterministic commitment using base:
$\mathsf{passCommit} \leftarrow \left\langle g_0^{c_0} \cdot g_1^{c_1} \cdot g_2^{c_2} \cdot \ldots \cdot g_{N-1}^{c_{N-1}} \right\rangle$.

3. The voting client sends $[\![\mathsf{passCommit}]\!]$ to the EA, which is an encryption of $\mathsf{passCommit}$ under the EA's threshold encryption scheme.

4. The EA officer issues a challenge like: "Reveal Character 4." The voter responds "F." The EA client computes $\mathsf{disclosedChar} \leftarrow ([\![\mathsf{passCommit}]\!]/g_4^{\mathsf{F}})$. The voting client proves knowledge of a representation of $\mathsf{disclosedChar}$ using a $\Sigma-$Protocol. This step is repeated to build confidence that the voter knows the passphrase, but bounded in repetitions to protect the passphrase.

5. The EA client posts $\left\langle \mathsf{VoterID}, [\![\mathsf{passCommit_{yes}}]\!], [\![\mathsf{passCommit_{no}}]\!] \right\rangle$ to the BB.

Registration Finalization.

1. After the registration period, the EA takes the list of $\left\langle \mathsf{VoterID}, [\![\mathsf{passCommit_{yes}}]\!], [\![\mathsf{passCommit_{no}}]\!] \right\rangle$ entries, removes the $\mathsf{VoterID}$ component, and verifiably shuffles, threshold-decrypts, and posts $\left\langle \mathsf{passCommit_{yes}}, \mathsf{passCommit_{no}} \right\rangle$ for each (now anonymous) voter.

2. Each trustee $T, T', T'', \ldots$ reveals their values producing $\mathsf{trapdoor}$.

3. Each voter uses $\mathsf{trapdoor}$ to reformat their two $\mathsf{passCommit}$ values into key pairs $\langle \mathsf{sk}, \mathsf{pk} \rangle$ such that $\mathsf{pk} = \mathsf{passCommit} = g_0^{\mathsf{sk}}$ as follows. Consider generator g_i and let $\alpha_i = a_i + a_i' + \ldots$. With this notation, $\mathsf{sk} = c_0 + \alpha_1 \cdot c_1 + \alpha_2 \cdot c_2 \ldots$.

4. Given that $\left\langle \mathsf{passCommit_{yes}}, \mathsf{passCommit_{no}} \right\rangle = \left\langle \mathsf{pk_{yes}}, \mathsf{pk_{no}} \right\rangle$, the EA holds an anonymized list, which we call the Roster, of $\left\langle \mathsf{pk_{yes}}, \mathsf{pk_{no}} \right\rangle$ keys for each registered voter.

Protocol 1: Registration Protocol.

Voting.
Each voter completes voting online. At completion, each voter will have submitted their ballot using a passphrase from registration.

1. The value nonce is a parameter of the election.

2. To mark a ballot for YES, the voter uses their YES passphrase to generate sk_{yes} and uses this key to sign n_0: $\sigma_{yes} \leftarrow \mathsf{Sign}(\mathsf{nonce})$. Corresponding values are used to vote NO.

3. The voter uses the EA's threshold encryption scheme to compute $\mathsf{ballot} \leftarrow \left\langle [\![pk_{yes}]\!], [\![\sigma_{yes}]\!], \pi_{ppk} \right\rangle$, where each group element of σ is individually encrypted and π_{ppk} is a proof of plaintext knowledge using the Chaum-Pedersen $\Sigma-$Protocol.

4. The voter submits ballot over the ACS to the BB. The EA marks it as invalid if it is an exact duplicate or if the proofs are invalid.

Protocol 2: Voting Protocol.

significant advantages in authentication strength with in-person registration, but other choices can be made.

Voting performs a straightforward signature using a registered key (see Protocol 2). At the end of registration, voter keys are unlinked from their identity. Until the election closes, votes are encrypted to preserve the secrecy of the tally, and ballots are submitted through the *ACS* to unlink them from the voter communication metadata.

6 Nullification

We explain the nullification protocol in detail. First, we present an overview. Second, we give the construction of the nullification ZKP and propose a novel succinct ZKP with $O(\log n)$ proof size, where n is the number of total ballots.

6.1 Overview

The tallying process (Protocol 3) includes our novel nullification technique. Consider a list of public keys that voted YES and assume the hedgehog wants to nullify one of them. It cannot point out which key it wants to nullify or the coercer would know the voter is working with (or is personally acting as) a hedgehog to intervene. So the hedgehog must hide its flag ($[\![1]\!]$) in a set of false flags ($[\![0]\!]$) for each YES key in the tally. We could allow the hedgehog to choose a fixed-sized subset of β keys at random to serve as an anonymity set, which improves performance but sacrifices full anonymity (cf. [CH11]). For simplicity, the protocol boxes do not explain that, for nullification, we use exponential ElGamal [CGS97] instead of standard ElGamal used in registration and voting (under the same election master key).

If a hedgehog flags a key with ($[\![1]\!]$), it must know the associated private key; otherwise, any hedgehog could nullify any vote. However, if it submits a false flag ($[\![0]\!]$), it does not need to know the associated key. Anyone can serve as a dummy hedgehog by submitting a full set of false flags. To enforce these constraints, the hedgehog must construct a ZK proof to prove that: [for each flag, (it is an encryption of 0) or (it is an encryption of 1 and I know sk_{no} corresponding to this pk_{no})]. We will describe the Non-Interactive Zero-Knowledge (NIZK) proof for the above statement in Section 6.2.

Once a hedgehog computes and submits a set of flags (along with the NIZK proof Π), Protocol 3 simplifies the description by having the EA wait to perform Steps 1–2 after the nullification period. In practice, it should not wait—the process is quadratic work (number of hedgehogs times number of voters) and subject to "board-flooding"

Provisional Tally.
After the voting period ends, the EA produces a verifiable provisional tally.

1. The EA takes the list of $\langle [\![pk]\!], [\![\sigma]\!] \rangle$, verifiably shuffles them, then threshold-decrypts them: $\langle pk, \sigma \rangle$.

2. For each ballot, the ballot is marked invalid if σ does not verify under its corresponding pk.

3. For each valid signature, pk is matched to its entry on the Roster. The EA determines if it is a YES or NO key, and counts the vote only if it is the only ballot cast that corresponds to that roster entry.

Nullification.
The goal of nullification is to allow voters to modify their cast ballots, particularly in the case of coercion. Unlike other protocols, voters can enlist the help of others parties, called hedgehogs. The nullification period runs after the provisional tallying. If the provisional tally contains pk_{no}, it can be nullified using sk_{yes} (the "opposite" key). In other words, casting a YES and nullifying a NO vote use the *same* key, as these two actions are aligned in their intention.

1. At any convenient time, before or after voting, the voter covertly communicates with a hedgehog to develop a coercion-resistant strategy. For example, assume the following strategy: the voter wants to vote YES and reveals sk_{yes} to the hedgehog, along with $\langle pk_{yes}, pk_{no} \rangle$. They request the hedgehog engage in nullification if pk_{no} is in the provisional tally.

2. Using the Roster and set of valid signatures from the provisional tally, the EA reformats the election data into two lists. The first list establishes, in arbitrary order, the set of pk_{no} keys from voters who cast valid votes for YES (call it yesVotes). The second list contains pk_{yes} from voters who voted NO.

3. For example, assume YES received six votes in the provisional tally. yesVotes consists of six pk_{no} keys. If the hedgehog wants to nullify the fourth key, it prepares a list of encrypted "flags" marking the ballot it wants to nullify: $\langle [\![0]\!], [\![0]\!], [\![0]\!], [\![1]\!], [\![0]\!], [\![0]\!] \rangle$.

4. The first encrypted flag corresponds to the first pk_{no} in yesVotes. The hedgehog adds a proof to this list using the nullification ZK protocol. Concisely, the proof statement is: [for each flag, (it is an encryption of 0) or (it is an encryption of 1 and I know sk_{no} corresponding to this pk_{no})].

Final Tally.
After the nullification period ends, the EA produces a verifiable final tally.

1. The EA takes all the encrypted flags for the first pk_{no} key in yesVotes and computes its logical OR under encryption using the mix and match secure function evaluation (SFE) protocol [JJ00]. It repeats this process for the remaining pk_{no} keys.

2. The EA takes the list of encrypted OR-ed flags, sums them under encryption, and verifiably threshold-decrypts the result. The EA subtracts this value from the number of YES votes in the provisional tally to produce the final tally for YES votes.

3. The EA repeats Steps 1–2 for each pk_{yes} key in noVotes.

Protocol 3: Tallying Protocol (including nullification).

attacks [KHF11]. The EA must process the nullifications as they arrive; that is, use "concurrent authorization" [ECH12]. Doing so is possible. When a new set of flags arrives, the EA checks each proof and computes the OR between the submitted flag. The EA also computes the accumulation of previous flags—each of these two steps is parallelizable for each flag. Thus, when nullification closes, the only remaining task is to threshold decrypt the accumulation of flags, which process is linear in the number of votes.

6.2 The Nullification ZKP

We provide a formal description of the nullification ZK proof. It is well known that $\Sigma-$Protocols can be *stacked through conjunction and disjunction (CDS)* [CDS94, FS90]. We first present the CDS-composition ZKP and then propose a novel succinct ZKP.

The CDS-composition ZKP. The CDS-composition ZKP takes a voter's public key *pk* and makes a disjunctive proof that either Case 1 OR Case 2 is true: In Case 1, the hedgehog proves (flag $= [\![0]\!]$). For exponential ElGamal, assume $\langle c_1, c_2 \rangle = \mathsf{Enc}(m) = \langle g^r, g^m y^r \rangle$ for generator g, public key y, and message m. A proof it encrypts $\hat{m}$ is equivalent to proving $\langle g, c_1, y, c_2\hat{m}^{-1} \rangle$ is a DDH tuple, which can be done with the Chaum-Pedersen $\Sigma-$Protocol. Call this subproof A. In $\Sigma-$Protocol format, its transcript is $\langle a_A, e_A, z_A \rangle$.

In Case 2, the hedgehog proves a conjunctive statement: (flag $= [\![1]\!]$) and it knows *sk*, which corresponds to *pk* for the associated voter's public key. Call the subproof that (flag $= [\![1]\!]$) B. It is implemented the same as in subproof A, with transcript $\langle a_B, e_B, z_B \rangle$. Call the proof of knowledge of *sk* subproof C, which can be implemented with a $\Sigma-$Protocol due to Schnorr: $\langle a_C, e_C, z_C \rangle$. To summarize, the hedgehog proves: $\Pi := [\text{A OR (B AND C)}]$ for each flag.

Further, the resulting proof can be made non-interactive (typically in the random oracle model with the Fiat-Shamir heuristic [FS86], in its strong form [BPW12], but other heuristics exist [HL10]). Specifically, the prover generates a single challenge $\hat{e}$ for Π. To handle the conjunction within Case 2, $e_B = e_C$; for the disjunction across the cases, $\hat{e} = e_A + e_B$. In Case 1, the prover computes $\langle a_A, e_A, z_A \rangle$ and simulates $\langle a_B, e_B, z_B \rangle$ and $\langle a_C, e_B, z_C \rangle$. In Case 2, the prover simulates $\langle a_A, e_A, z_A \rangle$ and computes $\langle a_B, e_B, z_B \rangle$ and $\langle a_C, e_B, z_C \rangle$.

A new succinct ZKP. We propose a novel succinct nullification ZKP with $O(\log N)$ proof size, where N is the number of total ballots. Assuming that in the nullification phase, each nullification request nullifies only one ballot. To nullify a ballot, the hedgehog will form a list of encrypted "flags," where there is one encryption of 1 and the other flags are encryptions of 0. The hedgehog needs to prove in ZK that (i) there is one encrypted flag containing $[\![1]\!]$ and the others are $[\![0]\!]$, and (ii) I know the corresponding sk.

Formally, let h denote the ElGamal public key, and let $[\![x; r]\!]$ denote exponential ElGamal encryption with explicit randomness, i.e., $[\![x; r]\!] := (g^r, g^x h^r)$. Let ck denote the Pedersen commitment key, and let Com denote Pedersen commitment, i.e., $\mathsf{Com}(x; r) := g^x \mathsf{ck}^r$. Denote the public keys in yesVotes as $\mathsf{pk}_0, \ldots, \mathsf{pk}_{N-1}$. Denote the encrypted flags as $E_0, \ldots, E_{N-1}$. Let $n := \lceil \log N \rceil$. We will give a ZK protocol for the relation

$$\mathcal{R} = \{((\mathsf{pk}_0, \ldots, \mathsf{pk}_{N-1}, E_0, \ldots, E_{N-1}), r_0, \ldots, r_{N-1}, \ell, \mathsf{sk}) \mid$$
$$\ell \in \{0, \ldots, N-1\} \ \wedge \ \mathsf{pk}_\ell = g^{\mathsf{sk}} \ \wedge \ E_\ell = [\![1; r_\ell]\!] \ \wedge E_i = [\![0; r_i]\!], i \neq \ell\}. \tag{1}$$

Following the idea of [GK15, ZOB19], the prover first commits bit-wise to the binary representation of ℓ. The key observation is that there exists a data-oblivious algorithm that takes as input the binary representation of ℓ and generates a unit vector where the ℓth element is 1.

Concretely, the protocol can be split into two parts: The first part proves that there is one encrypted flag containing $[\![1]\!]$ and the others are $[\![0]\!]$, which is actually a unit vector proof [ZOB19]. The second part proves that the prover knows the corresponding sk, which can be proven by modifying the one-out-of-many proof [GK15]. The modification works

as follows. The prover first computes the commitment of sk, denoted as c. Then, the verifier can compute $c_i := \mathsf{pk}_i/c$. Now $\{c_i\}$ is a vector satisfying $c_\ell = \mathsf{Com}(0)$ so that the one-out-of-many proof [GK15] can be applied. The prover needs to additionally prove that he knows the opening of c.

