



Amun: Securing E-Voting Against Over-the-Shoulder Coercion

Riccardo Longo¹ ^a and Chiara Spadafora² ^b

¹Fondazione Bruno Kessler, Center for Cybersecurity, Trento, Italy

²Department of Mathematics, University of Trento, 38123 Povo, Trento, Italy

Keywords: E-Voting, Over-the-Shoulder Attack, Coercion-Resistance, End-to-End Verifiability, Formal Proof of Security.

Abstract: In an election where each voter may express P preferences among M possible choices, the Amun protocol allows to secure vote casting against over-the-shoulder adversaries, retaining privacy, fairness, end-to-end verifiability, and correctness. We prove the security of the construction under the standard Decisional Diffie Hellman assumption in the random oracle model.

1 INTRODUCTION

Remote voting, enhanced by advanced cryptographic techniques, offers more security than traditional paper-based voting but raises privacy concerns when voting outside secure booths.

Although many systems protect against various adversaries who try to bribe electors, we found that it is more difficult to counter opponents that closely monitor voters during the voting phase (over-the-shoulder attacks). The main mitigation technique against coercion is the usage of fake credentials (Juels et al., 2010), which are indistinguishable from real ones but that do not produce valid votes. However, if the adversary keeps the voter under control until the end of the voting period, it becomes impossible to re-vote with the valid credential.

Here we present the Amun¹ protocol, which hides the real choice expressed by a voter even if an adversary is physically monitoring the elector during vote casting. This feature protects the elector against over-the-shoulder attacks without the need to re-vote. The Amun protocol aims to achieve end-to-end verifiability, universal verifiability, privacy, correctness, fairness, and coercion resistance.


The authors in (Spadafora et al., 2021) suggest a blockchain-based remote e-voting protocol for two candidates, where voters have two voting tokens (v -


tokens): one is valid and the other is a decoy, but only the voter knows which is which. The Amun protocol extends this concept to support multiple candidates and choices.


In a setting where the voter can choose P out of M candidates, extra care has to be taken in the design of the v -tokens. In fact, a naïve straightforward adaptation of (Spadafora et al., 2021) might reveal whether two v -tokens have the same validity, since there are multiple valid and decoy v -tokens. To deal with this we had to introduce extra masking steps and another authority to fully hide the v -tokens validity. So, in Amun, three authorities share the administration of the election: they setup the parameters, manage voters' registration, and compute the final tally at the end of the voting phase. Privacy is preserved even if an attacker colludes with one authority, limiting their power. As in (Spadafora et al., 2021), votes are cast by assigning to the candidates some "voting tokens" generated during registration. Among these tokens, only a few are valid and express the real preference of the voter, but they are indistinguishable from the other, decoy, tokens. This trick disguises the actual choice made, even if the adversary is watching.

1.1 Related Work

Protocols for electronic election systems have been abundantly proposed in recent years. Many have addressed the problem of coercion resistance, giving a plurality of definitions (Gardner et al., 2009; Haines and Smyth, 2019; Juels et al., 2010; Krips and Willemson, 2019; Kusters et al., 2010). *Civitas* (Clarkson et al., 2008), which derives from

^a  <https://orcid.org/0000-0002-8739-3091>

^b  <https://orcid.org/0000-0003-3352-9210>

¹Amun was a major ancient Egyptian deity. The name Amun  meant something like "the hidden one" or "invisible".

JCJ (Juels et al., 2010), deals with coercion by allowing voters to vote multiple times via a mechanism of *real* and *fake* credentials. *Selene* (Ryan et al., 2016) associates to every vote a unique tracker: the idea is that, in case of an attack, every voter is able to open up its commitment to a fake tracker in order to deceive the attacker. *Belenios* (Cortier et al., 2019) itself is not coercion resistant: voters can keep the randomness used to encrypt the ballot to prove how they voted. This limitation has been overcome with *BeleniosRF* (Chaidos et al., 2016).

Organization. We present some preliminaries in Section 2. We describe our protocol in Section 3 and we provide a proof of security in Section 5. Finally, in Section 6 we draw some conclusions.

2 PRELIMINARIES

Most of the algebraic preliminaries we need to build the protocol, such as the *Decisional Diffie-Hellman Assumption* (DDH) and commitment schemes, can be found in (Spadafora et al., 2021). Here we report the additional tools required for the generalization, for more details see (Longo and Spadafora, 2021).

We use the following notation for the indexes: $[n] = \{i \in \mathbb{N} : 1 \leq i \leq n\}$, $(t_j)_{j \in [m]} = (t_1, \dots, t_m)$.

2.1 Equality of Discrete Logarithms

In the protocol it is used a Zero-Knowledge Proof (ZKP) for the equality of two discrete logarithms (Spadafora et al., 2021; Longo and Spadafora, 2021), which is a variation of the Schnorr interactive protocol (Schnorr, 1991; Shoup and Alwen, 2007).

2.2 Non-Interactive ZKP

Non-Interactive ZKP (NIZKP) are a special type of ZKPs that allow the prover to publish a proof that can be independently verified by all the relevant parties later on. The Fiat-Shamir technique (Iovino and Visconti, 2019) can be used to transform an interactive sigma protocol into a NIZKP by exploiting a hash function modeled as a random oracle (RO).

2.3 Designated-Verifier ZKP

Designated-Verifier Non-Interactive ZKP systems (DVNIZKPs (Jakobsson et al., 2001)) are protocols which retain most of the security properties of a NIZKP, but are not publicly verifiable: only the owner

of some secret (the designated verifier) can check the proof. This property is useful in the context of e-voting to achieve end-to-end verifiability while still preventing the voter from transferring some proofs. In (Longo and Spadafora, 2021) there is a complete description of the protocol.

2.4 General Requirements for Remote Voting Systems

A trustworthy e-voting protocol has to satisfy conflicting requirements: it should preserve both *integrity* of election results and *confidentiality* of votes. In (Longo and Spadafora, 2021) are defined the properties that a trustworthy e-voting protocol should fulfill (namely *Correctness*, *Fairness*, *Privacy*, *Verifiability*). We define here only *Vote-Coercion Resistance* since we adopt a non-standard definition, the proof that our proposed protocol satisfies it is in Section 5.2.

Coercion resistance (Haines and Smyth, 2019) requires that an adversary cannot learn any additional information about the votes other than what is revealed by the results of tabulation. In other words, voters cannot prove whether or how they voted, even if they can interact with the adversary while voting. The Amun protocol protects against coercers that wish to sway elections towards specific candidates, but is not very effective against randomization and forced abstention. In this simplified model, we use a slightly weaker adaptation of the definition of Coercion Resistance given in (Juels et al., 2010):

Definition 1 (Vote-Coercion Resistance). *Let \mathcal{A} be a coercer, V_c the set of coerced voters, and $(C_{i,1}, \dots, C_{i,P})$ the choices that \mathcal{A} wants to impose to the voter corresponding to $v_i \in V_c$. Let Ψ_1 be the scenario in which \mathcal{A} has access only to the final tally. Let Ψ_2 be the scenario in which \mathcal{A} has access to the whole Bulletin Board, and can see all the actions performed by the voters in V_c , with the exception of the ones carried out in a protected environment (or through an untappable channel). A voting protocol is Vote-Coercion Resistant if the probability of \mathcal{A} detecting that a voter in V_c has not followed its instruction is the same in Ψ_1 and Ψ_2 .*

3 MULTI-CANDIDATE E-VOTING

This section presents our proposal for a remote e-voting protocol that manages an election with N voters, where each one expresses P preferences among M candidates (obviously $P < M$). The basic idea is that every voter owns M voting tokens (*v-tokens*): P

are valid, the others are a decoy, but only the voter knows which is which. When voting, voters express their preferences assigning the valid v -tokens to the chosen candidates and the decoy ones to the others.

The protocol allows for re-voting: before tallying duplicate ballots (i.e. with the same v -tokens ignoring their order) are discarded, keeping only the most recent. After the voting phase, when counting the votes, the decoy v -tokens do not contribute to the tally, so only valid v -tokens are counted. The whole process is publicly auditable and fully verifiable, and preserves privacy as long as at most one authority is corrupt.

The protocol is divided into four phases:

- **Setup.** Three authorities, knowing a list of eligible voters, generate the values for the creation of both the v -tokens and the masks associated to the candidates. These masks guarantee the voters' privacy, and prevent early tallying.
- **Registrar Phase.** In this phase, the three authorities engage in a 5-step protocol (see Figure 1) to create M indistinguishable v -tokens (P are valid and $M - P$ are a decoy) employing masking and shuffling so that at the end the authorities will not be able to identify which tokens are valid. The voter can check the validity of these v -tokens thanks to DVNIZKPs issued by the authorities. These proofs are worthless for a coercer because the voter can forge them.




Figure 1: The main steps of the ballot generation procedure, which correspond to steps 3-7 of the Registrar Phase as described in Section 3.1.2.

- **Voting Phase.** During this phase the voter express their preferences by assigning each of their M v -tokens to the candidates. All v -tokens of a voter must be assigned together, each to a distinct candidate. After the v -tokens have been assigned, the voter gets a transcript that reports the assignment of the v -tokens to the candidates. This transcript is worthless for a coercer since the v -tokens are indistinguishable. Here we assume that every candidate receives at least one legitimate vote (with a valid v -token), otherwise it is trivial to discern the validity of some tokens from the election results.
- **Tallying.** The v -tokens are processed (see Figure 2), removing the candidate masks, which al-

lows to count the number of valid and decoy tokens assigned to each candidate. The results and the intermediate computations are published, alongside a set of NIZKPs that allow anyone to check that the results are correct and there has not been any tampering. Every voter can also check, by examining the bulletin board, that their v -tokens have been cast and counted correctly.




Figure 2: Example of voting and tallying. Each voter has two valid tokens \bullet and two decoy tokens \circ . After the tallying it is revealed that candidates C_3 and C_4 are elected having received more preferences (3) with respect to the other two candidates (who received only 2).

3.1 Protocol Description

The key components involved in the protocol are:

1. a finite set of voters $V = \{v_i\}_{i \in [N]}$ (where v_i is a pseudonymous id), with $N \in \mathbb{N}$ the number of eligible voters;
2. a finite set of candidates $C = \{c_\ell\}_{\ell \in [M]}$ with $M \in \mathbb{N}$ the number of candidates;
3. three trusted authorities² $\mathcal{A}_0, \mathcal{A}_1,$ and \mathcal{A}_2 .
4. one ballot b_i (comprising M v -tokens) for every $i \in [N]$, i.e. one for each eligible voter.

Throughout the protocol we implicitly assume that every public value (including a description of the key components presented above) are published in the BB. The protocol is divided into four phases.

3.1.1 Setup

The authority \mathcal{A}_0 selects and publishes:

1. a secure group \mathbb{G} of prime order p in which the DDH assumption holds, with a generator $g \in \mathbb{G}$;
2. a commitment scheme Comm to be used to commit to the values computed before publishing them, in order to improve security.

²We use a weak concept of trust here, since the conduct of these authorities can be checked by voters.

Then \mathcal{A}_0 performs the following operations:

1. chooses uniformly at random two values k and λ in \mathbb{Z}_p^* . \mathcal{A}_0 knows that the v -tokens computed using k are valid, while the ones computed using λ are decoys, but this information is kept secret;
2. chooses uniformly at random $N \cdot M$ distinct values $\bar{z}_{i,\ell} \in \mathbb{Z}_p^*$, with $i \in [N], \ell \in [M]$;
3. finally, \mathcal{A}_0 commits to the values g^k, g^λ , and, for every $i \in [N]$, it commits to $(v_i, (g^{\bar{z}_{i,\ell}})_{\ell \in [M]})$.

An honest authority \mathcal{A}_0 is supposed to keep private all the values $\bar{z}_{i,\ell}, k, \lambda$.

The authority \mathcal{A}_1 performs the following operations:

1. chooses uniformly at random M distinct values $\alpha'_\ell \in \mathbb{Z}_p^*$, with $\ell \in [M]$, these will be the first half of the candidates' masks;
2. chooses uniformly at random N distinct values $x'_i \in \mathbb{Z}_p^*$, with $i \in [N]$;
3. chooses uniformly at random two sets of $N \cdot M$ distinct values $z'_{i,\ell}, y'_{i,\ell} \in \mathbb{Z}_p^*$, with $i \in [N], \ell \in [M]$;
4. finally, \mathcal{A}_1 commits to the values $g^{\alpha'_\ell}, \forall \ell \in [M]$, and for every $i \in [N]$ it commits to the tuple $(v_i, g^{x'_i}, (g^{z'_{i,\ell}})_{\ell \in [M]}, (g^{y'_{i,\ell}})_{\ell \in [M]})$.

An honest authority \mathcal{A}_1 is supposed to keep private the values $\alpha'_\ell, x'_i, z'_{i,\ell}, y'_{i,\ell}$.