We specify the polynomial $p_i(x)$ used in the protocol. Following [GK15], we write $i = i_1 \ldots i_n$ and $\ell = \ell_1 \ldots \ell_n$ in binary, and we let δ_{ij} be Kronecker's delta, i.e., $\delta_{\ell\ell} = 1$ and $\delta_{i\ell} = 0$ for $i \neq \ell$. We let $f_j = \ell_j x + a_j$, let $f_{j,1} = f_j = \ell_j x + a_j = \delta_{1\ell_j} x + a_j$ and $f_{j,0} = x - f_j = (1 - \ell_j)x - a_j = \delta_{0\ell_j} x - a_j$. Then, $p_i(x) = \prod_{j=1}^n f_{j,i_j}$ has the form:

$$p_i(x) = \prod_{j=1}^n (\delta_{i_j \ell_j} x) + \sum_{k=0}^{n-1} p_{i,k} x^k = \delta_{i\ell} x^n + \sum_{k=0}^{n-1} p_{i,k} x^k. \tag{2}$$

Fig. 3 shows the ZK protocol for relation $\mathcal{R}$. By the Fiat-Shamir heuristic [FS86], it can be transformed into a NIZK proof.

Theorem 1. *Assume that the DDH problem is hard. The protocol in Fig. 3 for relation $\mathcal{R}$ is a 4-move public coin ZK protocol with completeness, soundness, and special honest verifier ZK.*

Proof. For completeness, it is easy to see that $c_{\ell_j}^x c_{a_j} = \mathsf{Com}(f_j; z_{a_j})$ and $c_{\ell_j}^{x-f_j} c_{b_j} = \mathsf{Com}(0; z_{b_j})$ hold for $j \in \{1, \ldots, n\}$ and $g^{v_1} h^{v_2} = m \cdot c^x$ holds. Then, observe that $\prod_{j=1}^n f_{j,i_j}$ is a polynomial in the challenge x of the form $p_i(x) = \delta_{i\ell} x^n + \sum_{k=0}^{n-1} p_{i,k} x^k$. By the additive homomorphism of Pedersen commitment, $\prod_{i=0}^{N-1} c_i^{\prod_{j=1}^n f_{j,i_j}} \cdot \prod_{k=0}^{n-1} c_{d_k}^{-x^k} = \mathsf{Com}(0; z_d)$ always holds since c_ℓ is a commitment to 0. Similarly, denote $E_i = [\![e_i; r_i]\!]$, we have
$\prod_{i=0}^{N-1} \left((E_i)^{x^n} \cdot [\![- \prod_{j=1}^n f_{j,i_j}; 0]\!] \right)^{y^i} \cdot \prod_{k=0}^{n-1} (D_k)^{x^k} = [\![\sum_{i=0}^{N-1} (e_i \cdot x^n - p_i(x) + \sum_{k=0}^{n-1} p_{i,k} x^k) \cdot y^i; R]\!] = [\![0; R]\!]$. Thus, the protocol is perfectly complete.

To prove that the protocol is *sound*, suppose the adversary creates $n+1$ accepting responses $f_1^{(0)}, \ldots, v_2^{(0)}, \ldots, f_1^{(n)}, \ldots, v_2^{(n)}$ to $n+1$ different challenges $x^{(0)}, \ldots, x^{(n)}$ on the same initial message $c, \ldots, m$.

We first show that $\ell_j \in \{0, 1\}$ for $j \in [1, n]$. Pick two responses $f_j^{(0)}, z_{a_j}^{(0)}, z_{b_j}^{(0)}$ and $f_j^{(1)}, z_{a_j}^{(1)}, z_{b_j}^{(1)}$ to challenges $x^{(0)}, x^{(1)}$ on the commitments c_{a_j}, c_{b_j}. By combining the verification equations we obtain $c_{\ell_j}^{x^{(0)} - x^{(1)}} = \mathsf{Com}(f_j^{(0)} - f_j^{(1)}; z_{a_j}^{(0)} - z_{a_j}^{(1)})$ and $c_{\ell_j}^{x^{(0)} - f_j^{(0)} - x^{(1)} + f_j^{(1)}} = \mathsf{Com}(0; z_{b_j}^{(0)} - z_{b_j}^{(1)})$. Defining $\ell_j = \frac{f_j^{(0)} - f_j^{(1)}}{x^{(0)} - x^{(1)}}$ and $\gamma_j = \frac{z_{a_j}^{(0)} - z_{a_j}^{(1)}}{x^{(0)} - x^{(1)}}$ we extract an opening of $c_{\ell_j} = \mathsf{Com}(\ell_j; \gamma_j)$.
Furthermore, since $c_{\ell_j}^{x^{(0)} - f_j^{(0)} - x^{(1)} + f_j^{(1)}} = c_{\ell_j}^{(1 - \ell_j)(x^{(0)} - x^{(1)})} = \mathsf{Com}(\ell_j (1 - \ell_j)(x^{(0)} - x^{(1)}); \gamma_j (1 - \ell_j)(x^{(0)} - x^{(1)})) = \mathsf{Com}(0; z_{b_j}^{(0)} - z_{b_j}^{(1)})$, either $\ell_j(1 - \ell_j) = 0$ or the binding property of Pedersen commitment is broken. Thus, we have $\ell_j \in \{0, 1\}$ and extract $\ell = \ell_1 \ldots \ell_n$.

Then, the soundness is two-fold. In the first part, we prove $c = \mathsf{Com}(\mathsf{sk}) \wedge pk_\ell = g^{\mathsf{sk}}$ and extract sk. In the second part, we prove that $E_\ell = [\![1; r_\ell]\!] \wedge E_i = [\![0; r_i]\!], i \neq \ell$.

Let a_j be the number committed in c_{a_j}, from the verification equation $c_{\ell_j}^x c_{a_j} = \mathsf{Com}(f_j; z_{a_j})$ we conclude that $f_j^{(0)} = \ell_j x^{(0)} + a_j, \ldots, f_j^{(n)} = \ell_j x^{(n)} + a_j$ for all $j = 1, \ldots, n$ unless the adversary breaks the binding property of Pedersen commitment.

From the form of f_j's we have $f_{j,1} = \ell_j x + a_j$ and $f_{j,0} = (1 - \ell_j)x - a_j$. For $i \neq \ell$, it follows that $p_i(x) = \prod_{j=1}^n f_{j,i_j}$ is a polynomial of degree at most $n - 1$, and for $i = \ell$ it is a polynomial of the form $p_\ell(x) = x^n + \ldots$. Therefore we can rewrite $\prod_{i=0}^{N-1} c_i^{\prod_{j=1}^n f_{j,i_j}} \cdot \prod_{k=0}^{n-1} c_{d_k}^{-x^k} = \mathsf{Com}(0; z_d)$ as $c_\ell^{x^n} \cdot \prod_{k=0}^{n-1} c_{*_k}^{x^k} = \mathsf{Com}(0; z_d)$,

---

ZK protocol for relation $\mathcal{R}$

CRS: the ElGamal public key h, the Pedersen commitment key ck;
Statement: $\mathsf{pk}_0, \ldots, \mathsf{pk}_{N-1}, E_0, \ldots, E_{N-1}$;
Witness: $r_0, \ldots, r_{N-1}, \ell, \mathsf{sk}$ such that $\ell \in \{0, \ldots, N-1\} \wedge \mathsf{pk}_\ell = g^{\mathsf{sk}} \wedge E_\ell = [\![1; r_\ell]\!] \wedge E_i = [\![0; r_i]\!], i \neq \ell$.

Verifier:
- $V \to P$: Random $y \leftarrow \mathbb{Z}_q$.

Prover:
- Randomly pick $t \leftarrow \mathbb{Z}_q$, compute $c := \mathsf{Com}(\mathsf{sk}; t)$;
- For $i = 0, \ldots, N-1$, compute $c_i := \mathsf{pk}_i/c$.
- For $j = 1, \ldots, n$
    - Randomly pick $\tau_j, a_j, s_j, t_j, \rho_k \leftarrow \mathbb{Z}_q$;
    - Compute $c_{\ell_j} := \mathsf{Com}(\ell_j; \tau_j)$; $c_{a_j} := \mathsf{Com}(a_j; s_j)$; $c_{b_j} := \mathsf{Com}(\ell_j a_j; t_j)$;
    - Compute $c_{d_k} := \prod_{i=0}^{N-1} c_i^{p_{i,k}} \mathsf{Com}(0; \rho_k)$ (using $k = j - 1$ and $p_{i,k}$ from Eq. 2);
    - Pick random $R_k \leftarrow \mathbb{Z}_q$ and compute $D_k := [\![\sum_{i=0}^{N-1}(p_{i,k} \cdot y^i); R_k]\!]$ (using $k = j-1$ and $p_{i,k}$ from Eq. 2);
- Randomly pick $s', t' \leftarrow \mathbb{Z}_q$, compute $m := \mathsf{Com}(s', t')$;
- $P \to V$: $(c, c_{\ell_1}, c_{a_1}, c_{b_1}, c_{d_0}, D_0, \ldots, c_{\ell_n}, c_{a_n}, c_{b_n}, c_{d_{n-1}}, D_{n-1}, m)$.

Verifier:
- $V \to P$: Random $x \leftarrow \mathbb{Z}_q$.

Prover:
- For $j = 1, \ldots, n$
    - Compute $f_j := \ell_j x + a_j$; $z_{a_j} = \tau_j x + s_j$; $z_{b_j} = \tau_j(x - f_j) + t_j$;
- Compute $z_d = (-t)x^n - \sum_{k=0}^{n-1} \rho_k x^k$;
- Compute $R := \sum_{i=0}^{N-1}(r_i \cdot x^n \cdot y^i) + \sum_{k=0}^{n-1}(R_k \cdot x^k)$;
- Compute $v_1 := s' + x \cdot \mathsf{sk}$, $v_2 := t' + x \cdot t$;
- $P \to V$: $(f_1, z_{a_1}, z_{b_1}, \ldots, f_n, z_{a_n}, z_{b_n}, z_d, R, v_1, v_2)$.

Verifier:
- For $i = 0, \ldots, N-1$, compute $c_i = \mathsf{pk}_i/c$.
- For all $j \in \{1, \ldots, n\}$, check $c_{\ell_j}^x c_{a_j} = \mathsf{Com}(f_j; z_{a_j})$ and $c_{\ell_j}^{x-f_j} c_{b_j} = \mathsf{Com}(0; z_{b_j})$;
- Check $\prod_{i=0}^{N-1} c_i^{\prod_{j=1}^{n} f_{j,i_j}} \cdot \prod_{k=0}^{n-1} c_{d_k}^{-x^k} = \mathsf{Com}(0; z_d)$, using $f_{j,1} = f_j$ and $f_{j,0} = x - f_j$;
- Check $\prod_{i=0}^{N-1} \left((E_i)^{x^n} \cdot [\![-\prod_{j=1}^{n} f_{j,i_j}; 0]\!]\right)^{y^i} \cdot \prod_{k=0}^{n-1}(D_k)^{x^k} = [\![0; R]\!]$, using $f_{j,1} = f_j$ and $f_{j,0} = x - f_j$.
- Check $g^{v_1} \mathsf{ck}^{v_2} = m \cdot c^x$;
- Output 1 iff all the checks pass.

Figure 3: ZK protocol for relation $\mathcal{R}$.

for some fixed $c_{*_0}, \ldots, c_{*_{n-1}}$, which can be computed from commitments in $\{c_i\}_{i \in [0,N-1]}$ and the initial message.

Observe that the vectors $(1, x^{(\beta)}, \ldots, (x^{(\beta)})^n)$ can be viewed as rows in a Vandermonde matrix because $x^{(0)}, \ldots, x^{(n)}$ are all different. The matrix is invertible and we can therefore find a linear combination $(\alpha_0, \ldots, \alpha_n)$ of the rows that give us the vector $(0, \ldots, 0, 1)$. Combining the $n+1$ accepting verification equations, it follows that

$$c_\ell = \prod_{\beta=0}^{n} (c_\ell^{(x^{(\beta)})^n} \cdot \prod_{k=0}^{n-1} c_{*_k}^{(x^{(\beta)})^k})^{\alpha_\beta} = \mathsf{Com}(0; \sum_{\beta=0}^{n} \alpha_\beta z_d^{(\beta)}). \tag{3}$$

This equation gives us an extracted opening of c_ℓ to 0. Since $c_i = \mathsf{pk}_i/c$, and denoting $\mathsf{pk}_\ell = g^{\mathsf{sk}}$, we have $c = g^{\mathsf{sk}}\mathsf{ck}^t$, where $t = -\sum_{\beta=0}^{n} \alpha_\beta z_d^{(\beta)}$. Then, by $g^{v_1}\mathsf{ck}^{v_2} = m \cdot c^x$ we have that

$$c = g^{(v_1^{(0)} - v_1^{(1)})(x^{(0)} - x^{(1)})^{-1}} \mathsf{ck}^{(v_2^{(0)} - v_2^{(1)})(x^{(0)} - x^{(1)})^{-1}}. \tag{4}$$

This equation extracts $\mathsf{sk} = (v_1^{(0)} - v_1^{(1)})(x^{(0)} - x^{(1)})^{-1}$. Next, we start to prove that $E_\ell = [\![1; r_\ell]\!] \wedge E_i = [\![0; r_i]\!], i \neq \ell$. Denote $E_i = [\![e_i; r_i]\!]$. Since x is randomly chosen after the D_k's are committed, by the verification equation $\prod_{i=0}^{N-1} \left((E_i)^{x^n} \cdot [\![-\prod_{j=1}^{n} f_{j,i_j}; 0]\!]\right)^{y^i} \cdot \prod_{k=0}^{n-1} (D_k)^{x^k} = [\![0; R]\!]$, we have that $\prod_{i=0}^{N-1} \left((E_i)^{x^n} \cdot [\![-\prod_{j=1}^{n} f_{j,i_j}; 0]\!]\right)^{y^i} \cdot \prod_{k=0}^{n-1} (D_k)^{x^k}$ encrypts a zero polynimial w.r.t x with overwhelming probability (by the Schwartz-Zippel Lemma). Therefore, by denoting $i_{j,1} = \ell_j, i_{j,0} = 1 - \ell_j$, we obtain $Q(y) = \sum_{i=0}^{N-1}(e_i - \prod_{j=1}^{n} i_{j,i_j}) \cdot y^i = 0$. Since y is randomly chosen after the E_i's are encrypted, $Q(y)$ is a zero polynomial w.r.t y with overwhelming probability (by the Schwartz-Zippel Lemma). Hence, we have $e_i = \prod_{j=1}^{n} i_{j,i_j}$ for $i \in [0, N-1]$.

To prove that the protocol is *special honest verifier SK*, we build a simulator that is given $x, y \leftarrow \mathbb{Z}_q$. It first randomly picks $f_1, \ldots, v_2 \leftarrow \mathbb{Z}_q$. It then picks $c \leftarrow \mathbb{G}$ at random and $c_{\ell_1}, \ldots, c_{\ell_n}, c_{d_1}, \ldots, c_{d_{n-1}} \leftarrow \mathsf{Com}(0)$ as random commitments to 0. Next, it picks $U_i, R_i \leftarrow \mathbb{Z}_q$ at random and computes $D_i := [\![U_i; R_i]\!]$ for $i \in [1, n-1]$. After the random selection, it computes $c_i := \mathsf{pk}_i/c$; $c_{a_j} := c_{\ell_j}^{-x}\mathsf{Com}(f_j; z_{a_j})$, $c_{b_j} := c_{\ell_j}^{f_j - x}\mathsf{Com}(0; z_{b_j})$, $m := g^{v_1}h^{v_2}c^{-x}$, and

$$c_{d_0} := \prod_{i=0}^{N-1} c_i^{\prod_{j=1}^{n} f_{j,i_j}} \cdot \prod_{k=1}^{n-1} c_{d_k}^{-x^k} \cdot \mathsf{Com}(0; -z_d) \tag{5}$$

and

$$D_0 := \frac{[\![0; R]\!]}{\prod_{i=0}^{N-1} \left((E_i)^{x^n}[\![-\prod_{j=1}^{n} f_{j,i_j}; 0]\!]\right)^{y^i} \cdot \prod_{k=1}^{n-1} (D_k)^{x^k}}. \tag{6}$$

The simulator outputs the transcript $(y, c, \ldots, m, x, f_1, \ldots, v_2)$.

We argue that the adversary cannot distinguish the simulation from a real argument. First, in both real proofs and simulated proofs, $f_1, \ldots, v_2$ are uniformly random in $\mathbb{Z}_q$; c is uniformly random in $\mathbb{G}$. Furthermore, by the verification equations, $c_{a_1}, c_{b_1}, \ldots, c_{a_n}, c_{b_n}, m$, c_{d_0}, D_0 are determined by $f_1, \ldots, v_2$ and $c, c_{\ell_1}, \ldots, c_{\ell_n}, c_{d_1}, \ldots, c_{d_{n-1}}, D_1, \ldots, D_{n-1}$ both in real and in simulated proofs. The adversary's advantage must come from being able to distinguish $c, c_{\ell_1}, \ldots, c_{\ell_n}, c_{d_1}, \ldots, c_{d_{n-1}}, D_1, \ldots, D_{n-1}$ in real and simulated proofs. To do so, the adversary must either break the binding property of Pedersen commitment or break the IND-CPA property of ElGamal encryption by a standard hybrid argument. $\quad\square$

7 Implementation

We built proof-of-concept implementations for all components of VoteXX, included as an artifact in our submission. We wrote the EA and auditor in Java with the BouncyCastle [Leg00] library. We wrote the nullification and proofs in C++ with the cryptopp [Com95] library. Both implementations use the secp256k1 group. We can send inputs and outputs through ACS clients written in Golang using the native cryptographic libraries to interact with the BB, which currently is a simple file store utility.

We benchmarked specific operations on a PC using a AMD Ryzen 5 5600X 6-Core Processor with 2 16 GiB DIMMs at 2133 MT/s. The most expensive voting operation, tallying, took 15.34 seco nds for a simulated 2^{20} (1 million) voters. Assuming that a very high 2^{17} (25%) of voters nullify, it would take 9.21 minutes to verify the proofs. Each proof was 5.93 KB in size. See Section 8.3 for more general performance analysis.

8 Discussion

We now discuss our major design decisions, nullification options (cancel or flip), not revealing nullifications, performance analysis, extensions, and open problems.

8.1 Major Design Decisions

We made three major design decisions: (1) Nullification achieves the theoretically optimal *coercion resistance*, and using hedgehogs depends on a more realistic assumption than that assumed in previous work. (2) Our decentralized architecture provides *availability* and *malware resistance*. (3) In-person registration involving passphrases enhances voter authentication and supports key functionality for *malware resistance*.

Nullification and hedgehogs. Nullification allows the voter to share a passphrase anytime after they conceive of it. In-person registration ensures the voter knows their passphrases, providing ample opportunity even for captive voters (*e.g.,* a spouse or child) to signal a hedgehog. Because each passphrase can nullify a ballot only in one direction (the NO key can only vote NO or nullify YES; the YES key can only vote YES or nullify NO), voter intent matters and a signal to coordinate with a hedgehog can be optional. For example, a candidate who is a hedgehog might always nullify a ballot cast against them if possible.

Decentralized architecture. Routing all audit data through the ACS creates a special challenge to the adversary not present in traditional election systems: Any attack on the infrastructure must disable a much larger system, where there is an independent financial incentive for it to remain online. The BB, decentralized through the ACS, is not vulnerable to denial-of-service. Flooding the BB with data [KHF11] is limited as adversaries must pay for ACS bandwidth. Because all BB data are public and we use known E2E-voting constructions, the system meets the requirements for voter verifiable ballots, contestability, and auditability [PSNR20].

In-Person registration. Our registration design roots trust into passphrases known to the voter and written on physical paper associated with a specific person. This design provides a critical feature for the system's *malware resistance*: passphrases make it possible to detect and prove software misbehavior as all data posted to the BB can be regenerated with the passphrases on any device.

VoteXX greatly complicates undetectable wholesale attacks: the adversary must deploy malicious software across all devices controlled by checking with a passphrase. The deployment must go undetected forever, or at least until the election completes. If the attack is detected after the election, the adversary risks loss of confidence from a provably improper election outcome. It would be intractable for an adversary to remain undetected

for a useful period of time. Our design decisions allow VoteXX to prevent undetectable wholesale attacks at scale and provide detection and mitigation against retail attacks.

8.2 Cancel or Flip

Our design supports a variety of options for implementing the semantics of nullification, including what we call "cancel" or "flip." We recommend cancel, which is the default. Consider a vote that might have been nullified by one or more entities. We will describe the case for when there are two ballot choices (See Section 8.4 for the general case of k ballot choices). Assume that this vote selects from one of two ballot choices numbered $0, 1$. With *cancel*, the vote is cancelled if and only if at least one entity nullified it (and this idea can be generalized to at least t entities for some threshold t). With *flip*, the vote becomes $x + y \mod 2$, where x is the ballot choice of the vote, and y is the number of times the vote was nullified. Intuitively, cancel gives the voter the ability to cancel the vote, whereas flip gives the voter the ability to randomize the vote.

Each of these options can be implemented using different algebraic operations during Step 1 of the third phase (Final Tally) of Protocol 3: `AND` for cancel (realized with a homomorphic addition of the encrypted flags for each ballot followed by a plaintext equality test with $[\![0]\!]$), and `ADDITION` modulo 2 for flip (realized with mix and match. Step 2 replaces the final summation with a verifiable shuffle and threshold decryption of the flag set for each key). We view flip not as re-voting, but as "randomizing" the vote, which is a form of nullification.

As we point out below, under stronger assumptions, there are some use cases in which flip can be used to re-vote. Also, nullification can be used as an overlay in conjunction with re-voting strategies. A useful application of flip arises for a common form of low-intensity coercion. Suppose during remote voting at home, a coercer tells their spouse to vote for Alice and watches them comply, but the coercer does not collect the spouse's keys. The spouse can later flip their vote to Bob without the coercer knowing.

8.3 Performance Analysis

We analyze the running time of VoteXX for elections with T trustees, V voters, and H hedgehogs. If a passphrase is ℓ characters long with α possible characters, registration setup takes $\Theta(\ell \alpha T)$ work (comprised of modular exponentiations and $\Sigma-$Protocols). The proof size and verification time for the auditor is also $\Theta(\ell \alpha T)$. Example parameters might be $\alpha = 64$ characters of length $\ell = 20$ and $T = 10$ trustees. The shuffle proof dominates registration, generally taking $\Theta(VT \log V)$. Each vote has a constant amount of signatures, encryptions, and $\Sigma-$Protocols for the voter. Proof size and verification time for the auditor is $\Theta(V)$. The provisional tally consists of another shuffle, $\Theta(VT \log V)$, and decryption (subsumed in the shuffle), with the proof size and verification time of the same order for the auditor. Nullification is an involved protocol. As mentioned in Section 5, to avoid a quadratic bottleneck during the final tally, it is essential to process hedgehog flags as they arrive. Each hedgehog performs $\Theta(V)$ work (encryptions and $\Sigma-$Protocols) that an auditor must fetch and validate (space and time of $\Theta(V)$). For each of the V flags from one hedgehog, the trustees can precompute a logic gate (two-input gates are effectively constant time). Applying the gate to the inputs is $\Theta(T)$ (plaintext equality tests and $\Sigma-$Protocols). In total, nullification is $\Theta(HVT)$ work for the EA and auditors, with same order proof size on the BB. The final tally is fast: $\Theta(VT)$ work (consisting of decryption and $\Sigma-$Protocols) for the trustees and auditors, with same order proof size.

8.4 Extensions

We briefly describe several possible extensions of VoteXX.

Multiple candidates. VoteXX can be easily extended to support an election with multiple candidates. For example, for a k-candidate race, the voter can register k key pairs and then vote using the desired key. Without any major changes, nullification still operates as before. For example, to perform a flip, the system can use an addition modulo k to determine what flip to apply to the initially cast vote. Since the nullification protocol scales linearly in the number of voters and hedgehogs, introducing multiple candidates does not affect the overall performance of the nullification process.

Voting in person or by mail. To support the existing voting infrastructure, VoteXX can allow for a setting where the voting is accomplished by mail or in precincts using paper ballots. This capability can be achieved by incorporating a code-voting protocol, such as that used in Remotegrity [ZCIC⁺13].

Malware protection. Malicious software and can alter the operations performed by the voter. VoteXX allows for a two-phase voting process. In Phase 1, the user submits a vote or a vote commitment. In Phase 2, using a different device, the voter checks if the submission is correctly posted on the BB. Optionally, this extension can include an additional set of keys, where the user submits a payload signed with the additional keys and thereby "*locks in*" their submission.

Roster changes. If detected early in the election, it is possible to contest and remove a compromised passphrase. Providing proper documentation, the affected voter would rerun the in-person registration process.

Online registration. For lower-security elections, it is possible to replace the in-person registration with an online registration that follows appropriate identification mechanisms or uses an identity verification service [Wik22].

9   Conclusion

VoteXX provides a new, practical, and versatile solution to improper influence in elections against strong adversaries who learn the voter's voting keys. VoteXX works through the use of nullification supported by voter associates whom we call hedgehogs. In comparison with previous approaches, our solution makes fewer assumptions and protects against stronger adversaries. By separating our mechanism for mitigating improper influence from the mechanisms of ballot marking and collection, our technique works with a wide range of voting systems, including precinct voting with paper ballots, voting by mail, and Internet voting. For example, our mechanism works harmoniously with techniques for mitigating malware attacks, including allowing voters to check across multiple systems and devices. Also, our nullification mechanism can be used in addition to other mechanisms for mitigating improper influence.

Currently, election systems without voting booths are vulnerable to potential improper influence attacks. For example, a nation state, terrorist organization, billionaire, or anonymous hackers might offer significant amounts of money to vote for certain candidates. It could likely be impossible to know the extent to which such attacks succeeded. Such attacks would discredit the election, and re-running the election with the same technology would not resolve the issue. Our paper offers a solution to this threat and demonstrates that extreme coercion resistance is possible. Our work protects voting beyond the booth, and such voting is an essential enabler for the advance of democracy.

References

[Adi08]   Ben Adida. Helios: Web-Based open-audit voting. In *USENIX Security Symposium*, pages 335–348, 2008.

[AFT10]   Roberto Araujo, Sebastien Foulle, and Jacques Traoré. A practical and secure coercion-resistant scheme for Internet voting. *Toward Trustworthy Elections*, LNCS 6000, 2010.

[AOZZ15] Joël Alwen, Rafail Ostrovsky, Hong-Sheng Zhou, and Vassilis Zikas. Incoercible multi-party computation and universally composable receipt-free voting. In *Annual Cryptology Conference*, pages 763–780. Springer, 2015.

[ARR⁺10] Roberto Araujo, Narjes Ben Rajeb, Riadh Robbana, Jacques Traoré, and Souheib Yousfi. Towards practical and secure coercion-resistant electronic elections. In *CANS*, 2010.

[BN23]   Josh Benaloh and Michael Naehrig. Electionguard design specification version 2.0.0. Microsoft Research, 2023.

[BPW12]   David Bernhard, Olivier Pereira, and Bogdan Warinschi. How not to prove yourself: Pitfalls of the Fiat-Shamir heuristic and applications to Helios. In *ASIACRYPT*, 2012.

[Can01]   Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 136–145. IEEE, 2001.

[CCM08]   Michael R. Clarkson, Stephen Chong, and Andrew C. Myers. Civitas: Toward a secure voting system. In *IEEE Symposium on Security and Privacy*, pages 354–368, 2008.

[CDS94]   Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In *CRYPTO*, 1994.

[CGS97]   Ronald Cramer, Rosario Gennaro, and Berry Schoenmakers. A secure and optimally efficient multi-authority election scheme. In *EUROCRYPT*, 1997.

[CH11]   Jeremy Clark and Urs Hengartner. Selections: Internet voting with over-the-shoulder coercion-resistance. In *Financial Cryptography*, 2011.

[Cha12]   David Chaum. Random-Sample Voting. Online, 2012.

[CM22]   David Chaum and Thomas Moser. ecash 2.0 inalienably private and quantum-resistant to counterfeiting. Technical report, 2022.

[Com95]   Community Maintained, Originally written by Wei Dai. Crypto++ Library 8.7, 1995. Accessed: 2023-03-08. URL: https://www.cryptopp.com/.

[CP92]   David Chaum and Torben Pryds Pedersen. Wallet databases with observers. In *CRYPTO*, 1992.

[DKR06]   Stéphanie Delaune, Steve Kremer, and Mark Ryan. Coercion-Resistance and receipt-freeness in electronic voting. In *19th IEEE Computer Security Foundations Workshop (CSFW'06)*, pages 12+. IEEE, 2006.

[Dou10]   Douglas Wikström. Verificatum. https://www.verificatum.org/, 2010. Accessed on 05-01-2024.