The authority \mathcal{A}_2 performs the following operations:

1. chooses uniformly at random M distinct values $\alpha''_\ell \in \mathbb{Z}_p^*$, with $\ell \in [M]$, these will be the second half of the candidates' masks;
2. chooses uniformly at random N distinct values $x''_i \in \mathbb{Z}_p^*$, with $i \in [N]$;
3. chooses uniformly at random $N \cdot M$ distinct values $y''_{i,\ell} \in \mathbb{Z}_p^*$, with $i \in [N], \ell \in [M]$;
4. Finally \mathcal{A}_2 commits to the values $g^{\alpha''_\ell}, \forall \ell \in [M]$, and for every $i \in [N]$ it commits to the tuple $(v_i, g^{x''_i}, (g^{y''_{i,\ell}})_{\ell \in [M]})$.

An honest authority \mathcal{A}_2 is supposed to keep private all the values $\alpha''_\ell, x''_i, y''_{i,\ell}$.

Once that all the commitments have been published, the authorities can decommit the values:

- \mathcal{A}_0 publishes the decommitments for the values g^k, g^λ , alongside all the tuples $(v_i, (g^{\bar{z}_{i,\ell}})_{\ell \in [M]}) \forall i \in [N]$;
- \mathcal{A}_1 publishes the decommitments for the values $g^{\alpha'_\ell} \forall \ell \in [M]$, and the tuples $(v_i, g^{x'_i}, (g^{z'_{i,\ell}})_{\ell \in [M]}, (g^{y'_{i,\ell}})_{\ell \in [M]}) \forall i \in [N]$;

- \mathcal{A}_2 publishes the decommitments for the values $g^{\alpha''_\ell} \forall \ell \in [M]$, and the tuples $(v_i, g^{x''_i}, (g^{y''_{i,\ell}})_{\ell \in [M]}) \forall i \in [N]$.

All these published values are accompanied by NIZKPs which prove that the authority who published them knows the corresponding secret exponents. These NIZKPs can be constructed using the Schnorr protocol and the Fiat-Shamir transformation just like in Section 2.2.

To simplify notation we introduce some definitions for aggregate values for all $i \in [N]$ and $\ell \in [M]$:

$$\begin{aligned} x_i &= x'_i + x''_i, & \alpha_\ell &= \alpha'_\ell \cdot \alpha''_\ell, \\ z_{i,\ell} &= \bar{z}_{i,\ell} \cdot z'_{i,\ell}, & y_{i,\ell} &= y'_{i,\ell} \cdot y''_{i,\ell}. \end{aligned}$$

3.1.2 Registrar Phase

For every pseudonymous id $v_i \in V$ the following steps are performed:

1. Let Alice be the person associated to the pseudonymous id v_i , note that the authorities do not need to know this association. She goes in a safe and controlled environment (see Section 6 for further discussion on this requirement) where she is identified and authenticated as the eligible and not yet registered pseudonymous id v_i . In this environment she can interact with all three authorities without fear of eavesdropping or interference.
2. Alice creates a signing key-pair (s_i, K_i) , a designated verifier key-pair (e_i, D_i) , and gives K_i, D_i to the authorities proving the knowledge of s_i (e.g. by signing a challenge message), and of e_i via a NIZKP (which includes the challenge message among the public values). The authorities associate K_i, D_i to v_i in their respective voters lists.
3. \mathcal{A}_0 performs the following steps:
 - (a) \mathcal{A}_0 chooses, for every $i \in [N]$, a random subset $V_i \subset [M]$ with cardinality is exactly P , then sets:

$$\sigma_{i,\ell} = \begin{cases} k & \iff \ell \in V_i \\ \lambda & \iff \ell \notin V_i \end{cases}$$

i.e. the random choice of the V_i determines which tokens will be valid and which a decoy;

- (b) \mathcal{A}_0 takes the (publicly available) values $g^{x'_i}$ and $g^{x''_i}$ and creates the step 0 of the ballot $\bar{b}_{0,i} = (\bar{b}_{0,i,\ell})_{\ell \in [M]}$ where, $\forall \ell \in [M]$:

$$\bar{b}_{0,i,\ell} = \left(g^{\sigma_{i,\ell}} \cdot g^{x'_i} \cdot g^{x''_i} \right)^{\bar{z}_{i,\ell}} = g^{\bar{z}_{i,\ell}(\sigma_{i,\ell} + x_i)}$$

- (c) \mathcal{A}_0 sends to \mathcal{A}_1 the initial ballot $\bar{b}_{0,i}$ and sends to Alice $\bar{b}_{0,i}$ and V_i ;

(d) \mathcal{A}_0 proves its computations correct with multiple instances of the DVNIZKP of Section 2.3:

i. \mathcal{A}_0 proves that the $g^{\bar{z}_{i,\ell}\sigma_{i,\ell}}$ are correct (using $\sigma_{i,\ell} = k$ or $\sigma_{i,\ell} = \lambda$) with:

$$\omega = k, \quad u = g, \quad z = g^k,$$

$$\bar{u} = g^{\bar{z}_{i,\ell}}, \quad \bar{z} = g^{\bar{z}_{i,\ell}k}, \quad \forall \ell \in V_i,$$

$$\omega = \lambda, \quad u = g, \quad z = g^\lambda,$$

$$\bar{u} = g^{\bar{z}_{i,\ell}}, \quad \bar{z} = g^{\bar{z}_{i,\ell}\lambda} \quad \forall \ell \in [M] \setminus V_i,$$

ii. then \mathcal{A}_0 proves that the $\bar{b}_{0,i,\ell}$ are correct using for all $\ell \in [M]$:

$$\omega = \bar{z}_{i,\ell}, \quad u = g, \quad z = g^{\bar{z}_{i,\ell}},$$

$$\bar{u} = g^{\sigma_{i,\ell}} \cdot g^{x_i} \cdot g^{x_i''}, \quad \bar{z} = \bar{b}_{0,i,\ell}.$$

4. \mathcal{A}_1 computes the step 1 of the ballot $\bar{b}_{1,i} = (\bar{b}_{1,i,\ell})_{\ell \in [M]}$ where:

$$\bar{b}_{1,i,\ell} = (\bar{b}_{0,i,\ell})^{z'_{i,\ell}} = g^{z_{i,\ell}(\sigma_{i,\ell} + x_i)} \quad \forall \ell \in [M]$$

and sends it to Alice and to \mathcal{A}_2 . Then \mathcal{A}_1 proves that the $\bar{b}_{1,i,\ell}$ are correct with the DVNIZKP of Section 2.3, using:

$$\omega = z'_{i,\ell}, \quad u = g, \quad z = g^{z'_{i,\ell}},$$

$$\bar{u} = \bar{b}_{0,i,\ell}, \quad \bar{z} = \bar{b}_{1,i,\ell} \quad \forall \ell \in [M].$$