[ECH12]   Aleksander Essex, Jeremy Clark, and Urs Hengartner. Cobra: Toward concurrent ballot authorization for Internet voting. In *EVT/WOTE*, 2012.

[Eya21]   Ittay Eyal. On cryptocurrency wallet design. In Vincent Gramoli, Hanna Hal-
          aburda, and Rafael Pass, editors, *3rd International Conference on Blockchain
          Economics, Security and Protocols, Tokenomics 2021, November 18-19, 2021,
          New York University, USA (Virtual Conference)*, volume 97 of *OASIcs*,
          pages 4:1–4:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
          URL: https://doi.org/10.4230/OASIcs.Tokenomics.2021.4, doi:10.4230/
          OASICS.TOKENOMICS.2021.4.

[FOO92]   Atsushi Fujioka, Tatsuaki Okamoto, and Kazuo Ohta. A practical secret voting
          scheme for large scale elections. In *International Workshop on the Theory and
          Application of Cryptographic Techniques*, pages 244–251. Springer, 1992.

[FS86]   Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to
          identification and signature problems. In *CRYPTO*, pages 186–194, 1986.

[FS90]   Amos Fiat and Adi Shamir. Witness indistinguishable and witness hiding
          protocols. In *ACM STOC*, 1990.

[GGS23]   Rosario Giustolisi, Maryam Sheikhi Garjan, and Carsten Schuermann. Thwarting
          last-minute voter coercion. Cryptology ePrint Archive, Paper 2023/1876, 2023.
          URL: https://eprint.iacr.org/2023/1876.

[GK15]   Jens Groth and Markulf Kohlweiss. One-out-of-many proofs: Or how to leak a
          secret and spend a coin. In *Advances in Cryptology-EUROCRYPT 2015: 34th
          Annual International Conference on the Theory and Applications of Cryptographic
          Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II*, pages 253–
          280. Springer, 2015.

[GRBR13] Gurchetan S Grewal, Mark D Ryan, Sergiu Bursuc, and Peter YA Ryan. Caveat
          coercitor: Coercion-evidence in electronic voting. In *2013 IEEE Symposium on
          Security and Privacy*, pages 367–381. IEEE, 2013.

[Hir01]   Martin Hirt. *Multi-Party Computation: Efficient Protocols, General Adversaries
          and Voting.* PhD thesis, ETH Zurich, 2001.

[HL10]   Carmit Hazay and Yehuda Lindell. *Efficient Secure Two-Party Protocols*.
          Springer, 2010.

[HMQA23] Thomas Haines, Johannes Müller, and Iñigo Querejeta-Azurmendi. Scalable
          coercion-resistant e-voting under weaker trust assumptions. In *ACM SAC*, 2023.

[HS00]   Martin Hirt and Kazue Sako. Efficient receipt-free voting based on homomorphic
          encryption. In *EUROCRYPT*, 2000.

[I2P03]   I2P. Invisible Internet project. https://www.geti2p.net, 2003. Accessed on
          01-05-2022.

[JCJ05]   Ari Juels, Dario Catalano, and Markus Jacobsson. Coercion-Resistant electronic
          elections. In *ACM WPES*, 2005.

[JJ00]   Markus Jakobsson and Ari Juels. Mix and match: Secure function evaluation
          via ciphertexts. In *ASIACRYPT*, 2000.

[KHF11]   Reto Koenig, Rolf Haenni, and Stephan Fischli. Preventing board flooding
          attacks in coercion-resistant electronic voting schemes. In *SEC*, 2011.

[KTV10]   Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Accountability: Definition
          and relationship to verifiability. In *ACM CCS*, 2010.

[Leg00]   Legion of the Bouncy Castle. Bouncy Castle crypto APIs, 2000. Accessed:
            2022-05-05. URL: https://www.bouncycastle.org/java.html.

[LQAT20]   Wouter Lueks, Iñigo Querejeta-Azurmendi, and Carmela Troncoso. Voteagain:
            A scalable coercion-resistant voting system. In *USENIX Security*, 2020.

[MN06]   Tal Moran and Moni Naor. Receipt-free universally-verifiable voting with ever-
            lasting privacy. In *CRYPTO*, 2006.

[Oka92]   T Okamoto. Provably secure and practical identification schemes and corre-
            sponding signature schemes. In *CRYPTO*, 1992.

[Oxe20]   Oxen. Privacy made simple. https://www.oxen.io, 2020. Accessed on 01-05-
            2022.

[Ped91]   Torben Pryds Pedersen. A threshold cryptosystem without a trusted party. In
            EUROCRYPT, 1991.

[PSNR20]   Sunoo Park, Michael Specter, Neha Narula, and Ronald L. Rivest. Going from
            bad to worse: From Internet voting to blockchain voting. Online, 2020.

[RED22]   REDACTED. Redacted (extended abstract) for anonymous submission, 2022.

[Sch91]   C P Schnorr. Efficient signature generation by smart cards. *Journal of Cryptog-
            raphy*, 4:161–174, 1991.

[SHD10]   Oliver Spycher, Rolf Haenni, and Eric Dubuis. Coercion-Resistant hybrid voting
            systems. In *EVOTE*, 2010.

[Smy19]   Ben Smyth. Surveying definitions of coercion resistance. Cryptology ePrint
            Archive, Report 2019/822, 2019.

[TOR02]   TOR. The TOR project. https://www.torproject.org, 2002. Accessed on
            01-05-2022.

[VG06]   M. Volkamer and R. Grimm. Multiple casts in online voting: Analyzing chances.
            In *EVOTE*, 2006.

[WB09]   R Wen and R Buckland. Masked ballot voting for receipt-free online elections.
            In *VOTE-ID*, 2009.

[Wik22]   Wikipedia. Identity verification service, 2022. Accessed: 2022-05-07. URL:
            https://en.wikipedia.org/wiki/Identity_verification_service.

[xx.21]   xx.network. A quantum leap in privacy. https://www.xx.network, 2021. Ac-
            cessed on 01-05-2022.

[ZCIC⁺13]   Filip Zagórski, Richard T. Carback III, David Chaum, Jeremy Clark, Alek-
            sander Essex, and Poorvi L. Vora. Remotegrity: Design and use of an end-to-end
            verifiable remote voting system. In *ACNS*, 2013.

[ZOB19]   Bingsheng Zhang, Roman Oliynykov, and Hamed Balogun. A treasury system
            for cryptocurrencies: Enabling better collaborative intelligence. In *The Network
            and Distributed System Security Symposium 2019*, 2019.

A   Acronyms and Abbreviations

ACS   anonymous communication system
BB   bulletin board
CDS   stacked through conjunction and disjunction
DDH   decisional Diffie-Hellman
DKG   distributed key generation
EA   Election Authority
E2E   end-to-end
NIZK   non-interactive zero knowledge
TA   Tally Authority
UC   universal composability
VA   Voting Authority
ZK   zero knowledge
ZKP   zero-knowledge proof

B   Explanation of Properties in Table 1

0. *System resists coercion when the influencer: acts before/during registration.* In
   a number of coercion resistance mechanisms, the voter is expected to register a
   user-chosen key, password, or be assigned a key. If registration can be done in the
   presence of the adversary or using inputs supplied by the adversary (and compliance
   can be checked by the adversary), without impacting coercion resistance, we award ●.
   Otherwise if the system assumes the adversary cannot be active during this process
   or places limitations on their actions in corrupting the registrants, we award ○.

1. *System resists coercion when the influencer: colludes with the EA.* A system should
   maintain coercion-resistance even when the coercer is able to corrupt a minority of
   the EA (●). Limitations on this assumption might result in ◐: for example, in fake
   credentials, it is assumed the coercer can corrupt any member of the EA but the
   voter must know which EA member has not been corrupted. Other systems fail to
   provide coercion resistance (○) under this assumption.

2. *System resists coercion when the influencer: colludes with hardware manufactures.*
   A system that does not rely on trusted hardware to provide coercion-resistance is
   awarded ●. By contrast, as system that makes hardware assumptions beyond typical
   computational equipment, such as a trusted execution environment [AOZZ15], is
   awarded ○ (along with systems that do not provide coercion resistance).

3. *System resists coercion when the influencer: acts at any time.* Assuming an influencer
   cannot act at all times (see Property 5), are there additional restrictions on when
   they can act? If not: ●. Systems receive ○ include ones that assume the influencer
   does not act before or during registration, and systems like re-voting that assume
   the coercer does not act at the very end of the voting period (blocking a re-vote).

4. *System resists coercion when the influencer: learns all information stored by the
   voter, including all keys required by the protocol.* Our main contribution is that
   VoteXX achieves coercion resistance even if all the voter's stored keys/secrets are
   leaked—extreme coercion resistance. This property is stronger than the literature,
   which generally assumes voters can establish and maintain secret keys or passwords
   (and lie convincingly about them as necessary) that will need to be recalled to cast a
   vote: ○. While non-verifiable re-voting does not require voter secrets or private keys,
   end-to-end verifiable systems do as a way to cryptographically link ballots and prevent
   multiple votes from the same voter. VoteAgain is designed as an exception to this

rules; in it, such keys exist but are maintained by a special election trustee so voters do not need to. However it must be completely trusted for coercion resistance (and in fact, must be trusted for ballot privacy and election integrity as well [HMQA23]). Use of a trusted third party also receives ○. Trusting a hardware enclave to maintain keys is awarded ◑. Systems that assume the influencer cannot impersonate the voter but do not provide a specific mechanism for online settings are also awarded ○.

5. *System resists coercion when the influencer: learns every action taken by the voter.* With reasonable assumptions on how voting works, this property is in fact shown to be impossible to achieve, as the voter can never act independently [HS00]. We include it to highlight this fact and as an open problem: perhaps some other trust model or assumptions on the voter would enable this property.

6. *Voter can undo coercion undetectably.* Coercion needs to be corrected when the voter's intent is different than the influencer's. Assuming the coercion resistance mechanism is allowed to work, if the voter is always able to vote their true intent, the mechanism is awarded ●. By contrast, systems are awarded ○ when voters cannot reliably vote their true intent. These systems, however, can still be considered coercion resistant if they do allow the voter to cancel the coercer's intent by spoiling, nullifying, or randomizing their ballot. In VoteXX, voters can vote their true intent if they can predict the influencer's actions and respond strategically. However, we cannot assume this ability will always be the case, and so, at best, voters can cancel or randomize their ballots. In VoteXX this choice depends on a system configuration discussed in Section 8.2.

7. *System is inexpensive.* A system that does not introduce new expenses beyond the EA running a server and voters having access to standard computational devices is awarded ●. A system that requires special equipment or hardware for either the EA or for the voters is awarded ○ (*e.g.,* special hardware for digital signatures [AOZZ15]).

8. *System has low cognitive burden.* If the human voters strategy for evading coercion is automated or does not require any cognitive effort, it is awarded ●. If the voter needs to remember passwords, it is awarded ◑. Any strenuous mental effort (*e.g.,* remembering an integer offset and performing mental arithmetic [WB09]) is awarded ○.

9. *System has security proof.* A universal composability proof is awarded ●, a game-based security definition and proof is awarded ◑, while informal security arguments and sketches are awarded ○.

C Security Analysis

We formally state and prove properties of VoteXX in the *Universal Composability (UC)* framework [Can01]. To begin, we define types of coercion and state the security of VoteXX. Next, we give a UC specification of VoteXX and state and prove a theorem that characterizes its security properties, using "cancel" nullification (see Section 8.2). We model our proofs in part from those of Alwen, Ostrovsky, Zhou, and Zikas [AOZZ15].

C.1 Preliminaries

In this section, we formally define the cryptographic primitives in VoteXX and their properties.

NIZK. A *non-interactive zero-knowledge proof (NIZK)* for relation $\mathcal{R}$ consists of four probabilistic polynomial 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. Completeness says that an honest prover can always 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 it does not leak any information 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. Soundness says that a prover cannot 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 probabilistic polynomial algorithms $(\mathsf{Keygen}, \mathsf{Enc}, \mathsf{Dec})$. We require the underlying encryption scheme to be *indistinguishable under chosen-plaintext attack (IND-CPA)*. Consider the following experiment:

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

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 $\mathsf{Adv}_{\mathsf{Enc}}^{\mathsf{IND\text{-}CPA}}(\mathcal{A}, \lambda) :=$ $|2 \cdot \Pr[\mathrm{EXPT}_{\mathcal{A},\mathsf{Enc}}^{\mathsf{IND\text{-}CPA}}(\lambda) = 1] - 1|$ is negligible in λ.

Signature. A *signature scheme* consists of three probabilistic polynomial algorithms ($\mathsf{Keygen}, \mathsf{Sign}, \mathsf{Verify}$). We require the underlying signature scheme to be *existentially unforgeable under chosen-message attack (EUF-CMA)*. Consider the following experiment:

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

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

Bulletin Board. A *bulletin board* is required for any voting scheme to record ballots and other related information. We model the bulletin board as a UC global functionality $\mathcal{G}_{\mathsf{BB}}$, which is depicted in Fig. 4. It has two interfaces: READ and WRITE, and records all the valid messages. $\mathcal{G}_{\mathsf{BB}}$ ensures that all the communication with the bulletin board is anonymous.

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

The ideal functionality $\mathcal{G}_{\mathsf{BB}}$ is globally available to all participants. It is parameterized with a predicate $\mathsf{Validate}$.
Upon initialization, set $\mathsf{Storage} := \emptyset$.
Upon receiving (READ, sid) from P:

- Let $\mathsf{val} := \mathsf{Storage}[\mathsf{sid}]$;

- Return (READ, $\mathsf{sid}, \mathsf{val}$) to the requestor.

Upon receiving (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 (RECEIPT, sid) to the requestor;

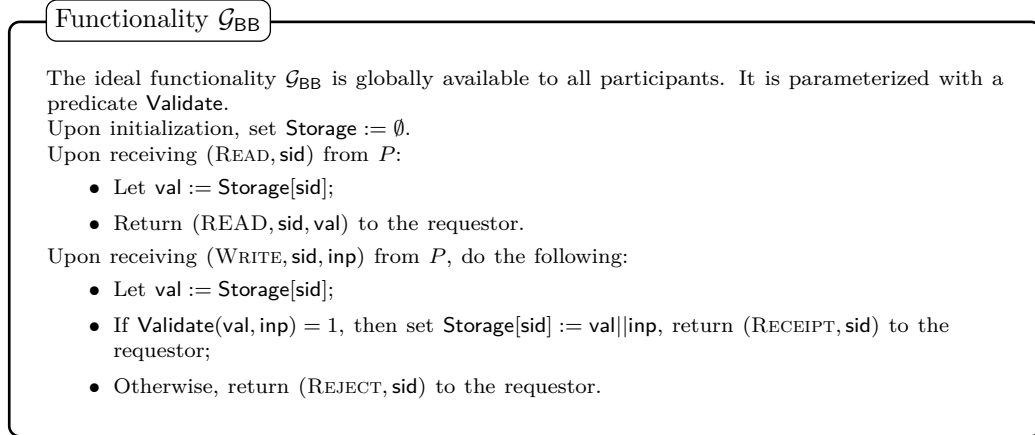- Otherwise, return (REJECT, sid) to the requestor.

Figure 4: Functionality $\mathcal{G}_{\mathsf{BB}}$

C.2   Types of Coercion

To articulate the capability of the coercer in VoteXX, we define a new type of coercion: *extreme coercion*, which differs from previous notions of *semi-honest coercion (receipt-freeness)* and *active coercion*.

Semi-Honest coercion. This type of coercion is most common in the literature; resistance against semi-honest coercion is called receipt-freeness. In this case, the coercer provides an input to the coerced party and expects evidence that such input was used, which evidence is called a "*receipt.*" For example, the receipt can be the entire view of the protocol execution.

Active coercion. Moran and Naor [MN06] proposed this stronger type of coercion. Instead of merely requiring a receipt, the coercer can query the current view interactively and send commands to the coerced party during the protocol execution.

Extreme coercion. We define this new stronger type of coercion, called "extreme coercion," which captures the real world more accurately. The coercer can obtain *all* the secret keys and passwords of the coerced party, and can perform operations in substitution for the coerced party. The coerced party, however, can secretly communicate with other people via some untappable channel.

Extreme coercion captures real world coercion more realistically because the coercer may ask the coerced party to hand over their device to extract the secret keys and monitor the coerced party's action. Because we consider it impossible to coerce a target throughout their entire life, they can recruit a hedgehog and agree on some secret action in advance.

C.3   Ballot privacy, coercion resistance and verifiability.

In this section, we give (informal) definitions of ballot privacy, coercion resistance, and verifiability. Then, we give intuition why VoteXX satisfies these properties. We formalize the secure definition in Section C.4 under the UC framework and argue that the UC definition implies these properties.

Ballot privacy [FOO92]. All votes must be secret.

The link between the voter and the corresponding public key in the roster is hidden by the verifiable shuffle in the registration phase. In addition, all the ballots are encrypted under the EA's threshold encryption scheme in the voting phase. Thus, VoteXX ensures ballot privacy assuming that the majority of EA trustees are honest.

Coercion resistance. No coercer can tell if the coerced party is trying to deceive.

The ballots and nullification requests are posted on the BB via an ACS to avoid identity leakage. In the nullification phase, the flags marking which ballots are to be nullified are encrypted and a ZKP establishes knowledge of the corresponding secret key. In addition, we assume that there is an untappable channel between the voter and his hedgehog(s) that cannot be blocked by the coercer. Therefore, the coercer cannot stop a coerced party from nullifying his vote and cannot know if the ballot is nullified.

Verifiability [FOO92]. No one can falsify the result of the voting.

We assume an honest BB and the messages posted on the BB cannot be deleted OR changed. In the provisional and final tallies, VoteXX uses ZKPs to ensure that the shuffle and decryption are performed correctly. The max-and-match SFE protocol [JJ00] in the final tally is publicly verifiable.

C.4   Security Definition.

We define the security of VoteXX in the UC framework [Can01]. A protocol is represented as a set of *interactive Turing machines (ITMs)*, where each ITM represents the program to be run by a participant. There are two additional entities: the *environment* $\mathcal{Z}$ and the *adversary* $\mathcal{A}$. The environment $\mathcal{Z}$ can communicate with $\mathcal{A}$ and provides inputs to the parties. We assume that each ITM is a *probabilistic polynomial-time (PPT)* machine.

Security is based on the indistinguishability between *real/hybrid world executions* and *ideal world executions*. Specifically, in the ideal world, all the participants are dummy parties and there is an ideal functionality $\mathcal{F}$ that serves as a trusted third party. We say that a *protocol π UC-realizes $\mathcal{F}$* if and only if, for any PPT adversary $\mathcal{A}$, there exists a PPT

simulator $\mathcal{S}$ such that no PPT environment $\mathcal{Z}$ can distinguish between the real/hybrid world and the ideal world.

Our protocol contains *pre-tally* and *nullification* phases. Therefore, by comparing the final tally with the pre-tally, the coercer can discern the amount by which nullification altered the results. The coercer, however, cannot attribute such difference to the voter. To capture this feature, we model the ideal world as follows.

Ideal deception. Our treatment of incoercibility in the ideal world is inspired by Alwen, Ostrovsky, Zhou, and Zikas [AOZZ15]. Alwen et al. define an ideal deception strategy DI as a mapping applied on the message given by the coercer to an intended choice, and they realize this ideal deception with the assumption of trusted hardware. In our system, since we make the minimal assumption of an untappable channel between the voter and their hedgehog(s), we cannot realize such a strong DI. Alternatively, we define a weaker DI: If DI chooses to obey, it forwards the coercer's input to the ideal functionality. If DI chooses to deceive, it forwards the coercer's input to the ideal functionality but sends a nullification request at the end of voting phase. Meanwhile, the ideal functionality accepts a special command—"nullify"—and leaks the provisional tally to the simulator $\mathcal{S}$, which captures the features of VoteXX. In this type of weaker ideal deception, the coercer may know that there are people deceiving him, but he cannot attribute such deception to the voter.

Definition 1. Let π be any protocol and let $\mathcal{F}$ be any ideal functionality. Let $\mathcal{A}$ be the adversary who has the power of corruption and coercion. We say that π *IUC realizes* $\mathcal{F}$ if, for every $i \in [n]$ and for every ideal deception strategy DI_i, there exists a real deception strategy DR_i such that, for every PPT adversary $\mathcal{A}$, there exist a simulator $\mathcal{S}$ such that, for any set $\mathsf{DI}_{\mathcal{J}} = \{\mathsf{DI}_i : i \in \mathcal{J}\}$ and any environment $\mathcal{Z}$,

$$\mathsf{EXEC}_{\mathcal{F},\mathsf{DI}_{\mathcal{J}},\mathcal{S},\mathcal{Z}} \approx \mathsf{EXEC}_{\pi,\mathsf{DR}_{\mathcal{J}},\mathcal{A},\mathcal{Z}}. \tag{7}$$

Theorem 2. *(Universal composition) Let π, ρ be any polynomial-time protocols, and let $\mathcal{F}$ be any ideal functionality. If π IUC realizes $\mathcal{F}$, then ρ^π IUC realizes $\rho^{\mathcal{F}}$.*

Following Alwen et al., it is easy to see that our framework remains universally composable with the same type of DI.

The ideal functionality. The voting ideal functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ has four phases: preparation, registration, voting, and tally. In the voting phase, $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ receives ballots from the voters and records them. In particular, $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ accepts a special type of request: "nullify." Upon receiving a nullify request, the former choice of the voter will not be counted in the final tally. Fig. 5 is a formal description of $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$.

Extreme coercion. In our UC model, the adversary has the power of extreme coercion. It is modeled as follows. When the adversary $\mathcal{A}$ sends a "extreme coercion" request to a voter, V_i, V_i will hand his state to $\mathcal{A}$ and follow $\mathcal{A}$'s instructions, but V_i can still communicate with his hedgehog(s) H_i secretly, i.e., the communication between V_i and H_i is not controlled by the adversary.

Connection with the properties. It is easy to see that our UC definition implies the basic properties of a secure voting scheme. First, $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ does not leak the ballot of a voter to anyone else, so it implies ballot privacy. Second, as mentioned above, the ideal deception is able to nullify the ballot and the coercer cannot know if the coercion was successful, so our definition implies coercion resistance. Third, $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ ensures that the tally procedure is performed correctly, so it implies verifiability.

C.5   UC Specification of VoteXX

Before we give a UC proof for VoteXX, we give a UC description of the VoteXX protocol. We assume that the protocol uses "cancel" nullification (Cf. 8.2). We will discuss the
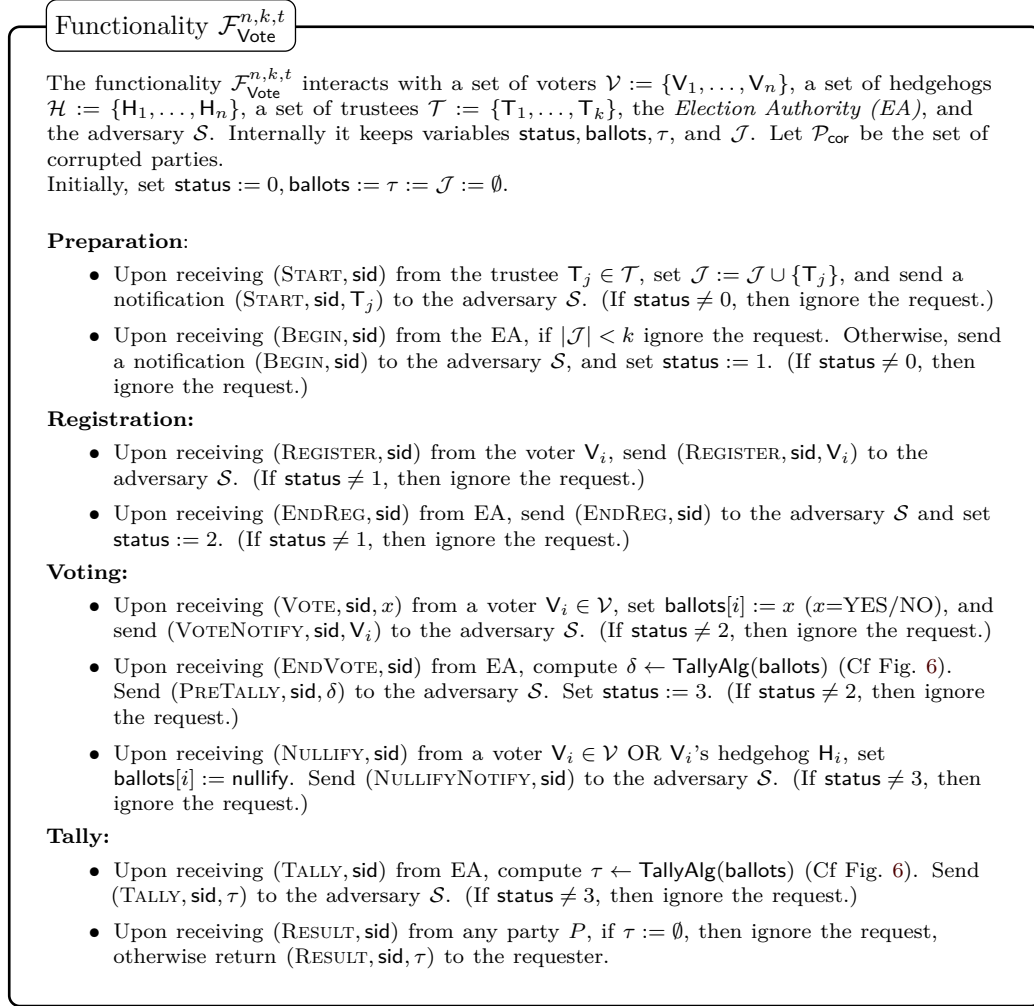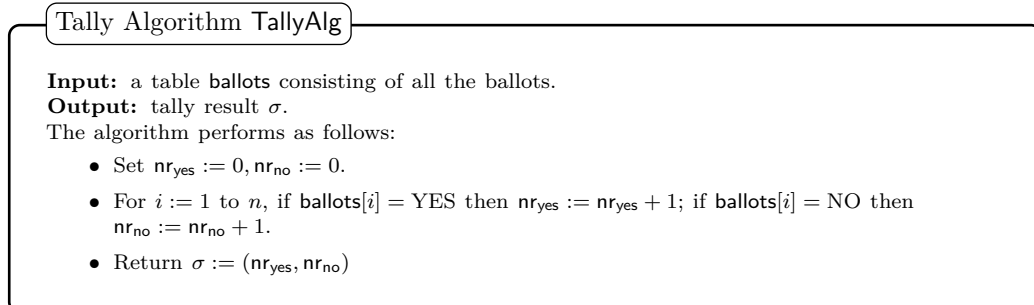

Functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$

The functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ interacts with a set of voters $\mathcal{V} := \{\mathsf{V}_1, \ldots, \mathsf{V}_n\}$, a set of hedgehogs $\mathcal{H} := \{\mathsf{H}_1, \ldots, \mathsf{H}_n\}$, a set of trustees $\mathcal{T} := \{\mathsf{T}_1, \ldots, \mathsf{T}_k\}$, the *Election Authority (EA)*, and the adversary $\mathcal{S}$. Internally it keeps variables $\mathsf{status}, \mathsf{ballots}, \tau$, and $\mathcal{J}$. Let $\mathcal{P}_{\mathsf{cor}}$ be the set of corrupted parties.
Initially, set $\mathsf{status} := 0, \mathsf{ballots} := \tau := \mathcal{J} := \emptyset$.

Preparation:
- Upon receiving (START, sid) from the trustee $\mathsf{T}_j \in \mathcal{T}$, set $\mathcal{J} := \mathcal{J} \cup \{\mathsf{T}_j\}$, and send a notification (START, sid, T_j) to the adversary $\mathcal{S}$. (If $\mathsf{status} \neq 0$, then ignore the request.)
- Upon receiving (BEGIN, sid) from the EA, if $|\mathcal{J}| < k$ ignore the request. Otherwise, send a notification (BEGIN, sid) to the adversary $\mathcal{S}$, and set $\mathsf{status} := 1$. (If $\mathsf{status} \neq 0$, then ignore the request.)

Registration:
- Upon receiving (REGISTER, sid) from the voter V_i, send (REGISTER, sid, V_i) to the adversary $\mathcal{S}$. (If $\mathsf{status} \neq 1$, then ignore the request.)
- Upon receiving (ENDREG, sid) from EA, send (ENDREG, sid) to the adversary $\mathcal{S}$ and set $\mathsf{status} := 2$. (If $\mathsf{status} \neq 1$, then ignore the request.)