5. \mathcal{A}_2 chooses uniformly at random a permutation $\pi_i \in \text{Sym}([M])$ and computes the step 2 of the ballot $\bar{b}_{2,i} = (\bar{b}_{2,i,\ell})_{\ell \in [M]}$ where, $\forall \ell \in [M]$:

$$\bar{b}_{2,i,\ell} = (\bar{b}_{1,i,\ell})^{y''_{i,\pi_i^{-1}(\ell)}} = g^{z_{i,\ell} y''_{i,\pi_i^{-1}(\ell)} (\sigma_{i,\ell} + x_i)}$$

and sends it to Alice and to \mathcal{A}_0 , π_i is sent to Alice and \mathcal{A}_1 . Then \mathcal{A}_2 proves that the $\bar{b}_{2,i,\ell}$ are correct with the DVNIZKP of Section 2.3, using:

$$\omega = y''_{i,\pi_i^{-1}(\ell)}, \quad u = g, \quad z = g^{y''_{i,\pi_i^{-1}(\ell)}},$$

$$\bar{u} = \bar{b}_{1,i,\ell}, \quad \bar{z} = \bar{b}_{2,i,\ell} \quad \forall \ell \in [M].$$

6. \mathcal{A}_0 computes the step 3 of the ballot $\bar{b}_{3,i} = (\bar{b}_{3,i,\ell})_{\ell \in [M]}$ where, $\forall \ell \in [M]$:

$$\bar{b}_{3,i,\ell} = (\bar{b}_{2,i,\ell})^{\frac{1}{z_{i,\ell}}} = g^{z'_{i,\ell} y''_{i,\pi_i^{-1}(\ell)} (\sigma_{i,\ell} + x_i)}$$

and sends it to Alice and to \mathcal{A}_1 . Then \mathcal{A}_0 proves that the $\bar{b}_{3,i,\ell}$ are correct with the DVNIZKP of Section 2.3, using:

$$\omega = \frac{1}{z_{i,\ell}}, \quad u = g^{\bar{z}_{i,\ell}}, \quad z = g,$$

$$\bar{u} = \bar{b}_{2,i,\ell}, \quad \bar{z} = \bar{b}_{3,i,\ell} \quad \forall \ell \in [M].$$

7. \mathcal{A}_1 computes the final ballot $b_i = (b_{i,\ell})_{\ell \in [M]}$, with:

$$b_{i,\ell} = (\bar{b}_{3,i,\pi_i(\ell)})^{\frac{y'_{i,\ell}}{z'_{i,\pi_i(\ell)}}} = g^{y_{i,\ell}(\sigma_{i,\pi_i(\ell)} + x_i)} \quad \forall \ell \in [M]$$

and sends it to Alice and publishes on the BB the pair (K_i, b_i) . Then \mathcal{A}_1 proves that the $b_{i,\ell}$ are correct with the DVNIZKP of Section 2.3 and using:

$$\omega = \frac{y'_{i,\ell}}{z'_{i,\pi_i(\ell)}}, \quad u = g^{z'_{i,\pi_i(\ell)}}, \quad z = g^{y'_{i,\ell}},$$

$$\bar{u} = \bar{b}_{3,i,\pi_i(\ell)}, \quad \bar{z} = b_{i,\ell} \quad \forall \ell \in [M].$$

Note that Alice, thanks to the proofs and the knowledge of the intermediate values, knows which ones are a valid token (the ones with $\sigma_{i,\ell} = k$), but thanks to the random choices of V_i and π_i the authorities cannot distinguish the tokens unless they collude. Effectively, the DVNIZKPs prove to Alice that the ballot has been created by \mathcal{A}_0 with the correct number of valid and decoy tokens, and that it has been correctly shuffled by \mathcal{A}_2 . Moreover the properties of the DVNIZKP allow Alice to forge the transcript changing which tokens are valid, making them useless for proving the validity of a token. In fact, since Alice is in a protected environment, she can manipulate the received data without being able to prove or disprove any manipulation. So, given that she knows e_i , she can forge a proof that states the presumed validity of any P of the M tokens, making any proof worthless to a coercer.

3.1.3 Voting Phase

Voters assign their valid tokens to preferred candidates and decoy tokens to the others. Each voter signs this assignment with their private key and publishes it on the BB for verification. After voting, duplicate, incomplete, and forged ballots are filtered out (relying on public data only).

3.1.4 Tallying

Once the voting phase is over, the tallying can start.

In order to count the votes, the authorities have to process the tokens received by each candidate, substituting the *voter's masks* $y_{i,\ell}$ with the appropriate *candidate mask* α_ℓ . Suppose that $T \leq N$ participants voted. Without loss of generality, we can assume that only the participants with index $i \in [T]$ voted, while the remaining $N - T$ abstained from voting.

For every $i \in [T]$, let $\phi_i : [M] \rightarrow [M]$ be the bijective map that associates to each candidate index ℓ the index of the token $b_{i,\phi_i(\ell)}$ that the voter associated to v_i sent to the candidate C_ℓ . Then, for every $i \in [T], \ell \in [M]$, the authorities process the token $b_{i,\phi_i(\ell)}$ by performing the following steps:

1. \mathcal{A}_1 computes and publishes the preliminary vote $\bar{t}_{\ell,i}$ as:

$$\bar{t}_{\ell,i} = (b_{i,\phi_i(\ell)})^{\frac{\alpha'_\ell}{y'_{i,\phi_i(\ell)}}} = g^{\alpha'_\ell y''_{i,\phi_i(\ell)} (\sigma_{i,\pi_i(\phi_i(\ell))} + x_i)},$$

alongside a NIZKP that proves this computation correct. \mathcal{A}_1 proves that $\bar{t}_{\ell,i}$ is correct with the NIZKP version of Section 2.1 and using:

$$\begin{aligned} \omega &= \frac{\alpha'_\ell}{y'_{i,\phi_i(\ell)}}, & u &= g^{y'_{i,\phi_i(\ell)}}, & z &= g^{\alpha'_\ell}, \\ \bar{u} &= b_{i,\phi_i(\ell)}, & \bar{z} &= \bar{t}_{\ell,i}. \end{aligned}$$

2. \mathcal{A}_2 then computes and publishes the final vote $t_{\ell,i}$:

$$t_{\ell,i} = (\bar{t}_{\ell,i})^{\frac{\alpha''_\ell}{y''_{i,\phi_i(\ell)}}} = g^{\alpha''_\ell (\sigma_{i,\pi_i(\phi_i(\ell))} + x_i)},$$

alongside a NIZKP that proves this computation correct. \mathcal{A}_2 proves that $t_{\ell,i}$ is correct with the NIZKP version of Section 2.1 and using:

$$\begin{aligned} \omega &= \frac{\alpha''_\ell}{y''_{i,\phi_i(\ell)}}, & u &= g^{y''_{i,\phi_i(\ell)}}, & z &= g^{\alpha''_\ell}, \\ \bar{t} &= b_{\ell,i}, & \bar{z} &= t_{\ell,i}. \end{aligned}$$

Once that all final votes have been computed, the actual tallying is performed.