Voting:
- Upon receiving (VOTE, sid, x) from a voter $\mathsf{V}_i \in \mathcal{V}$, set $\mathsf{ballots}[i] := x$ (x=YES/NO), and send (VOTENOTIFY, sid, V_i) to the adversary $\mathcal{S}$. (If $\mathsf{status} \neq 2$, then ignore the request.)
- Upon receiving (ENDVOTE, sid) from EA, compute $\delta \leftarrow \mathsf{TallyAlg}(\mathsf{ballots})$ (Cf Fig. 6). Send (PRETALLY, sid, δ) to the adversary $\mathcal{S}$. Set $\mathsf{status} := 3$. (If $\mathsf{status} \neq 2$, then ignore the request.)
- Upon receiving (NULLIFY, sid) from a voter $\mathsf{V}_i \in \mathcal{V}$ OR V_i's hedgehog H_i, set $\mathsf{ballots}[i] := \mathsf{nullify}$. Send (NULLIFYNOTIFY, sid) to the adversary $\mathcal{S}$. (If $\mathsf{status} \neq 3$, then ignore the request.)

Tally:
- Upon receiving (TALLY, sid) from EA, compute $\tau \leftarrow \mathsf{TallyAlg}(\mathsf{ballots})$ (Cf Fig. 6). Send (TALLY, sid, τ) to the adversary $\mathcal{S}$. (If $\mathsf{status} \neq 3$, then ignore the request.)
- Upon receiving (RESULT, sid) from any party P, if $\tau := \emptyset$, then ignore the request, otherwise return (RESULT, sid, τ) to the requester.

Figure 5: Functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$

Tally Algorithm TallyAlg

Input: a table $\mathsf{ballots}$ consisting of all the ballots.
Output: tally result σ.
The algorithm performs as follows:
- Set $\mathsf{nr}_{\mathsf{yes}} := 0, \mathsf{nr}_{\mathsf{no}} := 0$.
- For $i := 1$ to n, if $\mathsf{ballots}[i] = \mathsf{YES}$ then $\mathsf{nr}_{\mathsf{yes}} := \mathsf{nr}_{\mathsf{yes}} + 1$; if $\mathsf{ballots}[i] = \mathsf{NO}$ then $\mathsf{nr}_{\mathsf{no}} := \mathsf{nr}_{\mathsf{no}} + 1$.
- Return $\sigma := (\mathsf{nr}_{\mathsf{yes}}, \mathsf{nr}_{\mathsf{no}})$

Figure 6: Tally Algorithm TallyAlg

security of "flip" variant nullification in Section C.7.

VoteXX protocol $\Pi_{\text{VoteXX}}^{n,k,t}$

Denote the voters as $\mathcal{V} := \{V_1, \ldots, V_n\}$, the hedgehogs as $\mathcal{H} := \{H_1, \ldots, H_n\}$, the trustees as $\mathcal{T} := \{T_1, \ldots, T_k\}$, and the Election Authority as EA. We assume that EA cannot be corrupted.

Preparation phase:

Upon receiving (START, sid) from the environment $\mathcal{Z}$, the trustee T_i performs the following:

- Privately choose one random values $\{a_{i,j}\}_{j \in [N-1]}$, reveal $g_0^{a_{i,j}}$, and prove knowledge of $a_{i,j}$ with a Schnorr Σ−Protocol.

Upon receiving (BEGIN, sid) from the environment $\mathcal{Z}$, the EA performs the following:

- Compute $g_j := \prod_{i=1}^{k} g_0^{a_{i,j}}, j \in [N-1]$.

- Set base $\leftarrow \langle g_0, g_1, g_2, \ldots, g_{N-1} \rangle$ and send (WRITE, sid, base) to $\mathcal{G}_{\text{BB}}$.

Registration phase:

Upon receiving (REGISTER, sid) from the environment $\mathcal{Z}$, the voter V_i performs the following:

- Send (READ, sid) to $\mathcal{G}_{\text{BB}}$ and get base.

- Generate two N-character passphrases (for YES and NO).

- Parse the passphrase as a sequence of Base64 characters $\langle c_0, c_1, c_2, \ldots, c_N \rangle$ and compute its deterministic commitment using
  base: passCommit $\leftarrow \langle g_0^{c_0} \cdot g_1^{c_1} \cdot g_2^{c_2} \cdot \ldots \cdot g_{N-1}^{c_{N-1}} \rangle$.

- Send (PASSCOMMIT, sid, $\langle$passCommit$_{\text{yes}}$, passCommit$_{\text{no}}\rangle$) to the EA.

Upon receiving (PASSCOMMIT, sid, $\langle$passCommit$_{\text{yes}}$, passCommit$_{\text{no}}\rangle$) from the voter V_i, the EA performs the following:

- Send (WRITE, sid, $\langle$VoterID, $[\![$passCommit$_{\text{yes}}]\!]$, $[\![$passCommit$_{\text{no}}]\!]\rangle$) to $\mathcal{G}_{\text{BB}}$, where $[\![$passCommit$]\!]$ is an encryption of passCommit under the EA's threshold encryption scheme.

- Prove to the voter client the correctness of the encryptions using the Chaum-Pedersen Σ−Protocol.

Upon receiving (ENDREG, sid) from the environment $\mathcal{Z}$, the EA performs the following:

- Take the list of $\langle$VoterID, $[\![$passCommit$_{\text{yes}}]\!]$, $[\![$passCommit$_{\text{no}}]\!]\rangle$ entries, remove the VoterID component, and verifiably shuffle, threshold-decrypt, and send
  (WRITE, sid, $\langle$passCommit$_{\text{yes}}$, passCommit$_{\text{no}}\rangle$) to $\mathcal{G}_{\text{BB}}$.

- Send (REVEAL, sid) to each trustee T_i.

Upon receiving (REVEAL, sid) from EA, trustee T_i sends
(WRITE, sid, $\{a_{i,j}\}_{j \in [N-1]}$) to $\mathcal{G}_{\text{BB}}$.

Upon receiving (ENDREG, sid) from the environment $\mathcal{Z}$, the EA sends (STARTVOTE, sid) to each voter.

Voting phase:

Upon receiving (STARTVOTE, sid) from EA, voter V_i performs the following:

- Send (READ, sid) to $\mathcal{G}_{\mathsf{BB}}$ and get $\{a_{i,j}\}_{i\in[k],j\in[N-1]}$.

- Compute $\mathsf{sk} := c_0 + \alpha_1 \cdot c_1 + \alpha_2 \cdot c_2 \ldots$, where $\alpha_i := a_{1,i} + a_{2,i} + \ldots$.

Upon receiving (VOTE, sid, x) from the environment $\mathcal{Z}$, voter V_i performs the following:

- If $x = \mathrm{YES}$,

  - Generate $\sigma_{\mathsf{yes}} \leftarrow \mathsf{Sign(nonce)}$ and use EA's threshold encryption scheme to compute $\mathsf{ballot} \leftarrow \langle [\![\mathsf{pk}_{\mathsf{yes}}]\!], [\![\sigma_{\mathsf{yes}}]\!], \pi_{\mathsf{ppk}} \rangle$, where each group element of σ is individually encrypted and $\pi_{\mathsf{ppk}} \leftarrow \mathsf{NIZK_{ballot}.prove}([\![\mathsf{pk}_{\mathsf{yes}}]\!], [\![\sigma_{\mathsf{yes}}]\!])$ is a proof of plaintext knowledge using the Chaum-Pedersen $\Sigma-$Protocol.

  - Send (SECRETKEY, sid, $\mathsf{sk}_{\mathsf{yes}}, \langle \mathsf{pk}_{\mathsf{yes}}, \mathsf{pk}_{\mathsf{no}} \rangle$) to his hedgehog H_i.

- If $x = \mathrm{NO}$,

  - Generate $\sigma_{\mathsf{no}} \leftarrow \mathsf{Sign(nonce)}$ and use EA's threshold encryption scheme to compute $\mathsf{ballot} \leftarrow \langle [\![\mathsf{pk}_{\mathsf{no}}]\!], [\![\sigma_{\mathsf{no}}]\!], \pi_{\mathsf{ppk}} \rangle$, where each group element of σ is individually encrypted and $\pi_{\mathsf{ppk}} \leftarrow \mathsf{NIZK_{ballot}.prove}([\![\mathsf{pk}_{\mathsf{no}}]\!], [\![\sigma_{\mathsf{no}}]\!])$ is a proof of plaintext knowledge using the Chaum-Pedersen $\Sigma-$Protocol.

  - Send (SECRETKEY, sid, $\mathsf{sk}_{\mathsf{no}}, \langle \mathsf{pk}_{\mathsf{yes}}, \mathsf{pk}_{\mathsf{no}} \rangle$) to his hedgehog H_i.

- Send (WRITE, sid, ballot) to $\mathcal{G}_{\mathsf{BB}}$.

Upon receiving (ENDVOTE, sid) from the environment $\mathcal{Z}$, the EA performs the following:

- Send (READ, sid) to $\mathcal{G}_{\mathsf{BB}}$ and get the list of $\langle [\![\mathsf{pk}]\!], [\![\sigma]\!] \rangle$, then verifiably shuffle and threshold-decrypt them: $\langle \mathsf{pk}, \sigma \rangle$. Send (WRITE, sid, $\langle \mathsf{pk}, \sigma \rangle, \pi$) to $\mathcal{G}_{\mathsf{BB}}$, where $\pi \leftarrow \mathsf{NIZK_{Dec}.prove}([\![\mathsf{pk}]\!], [\![\sigma]\!], \mathsf{pk}, \sigma)$ is the decryption NIZK.

- For each ballot, the ballot is marked invalid if σ does not verify under its corresponding pk.

- For each valid signature, determine if it is a YES OR NO key, and count the vote only if it is the only ballot cast that corresponds to that roster entry.

- Use the Roster and set of valid signatures from the provisional tally to reformat the election data into two lists. The first list establishes, in arbitrary order, the set of $\mathsf{pk}_{\mathsf{no}}$ keys from voters who cast valid votes for YES (call it yesVotes). The second list contains $\mathsf{pk}_{\mathsf{yes}}$ from voters who voted NO (call it noVotes).

- Send (WRITE, sid, $\langle$ yesVotes, noVotes $\rangle$) to $\mathcal{G}_{\mathsf{BB}}$.

- Send (ENDVOTE, sid) to every voter.

Upon receiving (NULLIFY, sid, 1) from the voter V_i, the hedgehog H_i does the following:

- Send (READ, sid) to $\mathcal{G}_{\mathsf{BB}}$ and find the key to be nullified in yesVotes OR noVotes. Denote the index of the key as ℓ.

- Prepare a list of encrypted "flags" L marking the ballot it wants to nullify, i.e., the i'th element is $[\![1]\!]$ and the other elements are $[\![0]\!]$.

- Add a proof to this list using the nullification $\Sigma-$Protocol. Concisely, the proof statement is: [(this flag is an encryption of 0) OR (this flag is an encryption of 1 and I know $\mathsf{sk}_{\mathsf{no}}$ corresponding to this $\mathsf{pk}_{\mathsf{no}}$)]. Denote the nullification proof as $\pi \leftarrow \mathsf{NIZK_{nul}.prove}(\mathsf{yesVotes/noVotes}, L)$.

- Send (WRITE, sid, $\langle L, \pi \rangle$) to $\mathcal{G}_{\mathsf{BB}}$.

Upon receiving ($\textsc{Nullify}$, sid, 0) from the voter V_i, the hedgehog H_i does the following:

- Prepare a list of encrypted "flags" L where all the elements are $[\![0]\!]$.

- Add a proof to this list using the nullification $\Sigma-$Protocol. Concisely, the proof statement is: [(this flag is an encryption of 0) OR (this flag is an encryption of 1 and I know $\mathsf{sk}_{\mathsf{no}}$ corresponding to this $\mathsf{pk}_{\mathsf{no}}$)]. Denote the nullification proof as $\pi \leftarrow \mathsf{NIZK}_{\mathsf{nul}}.\mathsf{prove}(\mathsf{yesVotes}/\mathsf{noVotes}, L)$.

- Send ($\textsc{Write}$, sid, $\langle L, \pi \rangle$) to $\mathcal{G}_{\mathsf{BB}}$.

Tally phase:
Upon receiving ($\textsc{Tally}$, sid) from the environment $\mathcal{Z}$, EA performs the following:

- Send ($\textsc{Read}$, sid) to $\mathcal{G}_{\mathsf{BB}}$ and collect all encrypted flags.

- For each $\mathsf{pk}_{\mathsf{no}}$ in $\mathsf{yesVotes}$, take all encrypted flags and compute their OR under encryption using the max-and-match SFE protocol [JJ00].

- Take the list of encrypted ORed flags, sum them under encryption (denote it as c), and verifiably threshold-decrypt the result (denote it as x).

- Subtract this value from the number of YES votes in the provisional tally to produce the final tally for YES votes.

- Repeat the above three steps for each $\mathsf{pk}_{\mathsf{yes}}$ key in $\mathsf{noVotes}$.

- Denote the final tally result as τ. Send ($\textsc{Write}$, sid, $\langle \tau, \pi \rangle$) to $\mathcal{G}_{\mathsf{BB}}$, where $\pi \leftarrow \mathsf{NIZK}_{\mathsf{Dec}}.\mathsf{prove}(c, x)$ is the decryption proof.

Upon receiving ($\textsc{Result}$, sid) from the environment $\mathcal{Z}$, the party P returns ($\textsc{Result}$, sid, τ).

C.6   UC Proof for VoteXX

We have the following theorem.

Theorem 3. *Assume that* $\mathsf{NIZK}_{\mathsf{ballot}}$, $\mathsf{NIZK}_{\mathsf{nul}}$ *and* $\mathsf{NIZK}_{\mathsf{Dec}}$ *are complete, sound and zero-knowledge; the encryption scheme is IND-CPA secure; and the signature scheme is existentially unforgeable against chosen-message attack. The VoteXX protocol* $\Pi_{\mathsf{VoteXX}}^{n,k,t}$ *IUC realizes* $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ *against static corruption and mobile extreme coercion in the* $\mathcal{G}_{\mathsf{BB}}$*-hybrid model.*

Proof. To prove the theorem, we need to 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 $\mathsf{EXEC}_{\Pi_{\mathsf{VoteXX}}^{n,k,t}, \mathsf{DR}, \mathcal{A}, \mathcal{Z}}^{\mathcal{G}_{\mathsf{BB}}}$ from the (ii) the ideal execution $\mathsf{EXEC}_{\mathcal{F}_{\mathsf{Vote}}^{n,k,t}, \mathsf{DI}, \mathcal{S}, \mathcal{Z}}^{\mathcal{G}_{\mathsf{BB}}}$.