Let R_ℓ be the number of valid tokens given to the ℓ -th candidate (i.e. the number of preferences received by said candidate), and let F_ℓ be the number of decoy tokens given to the ℓ -th candidate. Clearly $T = R_\ell + F_\ell \quad \forall \ell \in [M]$. The count R_ℓ can be computed with the following steps:

1. Both \mathcal{A}_1 and \mathcal{A}_2 can compute g^{α_ℓ} (as $(g^{\alpha'_\ell})^{\alpha'_\ell}$ and $(g^{\alpha''_\ell})^{\alpha''_\ell}$ respectively). \mathcal{A}_1 can prove the correctness of this value by publishing a NIZKP. \mathcal{A}_2 can prove the correctness of this value by publishing a NIZKP. In practice, each authority may publish half of the values.
2. \mathcal{A}_0 computes and publishes $g^{\alpha_\ell k} = (g^{\alpha_\ell})^k$ and $g^{\alpha_\ell \lambda} = (g^{\alpha_\ell})^\lambda$. Then \mathcal{A}_0 proves that $g^{\alpha_\ell k}$ and that $g^{\alpha_\ell \lambda}$ are correct by publishing two NIZKPs.
3. \mathcal{A}_1 computes $\sum_{i=1}^T x'_i$, and publishes $g^{\alpha_\ell \sum_{i=1}^T x'_i}$. Then \mathcal{A}_1 proves that $g^{\alpha_\ell \sum_{i=1}^T x'_i}$ is correct by publishing a NIZKP, noting that any observer can compute the value $g^{\sum_{i=1}^T x'_i} = \prod_{i=1}^T g^{x'_i}$.
4. Similarly, \mathcal{A}_2 computes $\sum_{i=1}^T x''_i$ and publishes $g^{\alpha_\ell \sum_{i=1}^T x''_i}$. Then \mathcal{A}_2 proves that $g^{\alpha_\ell \sum_{i=1}^T x''_i}$ is correct by publishing a NIZKP, again, anyone can compute $g^{\sum_{i=1}^T x''_i} = \prod_{i=1}^T g^{x''_i}$.

5. Given that any observer can compute the value:

$$g^{\alpha_\ell (\sum_{i=1}^T x_i + R_\ell k + F_\ell \lambda)} = \prod_{i=1}^T t_{\ell,i},$$

and that:

$$g^{\alpha_\ell \sum_{i=1}^T x_i} = g^{\alpha_\ell \sum_{i=1}^T (x'_i + x''_i)} = g^{\alpha_\ell \sum_{i=1}^T x'_i} \cdot g^{\alpha_\ell \sum_{i=1}^T x''_i},$$

then anyone can compute:

$$\begin{aligned} \mathfrak{T} &= \left(g^{\alpha_\ell \sum_{i=1}^T x_i} \right)^{-1} \cdot g^{\alpha_\ell (\sum_{i=1}^T x_i + R_\ell k + F_\ell \lambda)} \\ &= \left(g^{\alpha_\ell k} \right)^{R_\ell} \cdot \left(g^{\alpha_\ell \lambda} \right)^{F_\ell}. \end{aligned}$$

6. R_ℓ and F_ℓ can now be computed by brute force, giving the number of preferences received by the ℓ -th candidate.

Given a positive integer $T \in \mathbb{N}$, it is possible to represent it in $T + 1$ ways as a sum of two non-negative integers. Given that the number of valid and decoy votes must sum up to the number of actual voters T , it follows that the number of possible values for \mathfrak{T} is $T + 1$, so the effort of computing R_ℓ and F_ℓ is linear in the number of actual votes.

4 USABILITY

In order to cast a vote, the voter has to remember which are the P valid v -tokens among the M in their ballot. This can be an usability issue when P and M grow. To help the voter remembering the position of the valid tokens, we can exploit error correcting codes. We can see the information on which tokens are valid as a binary vector of \mathbb{F}_2^M with constant weight P . We can exploit constant-weight codes (Fischer and Stern, 1996) to encode these vectors as a vector of the space $\mathbb{F}_q^{\mathcal{Z}}$ and then use a $[n, \mathcal{Z}]_q$ shortened Reed-Solomon code (Roth, 2006) to add error-correction capabilities. With this approach the voter has only to remember n elements of \mathbb{F}_q , with the added bonus that up to $\frac{n-\mathcal{Z}}{2}$ errors can be automatically corrected.

Finally, we highlight that, with the DVNIZKPs and the permutation received during the registrar phase, the voter can check whether they remember correctly the positions of the valid tokens.

5 SECURITY ANALYSIS

The goal is to prove that an adversary cannot distinguish between valid and decoy v -tokens and guess

how voters cast their preferences. Since election results are obviously public, we have to avoid some trivial cases in which the adversary can deduce the votes by simply observing the results.

Therefore we assume that the adversary controls one authority and all but two voters, and that these two voters express distinct preferences. In particular, we let the adversary select two distinct sets of preferences, then we randomly assign to each of the two uncorrupted voters one set of these sets of preferences. The adversary wins the security game if it guesses correctly which voter expressed which set of preferences, i.e. guesses the random assignment.

5.1 Security Model

The security of the protocol will be proven in terms of vote indistinguishability (VI), as detailed in Definition 3. We will assume the presence of a malicious authority, so the simulator in the proof will take on the roles of the two honest authorities and of the two voters that the adversary does not control.

To simplify our analysis we assume that the adversary-controlled authority does not intentionally fail decommitments or (DV)NIZKPs, so the protocol does not abort. This is a reasonable assumption considering the application context, however it is not necessary to attain security. In fact, if the adversary wins the security game with non-negligible advantage, then it must run the protocol smoothly with non-negligible probability (since it outputs its guess only once the protocol has correctly terminated).