Real Deception Strategy. The real deception strategy DR_i internally runs DI_i, forwarding messages to and from the environment $\mathcal{Z}$. DR_i works as follows:
DR_i follows the coercer's instructions. Upon receiving ($\textsc{EndVote}$, sid) from the EA:

- If DI_i does not send a nullification request to $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$, then DR_i sends ($\textsc{Nullify}$, sid, 0) to the hedgehog H_i via the untappable channel.

- If DI_i sends ($\textsc{Nullify}$, sid) to $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$, then DR_i sends ($\textsc{Nullify}$, sid, 1) to the hedgehog H_i via the untappable channel.

Simulator. The simulator $\mathcal{S}$ internally runs $\mathcal{A}$, forwarding messages to and from the environment $\mathcal{Z}$. The simulator $\mathcal{S}$ works as follows:

In the preparation phase:

- Upon receiving (START, sid, T_i) from the functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$, $\mathcal{S}$ simulates the trustee T_i following the protocol $\Pi_{\mathsf{VoteXX}}^{n,k,t}$ as if T_i received (START, sid) from the environment $\mathcal{Z}$.

- Upon receiving (BEGIN, sid) from the functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$, $\mathcal{S}$ simulates the EA following the protocol $\Pi_{\mathsf{VoteXX}}^{n,k,t}$ as if EA received (BEGIN, sid) from the environment $\mathcal{Z}$.

In the registration phase:

- Upon receiving (REGISTER, sid, V_i) from the functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$, the simulator $\mathcal{S}$ simulates V_i following the protocol $\Pi_{\mathsf{VoteXX}}^{n,k,t}$ as if V_i received (REGISTER, sid) from the environment $\mathcal{Z}$.

- Upon receiving (ENDREG, sid) from the functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$, the simulator $\mathcal{S}$ simulates EA following the protocol $\Pi_{\mathsf{VoteXX}}^{n,k,t}$ as if the EA received (ENDREG, sid) from the environment $\mathcal{Z}$.

In the voting phase:

- Upon receiving (VOTENOTIFY, sid, V_i) from the functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$, send an encryption of 0 $\langle [\![0]\!], [\![0]\!], \pi_{\mathsf{ppk}} \rangle$ to $\mathcal{G}_{\mathsf{BB}}$.

- If a corrupted party V_i casts a valid ballot for an honest voter on $\mathcal{G}_{\mathsf{BB}}$, $\mathcal{S}$ will abort.

- When a corrupted party V_i casts a ballot on $\mathcal{G}_{\mathsf{BB}}$, decrypt the ballot to get the choice $x = \text{YES/NO}$ and send (VOTE, sid, x) to $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ on behalf of V_i in the ideal world.

- Upon receiving (PRETALLY, sid, δ) from the functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$, the simulator $\mathcal{S}$ simulates the pre-tally result by simulating the NIZK for decryption based on δ.

- Upon receiving (NULLIFYNOTIFY, sid) from the functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$, the simulator sends a *dummy nullification request* (a nullification request where the encrypted flags are all $[\![0]\!]$) to $\mathcal{G}_{\mathsf{BB}}$.

- If a corrupted party V_i sends a valid nullification request for an honest voter on $\mathcal{G}_{\mathsf{BB}}$, $\mathcal{S}$ will abort.

- When a corrupted party V_i sends a nullification request on $\mathcal{G}_{\mathsf{BB}}$, decrypt the nullification request. Otherwise, if it is not a dummy nullification request, send (NULLIFY, sid) to $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$ on behalf of V_i in the ideal world.

In the tally phase:

- Upon receiving (TALLY, sid, τ) from the functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$, $\mathcal{S}$ simulates EA doing the following:

  - For each pk in yesVotes and noVotes, take all the encrypted flags, compute its OR under encryption. Sum them under encryption for yesVotes and noVotes, respectively.

  - Simulate the decryption of the summed encrypted flags and the corresponding NIZK proof π based on τ.

– Send (WRITE, sid, τ, π) to $\mathcal{G}_{\mathsf{BB}}$.

Indistinguishability.

We prove indistinguishability through a series of hybrid worlds $\mathcal{H}_0, \ldots, \mathcal{H}_9$.

Hybrid $\mathcal{H}_0$: This object is the real world execution $\mathsf{EXEC}^{\mathcal{G}_{\mathsf{BB}}}_{\Pi^{n,k,t}_{\mathsf{VoteXX}},\mathsf{DR},\mathcal{A},\mathcal{Z}}$.

Hybrid $\mathcal{H}_1$: $\mathcal{H}_1$ is the same as $\mathcal{H}_0$ except the followings. During the pre-tally phase (upon receiving (ENDVOTE, sid) from the environment $\mathcal{Z}$) and tally phase, the EA's decryption proofs are generated by NIZK simulator.

Claim 1: If the decryption NIZK is zero-knowledge with adversary advantage $\mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{Dec}}}(\mathcal{A}, \lambda)$, then $\mathcal{H}_1$ and $\mathcal{H}_0$ are indistinguishable with distinguishing advantage at most $(2n+2) \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{Dec}}}(\mathcal{A}, \lambda)$.

Proof: Each ballot has two ciphertexts and there are n voters in total, so there are $2n$ ciphertexts to decrypt in pre-tally. In tally phase, there are 2 ciphertexts to decrypt. Therefore, the overall advantage is at most $(2n+2) \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{Dec}}}(\mathcal{A}, \lambda)$ by a standard hybrid argument.

Hybrid $\mathcal{H}_2$: $\mathcal{H}_2$ is the same as $\mathcal{H}_1$ except the followings. During the pre-tally phase (upon receiving (ENDVOTE, sid) from the environment $\mathcal{Z}$) and tally phase, the honest EA members' decryption shares are backward calculated from the pre-tally result and the tally result, respectively.

Claim 2: If the encryption scheme is backward calculatable, then $\mathcal{H}_2$ and $\mathcal{H}_1$ are perfectly indistinguishable.

Proof: The backward calculated decryption shares in $\mathcal{H}_2$ and the decryption shares in $\mathcal{H}_1$ have the same distribution .

Hybrid $\mathcal{H}_3$: $\mathcal{H}_3$ is the same as $\mathcal{H}_2$ except the followings. During the voting phase, the honest voters' ballot NIZK proofs are generated by the NIZK simulator.

Claim 3: If the ballot NIZK proof is zero-knowledge with adversary advantage $\mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{ballot}}}(\mathcal{A}, \lambda)$, then $\mathcal{H}_3$ and $\mathcal{H}_2$ are indistinguishable with distinguishing advantage at most $n \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{ballot}}}(\mathcal{A}, \lambda)$.

Proof: There are at most n honest voters, so the overall advantage is at most $n \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{ballot}}}(\mathcal{A}, \lambda)$ by a standard hybrid argument.

Hybrid $\mathcal{H}_4$: $\mathcal{H}_4$ is the same as $\mathcal{H}_3$ except the followings. During the voting phase, the honest voters' ballots are replaced with $\langle [\![0]\!], [\![0]\!] \rangle$.

Claim 4: If the encryption scheme is IND-CPA secure with advantage $\mathsf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathsf{Enc}}(\mathcal{A}, \lambda)$, then $\mathcal{H}_4$ and $\mathcal{H}_3$ are indistinguishable with distinguishing advantage at most $2n \cdot \mathsf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathsf{Enc}}(\mathcal{A}, \lambda)$.

Proof: There are at most $2n$ ciphertexts in total. Thus, the overall advantage is at most $2n \cdot \mathsf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathsf{Enc}}(\mathcal{A}, \lambda)$ by a standard hybrid argument.

Hybrid $\mathcal{H}_5$: $\mathcal{H}_5$ is the same as $\mathcal{H}_4$ except the followings. During the voting phase, honest parties' nullification NIZKs are generated by the NIZK simulator.

Claim 5: If the nullification NIZK proof is zero-knowledge with adversary advantage $\mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{nul}}}(\mathcal{A}, \lambda)$, then $\mathcal{H}_5$ and $\mathcal{H}_4$ are indistinguishable with distinguishing advantage at most $n \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{nul}}}(\mathcal{A}, \lambda)$.

Proof: There are at most n honest voters, so the overall advantage is at most $n \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{nul}}}(\mathcal{A}, \lambda)$ by a standard hybrid argument.

Hybrid $\mathcal{H}_6$: $\mathcal{H}_6$ is the same as $\mathcal{H}_5$ except the followings. During the voting phase, honest parties' nullification requests are replaced with $\langle [\![0]\!], \ldots, [\![0]\!] \rangle$.

Claim 6: If the encryption scheme is IND-CPA secure with advantage $\mathsf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathsf{Enc}}(\mathcal{A}, \lambda)$, then $\mathcal{H}_6$ and $\mathcal{H}_5$ are indistinguishable with distinguishing advantage at most $n \cdot \mathsf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathsf{Enc}}(\mathcal{A}, \lambda)$.

Proof: For each nullification request, there is an encryption of 1 replaced with encryption of 0, and there are at most n honest voters. Thus, the overall advantage is at most $n \cdot \mathsf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathsf{Enc}}(\mathcal{A}, \lambda)$ by a standard hybrid argument.

Hybrid $\mathcal{H}_7$: $\mathcal{H}_7$ is the same as $\mathcal{H}_6$ except that, if a corrupted voter generates a valid nullification request for an honest voter, the execution will abort.

Claim 7: If the nullification NIZK is sound with soundness error $\mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{nul}}}(\mathcal{A}, \lambda)$, then $\mathcal{H}_7$ and $\mathcal{H}_6$ are indistinguishable with distinguishing advantage at most $n \cdot \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{nul}}}(\mathcal{A}, \lambda)$.

Proof: There are at most n honest voters, so the probability of aborting is no more than $n \cdot \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{nul}}}(\mathcal{A}, \lambda)$ by a standard hybrid argument.

Hybrid $\mathcal{H}_8$: $\mathcal{H}_8$ is the same as $\mathcal{H}_7$ except that, if a corrupted voter generates a valid ballot for an honest voter, the execution will abort.

Claim 8: If the signature scheme is existentially unforgeable against chosen-message attack with adversary advantage $\mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig}}(\mathcal{A}, \lambda)$, then $\mathcal{H}_8$ and $\mathcal{H}_7$ are indistinguishable with distinguishing advantage at most $\mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig}}(\mathcal{A}, \lambda)$.

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

Hybrid $\mathcal{H}_9$: This object is the ideal execution $\mathsf{EXEC}^{\mathcal{G}_{\mathsf{BB}}}_{\mathcal{F}^{n,k,t}_{\mathsf{Vote}}, \mathsf{DI}, \mathcal{S}, \mathcal{Z}}$.

Claim 9: If the decryption NIZK is sound with soundness error $\mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{Dec}}}(\mathcal{A}, \lambda)$, the shuffle NIZK is sound with soundness error $\mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{shuffle}}}(\mathcal{A}, \lambda)$ and the max-and-match SFE protocol is robust with adversary advantage $\mathsf{Adv}^{\mathsf{robust}}_{\mathsf{SFE}}(\mathcal{A}, \lambda)$, then $\mathcal{H}_9$ and $\mathcal{H}_8$ are indistinguishable with distinguishing advantage at most $(n + 2) \cdot \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{Dec}}}(\mathcal{A}, \lambda) + n \cdot \mathsf{Adv}^{\mathsf{robust}}_{\mathsf{SFE}}(\mathcal{A}, \lambda) + \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{shuffle}}}(\mathcal{A}, \lambda)$.

Proof: To prove Claim 9, we will show that the real tally (pre-tally) and ideal tally (pre-tally) are indistinguishable.

We first show that the real pre-tally and ideal pre-tally are indistinguishable. If decryption correctness of honest voters' ballots holds, the number of yesVotes and noVotes are identical in both real pre-tally and ideal pre-tally. In the ideal execution $\mathsf{EXEC}^{\mathcal{G}_{\mathsf{BB}}}_{\mathcal{F}^{n,k,t}_{\mathsf{Vote}}, \mathsf{DI}, \mathcal{S}, \mathcal{Z}}$, the corrupted parties' ballots are honestly tallied, while honest and coerced parties' pre-tally are simulated by randomly choosing correct number of public keys. They are indistinguishable by the verifiable shuffle in registration phase. Thus, the overall advantage in pre-tally is no more than $n \cdot \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{Dec}}}(\mathcal{A}, \lambda) + \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{shuffle}}}(\mathcal{A}, \lambda)$.

We then show that the real tally and ideal tally are indistinguishable. If the max-and-match SFE protocol is sound and decryption correctness holds, the number of nullified ballots are identical in both real tally and ideal tally. The max-and-match protocol is performed for n times. Thus, the overall advantage in tally is no more than $n \cdot \mathsf{Adv}^{\mathsf{robust}}_{\mathsf{SFE}}(\mathcal{A}, \lambda) + 2 \cdot \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{Dec}}}(\mathcal{A}, \lambda)$.

In summary, the distinguishing advantage of $\mathcal{H}_9$ and $\mathcal{H}_8$ is at most $(n+2) \cdot \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{Dec}}}(\mathcal{A}, \lambda) + n \cdot \mathsf{Adv}^{\mathsf{robust}}_{\mathsf{SFE}}(\mathcal{A}, \lambda) + \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{shuffle}}}(\mathcal{A}, \lambda)$.