Definition 2 (Security Game). *The security game for the election protocol proceeds as follows:*

- **Init.** *The adversary \mathcal{A} chooses the authority and the $N - 2$ voters that it will control (i.e. the adversary knows which are the valid and decoy v -tokens of these voters). The remaining two voters are called free voters. The challenger \mathcal{C} takes the role of the other authorities and the free voters.*
- **Phase 0.** *\mathcal{A} and \mathcal{C} run the Setup and Registrar phases of the protocol, interacting as needed.*
- **Phase 1.** *The adversary votes with some or all of the voters it controls.*
- **Challenge.** *The challenge phase is articulated as follows:*
 1. *\mathcal{A} selects two distinct sets of preferences $\tilde{P}_0 \neq \tilde{P}_1$, with $\tilde{P}_i \subset [M]$, $\#\tilde{P}_i = P$ for $i = 0, 1$, and sends them to \mathcal{C} ;*
 2. *\mathcal{C} flips a random coin $\mu \in \{0, 1\}$ to determine which preference set the first free voter will use, i.e. $P_1 = \tilde{P}_\mu$, setting also $P_2 = \tilde{P}_{\mu \oplus 1}$;*

3. *\mathcal{C} constructs two random ballot assignment maps $\tilde{\phi}_1, \tilde{\phi}_2 : [M] \rightarrow [M]$ such that $\tilde{\phi}_i(\ell)$ refers to a valid token if and only if $\ell \in P_i$, for $i = 1, 2$;*
4. *finally, \mathcal{C} votes by sending to the candidate C_ℓ , $\forall \ell \in [M]$, the $\tilde{\phi}_1(\ell)$ -th and $\tilde{\phi}_2(\ell)$ -th tokens of the first and second free voter respectively.*

- **Phase 2.** *The adversary votes with some or all of the voters it controls.*
- **Phase 3.** *\mathcal{A} and \mathcal{C} run the Tallying phase of the protocol, and the election result is published.*
- **Guess.** *The adversary outputs a guess μ' of the coin flip that randomly assigned the voting preferences of the two free voters.*

\mathcal{A} wins if $\mu' = \mu$.

Definition 3 (Vote Indistinguishability). *An E-Voting Protocol with security parameter θ is VI-secure if, for every probabilistic polynomial-time adversary \mathcal{A} that outputs a guess μ' of the coin flip μ (as described in the security game of Definition 2), there exists a negligible function η such that $\mathbb{P}[\mu' = \mu] \leq \frac{1}{2} + \eta(\theta)$.*

In the following theorem we prove our voting protocol VI-secure under the DDH assumption in the security game defined above.

Theorem 1. *In the Random Oracle Model (ROM), if the DDH assumption holds, then the protocol of Section 3.1 is VI-secure, as per Definition 3.*

Proof. Suppose there exists a polynomial time adversary \mathcal{A} , that can attack the scheme with advantage ϵ . We claim that a simulator \mathcal{S} can be built to play the decisional DH game with advantage $\frac{\epsilon}{2}$. The simulator controls the random oracle that defines the hash function H , and starts by taking in a DDH challenge:

$$(g, A = g^a, B = g^b, \Xi),$$

with $\Xi = g^{ab}$ or $\Xi = R = g^{\xi}$.

First we consider the case in which the adversary controls \mathcal{A}_0 , the simulation proceeds as follows.

- **Init.** The adversary chooses the $N - 2$ voters to control. Without loss of generality we assume that the two free voters are associated to v_1 and v_2 .
- **Setup.** \mathcal{S} chooses uniformly at random in \mathbb{Z}_p^* the values \tilde{x}_i , $\tilde{\alpha}_\ell$, $\tilde{y}_{i,\ell}$, and $\tilde{z}_{i,\ell}$ for all $i \in [2]$, $\ell \in [M]$, and implicitly sets for all $i \in [2]$, $\ell \in [M]$:

$$\begin{aligned} x_i'' &= \tilde{x}_i + (-1)^i b, & \alpha'_\ell &= a \cdot \tilde{\alpha}_\ell, \\ y'_{i,\ell} &= a \cdot \tilde{y}_{i,\ell}, & z'_{i,\ell} &= a \cdot \tilde{z}_{i,\ell}. \end{aligned}$$

\mathcal{S} chooses the other values for authorities \mathcal{A}_1 and \mathcal{A}_2 following the protocol.

In the improbable case that $a = 0$, the DDH problem is trivially solvable ($g^a = g^{ab} = 1$). If $a \neq 0$,

since a and b come from an uniform distribution, then also these implicit values are uniformly distributed, so the choices of the simulator are indistinguishable from a real protocol execution.

Note that \mathcal{S} can compute all the values $g^{x_i''}$, g^{α_ℓ} , $g^{y_{i,\ell}'}$, $g^{z_{i,\ell}'}$, either normally (when the parameter has been explicitly chosen) or as follows:

$$\begin{aligned} g^{x_i''} &= g^{\tilde{x}_i} \cdot B^{(-1)^i}, & g^{\alpha_\ell} &= A^{\tilde{\alpha}_\ell}, \\ g^{y_{i,\ell}'} &= A^{\tilde{y}_{i,\ell}'}, & g^{z_{i,\ell}'} &= A^{\tilde{z}_{i,\ell}'}. \end{aligned}$$

for all $i \in [2]$, $\ell \in [M]$. Therefore, \mathcal{S} can simulate the setup phase, exploiting the RO to simulate the NIZKPs for x_i'' , α_ℓ , $y_{i,\ell}'$, $z_{i,\ell}'$ for $i \in [2]$, $\ell \in [M]$.

- **Registrar Phase.** For the voters associated to v_i with $3 \leq i \leq N$, \mathcal{S} can simulate this phase following the protocol normally (since all relevant parameters have been explicitly chosen), while for $i \in [2]$ \mathcal{S} does the following:

1. \mathcal{A} computes the initial step of the ballot $\bar{b}_{0,i}$ on behalf of \mathcal{A}_0 and proves its correctness with the appropriate DVNIZKPs. By rewinding \mathcal{A} and exploiting the control of the random oracle, \mathcal{S} is able to extract from the DVNIZKPs the values of k , λ , and $\tilde{z}_{i,\ell}$ for all $\ell \in [M]$ (see (Spadafora et al., 2021)). Moreover, since \mathcal{A}_0 communicates the set of indexes of valid tokens V_i to the voter associated to v_i (that is controlled by the simulator), \mathcal{S} can reconstruct the values of the $\sigma_{i,\ell}$ for all $\ell \in [M]$.
2. \mathcal{S} computes step 1 of the ballot $\bar{b}_{1,i} = (\bar{b}_{1,i,\ell})_{\ell \in [M]}$ as:

$$\begin{aligned} \bar{b}_{1,i,\ell} &= A^{\tilde{z}_{i,\ell} \tilde{z}_{i,\ell} (\sigma_{i,\ell} + x_i' + \tilde{x}_i)} \cdot \Xi^{\tilde{z}_{i,\ell} \tilde{z}_{i,\ell} (-1)^i} \\ &\stackrel{*}{=} g^{z_{i,\ell} (\sigma_{i,\ell} + x_i)} \quad \forall \ell \in [M] \end{aligned} \quad (1)$$

where $\stackrel{*}{=}$ of Equation (1) holds iff $\Xi = g^{ab}$ in the DDH challenge. Since it controls the voter associated to v_i , \mathcal{S} can forge the DVNIZKPs exploiting the value e_i . In order to hide from \mathcal{A} which tokens are valid, these DVNIZKPs are forged using random values.

3. \mathcal{S} can perform step 2 on behalf of \mathcal{A}_2 normally, then \mathcal{A} computes step 3 on behalf of \mathcal{A}_0 and proves its correctness.
4. Finally \mathcal{S} computes the final ballot $b_i = (b_{i,\ell})_{\ell \in [M]}$ as:

$$\begin{aligned} b_{i,\ell} &= A^{\tilde{y}_{i,\ell} y_{i,\ell}' (\sigma_{i,\ell} + x_i' + \tilde{x}_i)} \cdot \Xi^{\tilde{y}_{i,\ell} y_{i,\ell}' (-1)^i} \\ &\stackrel{*}{=} g^{y_{i,\ell} (\sigma_{i,\ell} + x_i)} \end{aligned} \quad (2)$$

where again $\stackrel{*}{=}$ of Equation (2) holds if and only if $\Xi = g^{ab}$ in the DDH challenge.

- **Voting:** Phases 1, 2, and the Challenge are performed as in Definition 2.

- **Tallying.** Without loss of generality, suppose that only the v_i with $i \in [T]$ have voted. For $\ell \in [M]$, \mathcal{S} carries on with the simulation as follows:

1. \mathcal{S} computes the preliminary and final votes on behalf of \mathcal{A}_1 and \mathcal{A}_2 following the protocol without problems. In fact, for $i \in [2]$, we have:

$$\frac{\alpha_\ell'}{y_{i,\tilde{\phi}_i(\ell)}'} = \frac{a \tilde{\alpha}_\ell}{a \tilde{y}_{i,\tilde{\phi}_i(\ell)'}} = \frac{\tilde{\alpha}_\ell}{\tilde{y}_{i,\tilde{\phi}_i(\ell)'}} \quad \forall \ell \in [M],$$

and these values are known to \mathcal{S} .

2. \mathcal{S} computes and publishes the values $g^{\alpha_\ell} = A^{\tilde{\alpha}_\ell \alpha_\ell'} \quad \forall \ell \in [M]$, and simulates the proofs of correctness.
3. Finally note that \mathcal{S} can compute:

$$\sum_{i=1}^T x_i'' = \tilde{x}_1 - b + \tilde{x}_2 + b + \sum_{i=3}^T x_i'' = \tilde{x}_1 + \tilde{x}_2 + \sum_{i=3}^T x_i''$$

so for the rest of the tallying phase \mathcal{S} can follow the protocol.

- **Guess** Eventually \mathcal{A} will output a guess μ' of the coin flip performed by \mathcal{S} during the Challenge. \mathcal{S} then outputs 0 to guess that $\Xi = g^{ab}$ if $\mu' = \mu$, otherwise it outputs 1 to indicate that Ξ is a random group element $R \in \mathbb{G}$.

The case in which the adversary controls \mathcal{A}_1 and the case in which the adversary controls \mathcal{A}_2 , proceed similarly. If \mathcal{A}_1 is corrupted, the main difference is that \mathcal{S} implicitly sets:

$$\alpha_\ell'' = a \cdot \tilde{\alpha}_\ell, \quad y_{i,\ell}'' = a \cdot \tilde{y}_{i,\ell}', \quad \tilde{z}_{i,\ell} = a \cdot \tilde{y}_{i,\ell}',$$

while α_ℓ' , $y_{i,\ell}'$, $z_{i,\ell}'$ are chosen normally.

If \mathcal{A}_2 is corrupted, the main difference is that \mathcal{S} implicitly sets:

$$x_i' = \tilde{x}_i + (-1)^i b,$$

while x_i'' is chosen normally.

Essentially, in all three cases when Ξ is not random the simulator \mathcal{S} gives a perfect simulation. This means that the advantage is preserved, so it holds that:

$$\mathbb{P}[\mathcal{S}(g, A, B, \Xi = g^{ab}) = 0] = \frac{1}{2} + \epsilon.$$

On the contrary, when Ξ is a random element $R \in \mathbb{G}$, every token and vote belonging to the free voters becomes independent from the values that would have been computed by following the protocol (since they are simulated using the random value R), so \mathcal{A} can gain no information about the votes from them, while the tally is always correct. Since the security game is structured in such a way that the tally and the tokens

of the other voters (i.e. the values where Ξ is not used in the computation by \mathcal{S}) do not give any information about the coin flip μ , we have that:

$$\mathbb{P}[\mathcal{S}(g, A, B, \Xi = R) = 0] = \frac{1}{2}.$$

Therefore, \mathcal{S} can play the DDH game with non-negligible advantage $\frac{\epsilon}{2}$. \square

5.2 General Properties of the Protocol

The general properties of a vote system introduced in Section 2.4, can all be proved for the protocol described in Section 3.1. Here we prove only *Vote-Coercion Resistance*, we refer to (Longo and Spadafora, 2021) for the other proofs.

Proposition 1 (Vote-Coercion Resistance). *In the ROM, if the DDH assumption holds, then the protocol is vote-coercion resistant, as per Definition 1.*

Proof. In order to comply with the coercer’s request, a voter associated to $v_i \in V_c$ has to assign the valid tokens to $(C_{i,1}, \dots, C_{i,p})$. Since the Registrar Phase is performed in a protected environment, only the voter associated to v_i knows which tokens are valid, and cannot give a meaningful proof of this fact to \mathcal{A} as discussed at the end of the registrar phase (Section 3.1.2).

Thanks to Theorem 1, in the ROM, if the DDH assumption holds, then the protocol has vote-indistinguishability and the only way to determine if a vote expresses a specific choice is to distinguish valid and decoy tokens. Since \mathcal{A} cannot do so, all the information that can be gained from the votes is given by the final tally. This means exactly that the probability of \mathcal{A} detecting that a voter in V_c has not followed its instruction is the same in Ψ_1 and Ψ_2 . \square

6 CONCLUSIONS

In this paper we have generalized the *two-candidates-one-preference* e-voting protocol of (Spadafora et al., 2021) into an *M-candidates-P-preferences* protocol. We have tweaked the system of ZKPs that ensure transparency and full auditability of the process by using non-interactive proofs to enhance efficiency, exploiting designated-verifier proofs to preserve plausible deniability against coercers. Moreover, we have abandoned the blockchain infrastructure in favor of a more traditional bulletin board.

Compared with the two-candidates protocol, our generalization introduces an additional authority, that is required in order to properly mask the multiple *valid* and *decoy* tokens in each ballot, so that the system remains secure even if one authority is corrupt.