Consequently, the real execution $\mathsf{EXEC}^{\mathcal{G}_{\mathsf{BB}}}_{\Pi^{n,k,t}_{\mathsf{VoteXX}}, \mathsf{DR}, \mathcal{A}, \mathcal{Z}}$ and ideal execution $\mathsf{EXEC}^{\mathcal{G}_{\mathsf{BB}}}_{\mathcal{F}^{n,k,t}_{\mathsf{Vote}}, \mathsf{DI}, \mathcal{S}, \mathcal{Z}}$

are indistinguishable with distinguishing advantage no more than

$$(2n + 2) \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{Dec}}}(\mathcal{A}, \lambda) + n \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{ballot}}}(\mathcal{A}, \lambda) +$$
$$3n \cdot \mathsf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathsf{Enc}}(\mathcal{A}, \lambda) + n \cdot \mathsf{Adv}^{\mathsf{zk}}_{\mathsf{NIZK}_{\mathsf{nul}}}(\mathcal{A}, \lambda) + n \cdot \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{nul}}}(\mathcal{A}, \lambda) +$$
$$n \cdot \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig}}(\mathcal{A}, \lambda) + (n + 2) \cdot \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{Dec}}}(\mathcal{A}, \lambda) +$$
$$n \cdot \mathsf{Adv}^{\mathsf{robust}}_{\mathsf{SFE}}(\mathcal{A}, \lambda) + \mathsf{Adv}^{\mathsf{sound}}_{\mathsf{NIZK}_{\mathsf{shuffle}}}(\mathcal{A}, \lambda)$$

This argument concludes the proof. $\square$

$\square$

C.7   Security of "Flip" Variant

In this section, we will illustrate that the above UC proof can be adapted to "flip" nullification. To prove security of "flip" nullification, the ideal functionality $\mathcal{F}^{n,k,t}_{\mathsf{Vote}}$ has interface FLIP instead of NULLIFY. Upon receiving (FLIP, sid) from a voter V_i OR V_i's hedgehog, $\mathcal{F}^{n,k,t}_{\mathsf{Vote}}$ flips V_i's vote and sends the notification to $\mathcal{S}$. The ideal deception strategy DI also needs to be modified accordingly: If DI chooses to deceive, it forwards the coercer's input to the ideal functionality and has 50% probability sending a flip request to the ideal functionality.

The construction of the real deception strategy DR_i is straightforward. If DI_i sends a flip request to the ideal functionality, DR_i will follow the protocol to flip the vote. Otherwise, DR_i sends a dummy flip request to BB. The construction of simulator $\mathcal{S}$ is completely the same except that "cancel" request is replaced by "flip" request. Then, through a series of same hybrid worlds, we can prove that "flip" protocol UC realizes $\mathcal{F}^{n,k,t}_{\mathsf{Vote}}$.

D   VoteXX with Decoy

In this section, we will show that we can achieve better coercion resistance by adding decoy ballots to VoteXX.

D.1   Intuition

A decoy ballot is a ballot that will not be counted in the final tally, but is indistinguishable from a real ballot in the coercer's view. In the registration phase, if the voter knows that he will be coerced, he will register two (or more) public keys to the distributed EA and only one of them is real, while the others are decoys. In the voting phase, the voter can use the decoy secret key to submit a ballot that will not be counted in the final tally, and use the real secret key to vote as he wants. However, if the voter cannot hide the real secret key from the coercer, he can still use VoteXX's nullification to nullify the vote. **In conclusion, adding decoy allows a voter to vote as he wants, if he can keep the real secret key from the coercer, and to nullify if he cannot hold the secret.**

To add decoy ballots into VoteXX protocol, we modify the registration phase and the final tally phase while keep the other parts same as VoteXX. We assume that the VoteXX protocol uses "cancel" nullification.

We assume that there is a public roster consisting of commitments of voters' credentials, and each voter holds his credential σ. In the registration phase, if a voter knows that he will be coerced, he (and his hedgehog) sends two (or more) $\langle \mathsf{pk}_{\mathsf{yes}}, \mathsf{pk}_{\mathsf{no}}, \rho \rangle$ tuples to the distributed EA, where $\rho \leftarrow \mathsf{Sign}_\sigma(\mathsf{pk}_{\mathsf{yes}} || \mathsf{pk}_{\mathsf{no}})$. At the end of the registration phase, the distributed EA will a perform *secure multi-party computation (MPC)* to generate the table of public keys, which is of the form $\langle \mathsf{pk}_{\mathsf{yes}}, \mathsf{pk}_{\mathsf{no}}, i \rangle$, where $i \in \{[\![0]\!], [\![1]\!]\}$ is called a "decoy flag". If $i = [\![1]\!]$, it means that the corresponding public key is decoy. For the public keys

signed by the same credential, the distributed EA uses a public function to determine which public key is real. The secure MPC ensures that none of the EA trustees know which of the public keys are decoys.

In the final tally phase, for each key in yesVotes OR noVotes, the EA takes all the encrypted flags and the corresponding "decoy flags" to compute them OR under encryption using the max-and-match SFE protocol [JJ00]. In this way, if the ballot is a decoy ballot, it will be nullified automatically since the "decoy flag" is an encryption of 1; if the ballot is a real ballot, it can be nullified the same way as in VoteXX.

D.2 Security Definition

We analyze security of VoteXX with decoy under the UC framework [Can01]. Comparing with VoteXX, the only difference is the ideal deception strategy.

Ideal deception. VoteXX with decoy realizes a stronger DI, which performs as follows: when DI receives an input x from the simulator (ideal coercer) $\mathcal{S}$, it maps x to x' and sends (VOTE, sid, x') to the ideal functionality $\mathcal{F}$. (x' can be equal OR not equal to x, representing obeying and deceiving, respectively.)

D.3 UC Specification of VoteXX with Decoy

VoteXX with decoy protocol $\Pi_{\mathsf{VoteXX-decoy}}^{n,k,t}$

Denote the voters as $\mathcal{V} := \{\mathsf{V}_1, \ldots, \mathsf{V}_n\}$, the hedgehogs as $\mathcal{H} := \{\mathsf{H}_1, \ldots, \mathsf{H}_n\}$, the trustees as $\mathcal{T} := \{T_1, \ldots, T_k\}$, and the *Election Authority* as EA. We assume that EA cannot be corrupted. At the beginning of the protocol, each voter holds his credential σ and $\mathcal{G}_{\mathsf{BB}}$ contains commitments of σ.

Preparation phase:

Upon receiving (BEGIN, sid) from the environment $\mathcal{Z}$, the EA performs the initialization procedure of the secure multi-party computation.

Registration phase:

Upon receiving (REGISTER, sid) from the environment $\mathcal{Z}$, the voter V_i performs the following:

- Send $\langle \mathsf{pk}_{\mathsf{yes}}, \mathsf{pk}_{\mathsf{no}}, \rho \rangle$ to the EA, where $\rho \leftarrow \mathsf{Sign}_\sigma(\mathsf{pk}_{\mathsf{yes}} || \mathsf{pk}_{\mathsf{no}})$, and hold the corresponding $\langle \mathsf{sk}_{\mathsf{yes}}, \mathsf{sk}_{\mathsf{no}} \rangle$.

- (If V_i will be coerced, he sends $\langle \mathsf{pk}_{\mathsf{yes}}, \mathsf{pk}_{\mathsf{no}}, \rho \rangle$ and $\langle \mathsf{pk}_{\mathsf{yes}}', \mathsf{pk}_{\mathsf{no}}', \rho' \rangle$ to the EA, where $\rho \leftarrow \mathsf{Sign}_\sigma(\mathsf{pk}_{\mathsf{yes}}, \mathsf{pk}_{\mathsf{no}})$ and $\rho' \leftarrow \mathsf{Sign}_\sigma(\mathsf{pk}_{\mathsf{yes}}' || \mathsf{pk}_{\mathsf{no}}')$. He holds the corresponding $\langle \mathsf{sk}_{\mathsf{yes}}, \mathsf{sk}_{\mathsf{no}} \rangle$ and $\langle \mathsf{sk}_{\mathsf{yes}}', \mathsf{sk}_{\mathsf{no}}' \rangle$. Let $\langle \mathsf{sk}_{\mathsf{yes}}, \mathsf{sk}_{\mathsf{no}} \rangle$ be the real key and $\langle \mathsf{sk}_{\mathsf{yes}}', \mathsf{sk}_{\mathsf{no}}' \rangle$ be the decoy key.)

Upon receiving (ENDREG, sid) from the environment $\mathcal{Z}$, the EA performs the following:

- Perform secure multi-party computation to generate a list of the form $\langle \mathsf{pk}_{\mathsf{yes}}, \mathsf{pk}_{\mathsf{no}}, i \rangle$, where $i \in \{[\![0]\!], [\![1]\!]\}$. For the public keys signed by the same credential, the EA uses a public function h to determine which public key is real. The real keys have $i = [\![0]\!]$ while the decoy keys have $i = [\![1]\!]$.

- Send (WRITE, sid, $\{\langle \mathsf{pk}_{\mathsf{yes}}, \mathsf{pk}_{\mathsf{no}}, i \rangle\}$) to $\mathcal{G}_{\mathsf{BB}}$.

Voting phase is completely the same as $\Pi_{\mathsf{VoteXX}}^{n,k,t}$.

Tally phase:

Upon receiving (TALLY, sid) from the environment $\mathcal{Z}$, EA performs the following:

- Send (READ, sid) to $\mathcal{G}_{\mathsf{BB}}$ and collect all encrypted flags and decoy flags.

- For each $\mathsf{pk}_{\mathsf{no}}$ in yesVotes, take all encrypted flags **along with the decoy flag** and compute their OR under encryption using the max-and-match SFE protocol [JJ00].

- Take the list of encrypted ORed flags, sum them under encryption, and verifiably threshold-decrypt the result.

- Subtract this value from the number of YES votes in the provisional tally to produce the final tally for YES votes.

- Repeat the above three steps for each $\mathsf{pk}_{\mathsf{yes}}$ key in noVotes.

- Denote the final tally result as τ. Send (WRITE, sid, τ, π) to $\mathcal{G}_{\mathsf{BB}}$, where π is the NIZK for SFE and decryption.

Upon receiving (RESULT, sid) from the environment $\mathcal{Z}$, the party P returns (RESULT, sid, τ).

D.4   UC Proof for VoteXX with Decoy

By constructing a similar simulator, we have the following theorem. Since we need to assume the voter knows he will be coerced in the registration phase, the VoteXX with decoy protocol is secure against static active coercion.

Theorem 4. *Assume that the NIZKs are complete, sound and zero-knowledge; the encryption scheme is IND-CPA secure; and the signature scheme is secure against existential forgery. The "VoteXX with decoy" protocol $\Pi^{n,k,t}_{\mathsf{VoteXX\text{-}decoy}}$ IUC realizes $\mathcal{F}^{n,k,t}_{\mathsf{Vote}}$ against static corruption and static active coercion in the $\mathcal{G}_{\mathsf{BB}}$-hybrid model.*

Proof sketch. To prove the theorem, we need to 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 $\mathsf{EXEC}^{\mathcal{G}_{\mathsf{BB}}}_{\Pi^{n,k,t}_{\mathsf{VoteXX\text{-}decoy}},\mathsf{DR},\mathcal{A},\mathcal{Z}}$ from the (ii) the ideal execution $\mathsf{EXEC}^{\mathcal{G}_{\mathsf{BB}}}_{\mathcal{F}^{n,k,t}_{\mathsf{Vote}},\mathsf{DI},\mathcal{S},\mathcal{Z}}$.

Real Deception Strategy. The real deception strategy DR_i internally runs DI_i, forwarding messages to and from the outside. DR_i works as follows:

In the registration phase, upon receiving (REGISTER, sid) from the environment $\mathcal{Z}$, V_i sends $\langle \mathsf{pk}_{\mathsf{yes}}, \mathsf{pk}_{\mathsf{no}}, \rho \rangle$ and $\langle \mathsf{pk}'_{\mathsf{yes}}, \mathsf{pk}'_{\mathsf{no}}, \rho' \rangle$ to the EA, where $\rho \leftarrow \mathsf{Sign}_\sigma(\mathsf{pk}_{\mathsf{yes}}, \mathsf{pk}_{\mathsf{no}})$ and $\rho' \leftarrow \mathsf{Sign}_\sigma(\mathsf{pk}'_{\mathsf{yes}} || \mathsf{pk}'_{\mathsf{no}})$, and holds the corresponding $\langle \mathsf{sk}_{\mathsf{yes}}, \mathsf{sk}_{\mathsf{no}} \rangle$ (real key) and $\langle \mathsf{sk}'_{\mathsf{yes}}, \mathsf{sk}'_{\mathsf{no}} \rangle$ (decoy key).

In the voting phase:

- When coerced, DR_i provides the decoy key $\mathsf{sk}'_{\mathsf{yes}}, \mathsf{sk}'_{\mathsf{no}}$ to the coercer and simulates the transcript of using $\mathsf{sk}'_{\mathsf{yes}}, \mathsf{sk}'_{\mathsf{no}}$ to register. DR_i follows the coercer's instructions.

- If DI_i sends (VOTE, sid, x') to the ideal functionality, DR_i performs as if V_i receives (VOTE, sid, x') from the environment $\mathcal{Z}$.

Simulator. The simulator $\mathcal{S}$ is almost the same as the simulator for $\Pi^{n,k,t}_{\mathsf{VoteXX}}$ since we only modified the registration phase and tally phase.

In the registration phase, the simulator $\mathcal{S}$ simulates the parties following the protocol:

- Upon receiving (REGISTER, sid, V_i) from the functionality $\mathcal{F}^{n,k,t}_{\mathsf{Vote}}$, the simulator $\mathcal{S}$ simulates V_i following the protocol $\Pi^{n,k,t}_{\mathsf{VoteXX\text{-}decoy}}$ as if V_i received (REGISTER, sid) from the environment $\mathcal{Z}$.

- Upon receiving (ENDREG, sid) from the functionality $\mathcal{F}^{n,k,t}_{\mathsf{Vote}}$, the simulator $\mathcal{S}$ simulates EA following the protocol $\Pi^{n,k,t}_{\mathsf{VoteXX\text{-}decoy}}$ as if the EA received (ENDREG, sid) from the environment $\mathcal{Z}$.

In the tally phase, the simulator $\mathcal{S}$ simulates the EA to perform the tally procedure, but uses the NIZK simulator to generate the NIZK proof:

- Upon receiving $(\textsc{Tally}, \mathsf{sid}, \tau)$ from the functionality $\mathcal{F}_{\mathsf{Vote}}^{n,k,t}$, $\mathcal{S}$ simulates EA doing the following:

  - For each pk in yesVotes and noVotes, take all the encrypted flags **along with the decoy flag**, compute its OR under encryption. Sum them under encryption for yesVotes and noVotes, respectively.
  - Simulate the decryption of the summed encrypted flags and the corresponding NIZK π based on τ.
  - Send $(\textsc{Write}, \mathsf{sid}, \tau, \pi)$ to $\mathcal{G}_{\mathsf{BB}}$.

Same as the proof of **Theorem A.2**, through a series of hybrid worlds, the real world and ideal world are indistinguishable.