Note that the authorities can perform the setup phase asynchronously, and possible DOS attacks may be mitigated with a long-lasting Registrar phase. We can also adopt the strategy of dividing the authorities in independent triplets that manage restricted pools of voters (much like how large-scale elections are divided in voting districts). This approach limits the damage in case that more than one authority is corrupted, speeds up the final step of tallying (whose computational cost is linear in the number of votes managed by a triplet of authorities), and enhances the overall efficiency by distributing the workload.

Security. The protocol fulfills all the security properties required for an e-voting protocol to be considered secure, proven in the random oracle model under the classical Decisional Diffie-Hellman Assumption.

Regarding coercion resistance, the differences between definitions are subtle. In its strongest form, coercion resistance includes protection against forced abstention attacks and randomized voting. Randomized vote attacks are less effective in swaying an election result with respect to other coercion attacks, while forced abstention may be more effective, but it would require more effort, since more voters have to be controlled in order to achieve an impacting result. In fact, in our protocol the attacker should identify every coerced voter by requesting a signature, in order to link the voter’s identity with a public key and its ballot, as published in the BB.

Although our definition of coercion resistance seems weaker, we remark that the most prominent e-voting protocols with stronger defence against coercion assume that there is a moment during the voting phase when the voter is not under control of the attacker. The Amun protocol, instead, protects the voter even if during the voting period there is constant surveillance from the coercer. Therefore, this may be preferable when the voting period is limited, since, in this scenario, it is more likely for the attacker to maintain continuous control.

To have any kind of anti-coercion resistance is essential that there is a moment where the voter receives some private information that can then be concealed from the coercer with plausible deniability. In the description of the protocol we have assumed that the communication between the voter and the authorities during the registrar phase happens in a safe and controlled environment, where the coercer has no power. This requirement is equivalent to exchanging information through untappable channels. This is a common assumption in coercion-resistant protocols (Juels et al., 2010; Clarkson et al., 2008).

In the scenario where the voting period is limited,

some voters may struggle to access a secure physical voting booth, potentially leading to disenfranchisement. However, the longer duration of the registration phase provides more opportunities for voters to reach a secure registration booth.

The authors of (Longo et al., 2022; Bitussi et al., 2023) propose a method to protect voters from coercion by exploiting surveillance gaps. They assume that an adversary cannot maintain constant surveillance over a voter, allowing the voter to act freely during these gaps. In particular, when a voter registers, the voting credential is not issued immediately, but after a random delay. A DVNIZKP is sent after another random wait to prove the credential's correctness. These random waiting periods allow a coerced voter to exploit surveillance gaps to forge a credential and its DVNIZKP. This approach can also be applied to our protocol, where the credential can be seen as the set of indexes of the valid v -tokens.

Final Remarks. Many election systems allow voters to cast a blank ballot or to leave some of the P possible preferences unexpressed. This feature can be easily added to the protocol presented here by simply adding P dummy candidates that represent blank choices.

ACKNOWLEDGMENTS

This work has been partially supported by project SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU. This work has been partially supported by a joint laboratory between Fondazione Bruno Kessler and the Italian Government Printing Office and Mint, Italy. The authors are members of the INdAM Research Group GNSAGA.

REFERENCES

Bitussi, M., Longo, R., Marino, F. A., Morelli, U., Sharif, A., Spadafora, C., and Tomasi, A. (2023). Coercion-resistant i-voting with short pin and oauth 2.0. In *E-Vote-ID 2023*, LNI. Herausgeber et al. Bonn. to appear.

Chaidos, P., Cortier, V., Fuchsbaue, G., and Galindo, D. (2016). Beleniosrf: A non-interactive receipt-free electronic voting scheme. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*.

Clarkson, M. R., Chong, S., and Myers, A. C. (2008). Civitas: Toward a secure voting system. In *IEEE Symposium on Security and Privacy*.

Cortier, V., Gaudry, P., and Glondou, S. (2019). Belenios: a simple private and verifiable electronic voting system. In *Foundations of Security, Protocols, and Equational Reasoning*. Springer.

Fischer, J.-B. and Stern, J. (1996). An efficient pseudo-random generator provably as secure as syndrome decoding. In *Advances in Cryptology—EUROCRYPT'96: International Conference on the Theory and Application of Cryptographic Techniques Saragossa, Spain, May 12–16, 1996 Proceedings 15*. Springer.

Gardner, R. W., Garera, S., and Rubin, A. D. (2009). Coercion resistant end-to-end voting. In *International Conference on Financial Cryptography and Data Security*. Springer.

Haines, T. and Smyth, B. (2019). Surveying definitions of coercion resistance. *Cryptology ePrint Archive*.

Iovino, V. and Visconti, I. (2019). Non-interactive zero knowledge proofs in the random oracle model. In *Codes, Cryptology and Information Security: Third International Conference, C2SI 2019, Rabat, Morocco, April 22–24, 2019, Proceedings-In Honor of Said El Hajji*. Springer.

Jakobsson, M., Sako, K., and Impagliazzo, R. (2001). Designated verifier proofs and their applications. In *Advances in Cryptology—EUROCRYPT'96: International Conference on the Theory and Application of Cryptographic Techniques Saragossa, Spain, May 12–16, 1996 Proceedings*. Springer.

Juels, A., Catalano, D., and Jakobsson, M. (2010). Coercion-resistant electronic elections. In *Towards Trustworthy Elections*. Springer.

Krips, K. and Willemson, J. (2019). On practical aspects of coercion-resistant remote voting systems. In *International Joint Conference on Electronic Voting*. Springer.

Kusters, R., Truderung, T., and Vogt, A. (2010). A game-based definition of coercion-resistance and its applications. In *2010 23rd IEEE Computer Security Foundations Symposium*.

Longo, R., Morelli, U., Spadafora, C., and Tomasi, A. (2022). Adaptation of an i-voting scheme to italian elections for citizens abroad. *University of Tartu Press*.

Longo, R. and Spadafora, C. (2021). Amun: Securing e-voting against over-the-shoulder coercion. *Cryptology ePrint Archive*.

Roth, R. M. (2006). Introduction to coding theory. *IET Communications*, 47(18-19).

Ryan, P. Y., Rønne, P. B., and Iovino, V. (2016). Selene: Voting with transparent verifiability and coercion-mitigation. In *International Conference on Financial Cryptography and Data Security*. Springer.

Schnorr, C. (1991). Efficient signature generation by smart cards. *Journal of Cryptology*, 4:161–174.

Shoup, V. and Alwen, J. (2007). Σ -Protocols Continued and Introduction to Zero Knowledge.

Spadafora, C., Longo, R., and Sala, M. (2021). A coercion-resistant blockchain-based E-voting protocol with receipts (2021). In *Advances in Mathematics of Communications*. AIMS, doi:10.3934/amc.2021005.