

Fear Not, Vote Truthfully: Secure Multiparty Computation of Score Based Rules

Lihi Dery^{a,b,*}, Tamir Tassa^c, Avishay Yanai^d

^a*Department of Industrial Engineering and Management, Ariel University, Ariel, Israel*

^b*Ariel Cyber Innovation Center, Ariel, Israel*

^c*Department of Mathematics and Computer Science, The Open University, Raanana, Israel*

^d*VMware Research, Herzliya, Israel*

Abstract

We propose a secure voting protocol for score-based voting rules, where independent talliers perform the tallying procedure. The protocol outputs the winning candidate(s) while preserving the privacy of the voters and the secrecy of the ballots. It offers perfect secrecy, in the sense that apart from the desired output, all other information – the ballots, intermediate values, and the final scores received by each of the candidates – is not disclosed to any party, including the talliers. Such perfect secrecy may increase the voters' confidence and, consequently, encourage them to vote according to their true preferences. The protocol is extremely lightweight, and therefore it can be easily deployed in real life voting scenarios.

Keywords: Electronic Voting, Secure Multiparty Computation, Perfect Ballot Secrecy, Voting Protocols, Computational Social Choice

*Corresponding author

Email addresses: lihid@ariel.ac.il (Lihi Dery), tamirta@openu.ac.il (Tamir Tassa), yanaia@vmware.com (Avishay Yanai)

1. Introduction

Ballot secrecy is an essential goal in the design of voting systems. When voters are concerned for their privacy, they might decide to vote differently from their real preferences, or even abstain from voting altogether. Our main goal here is achieving perfect ballot secrecy. The usual meaning of privacy in the context of secure voting is that the voters remain anonymous. Namely, even though the ballots are known (as is the case when opening the ballot box at the end of an election day), no ballot can be traced back to the voter who cast it. We go one step further and consider perfect ballot secrecy, or full privacy [14], i.e., given any coalition of voters, the protocol does not reveal any information on the ballots, beyond what can be inferred from the published results.

The mere anonymity of the ballots might not provide sufficient privacy and hence may encourage untruthful voting, as our next two examples show. Consider a group of faculty members who need to jointly decide which applicant to accept to the faculty out of a given list of candidates. To that end, each faculty member (voter) anonymously casts a ballot. A tallier counts the ballots and uses some voting rule to determine the elected candidate(s). The problem with this voting strategy is that even though the tallier cannot link voters to ballots, he does see the actual ballots. Hence, besides the final outcome, say, that Alice is the elected candidate, the tallier is exposed to additional information which may be sensitive; e.g., that the candidate Bob received no votes, even though some of the voters declared upfront that they are going to vote for him. The imperfect privacy of such a voting system may cause some voters to vote untruthfully. The protocol that we present herein offers perfect privacy and, thus, may encourage voters to vote truthfully.

As another example, consider the London Inter-Bank Offered Rate (LIBOR)¹ which is the benchmark interest rate at which banks can borrow from each other. The rate is computed daily; banks that are benchmark submitters contribute to setting the LIBOR by means of voting: each bank's "vote" is an interest rate and the LIBOR is determined by some averaging over the submitted votes. The bank's submitted rate may signal the bank's financial viability. Worrying about the signal which their submitted rate conveys, some banks may submit an untruthful rate. To prevent this, the bank's individual submissions (the ballots) are kept private and are published only three months after the submission date. However, the tallier is exposed to these ballots and may be able to link some ballots to banks by financial

¹See ICE LIBOR <https://www.theice.com/iba/libor>

analysis. Therefore, even anonymous ballots might not provide sufficient privacy. Securing the ballots, as we suggest herein, means that there is less incentive to misrepresent one’s ballot, and thus there is less incentive for strategic voting.

Contributions. We present a secure protocol with perfect ballot secrecy to compute election results for score-based voting rules. This is achieved by employing cryptographic multiparty computation techniques. Score-based voting rules are rules where a voter’s ballot consists of scores given to each of the candidates, and the winner is the candidate with the highest aggregated score [8]. This family includes rules such as PLURALITY, RANGE, APPROVAL, VETO, and BORDA. We follow what is known in cryptography as “the mediated model” [2], in the sense that our protocol involves a set of talliers who perform the aggregation of ballots and compute the final voting results, but they are not allowed to access the actual ballots or other computational results such as the final scores of candidates. Our protocol is secure under the assumption that the talliers have an *honest majority*; that means that if a majority of the talliers are honest and do not collude, then no tallier can infer any information on the private ballots nor on the aggregated scores of candidates. Such perfect ballot privacy, by which the ballots and aggregated scores are not disclosed even to the talliers, may increase the voters’ confidence and, consequently, encourage them to vote according to their true preferences. As the protocol is compliant with all desired properties of secure voting systems, and is very efficient, it can be readily implemented in real life voting scenarios.

The paper is organized as follows. In Section 2 we review related work. In Section 3 we provide the necessary preliminaries on score-based voting rules and on secret sharing schemes. Our protocol is presented and discussed in Section 4. We analyze the computational and communication costs of the protocol in Section 5, discuss the protocol’s compliance with essential electronic voting requirements in Section 6, and conclude in Section 7.

2. Related work

The issue of enhancing democratic elections is widely studied in the AI community and specifically in the computational social choice community. Some recent studies look at securing attacks during the recounting of ballots [26], optimal attack problems in voting (e.g. by deleting voters [24]), electoral bribery problem [16], election control through social influence [20] and the complexity of multi-winner voting rules [48]. A central goal in this context is the design of systems in which the election results reflect properly and truthfully the will of the individuals in the underlying society. An important vehicle towards achieving that

goal is to secure the voting system so that it provides desired properties such as anonymity/privacy, fairness, robustness, uniqueness, and uncoercibility.

Previous studies on secure voting focused on different desired properties, e.g. privacy, or anonymity (a ballot cannot be connected to the voter who cast it), uniqueness (every voter can vote once), correctness (the issued winners are the ones that should be selected by the underlying voting rule from the cast ballots), and fairness (all voters must cast their ballot without seeing other votes or intermediate voting results). See e.g. Chang & Lee [12], Gritzalis [29], Zagórski et al. [50].

Our focus is on preserving privacy and achieving perfect ballot secrecy. One way to achieve those goals is by using methods that allow the voters to compute the outcome themselves without relying on a tallier to aggregate and count the votes, e.g. [6]. Another way is to use a third-party, a.k.a *a tallier*. In order to secure the transition of the votes which are sent from the voters to the tallier, various cryptographic techniques were utilized in prior art.

Early studies used the notions of mix-nets and anonymous channels [14, 42, 44]. Blind signatures [13] were used in other secure e-voting protocols, e.g. [28, 30]. Chen et al. [15] proposed a secure e-voting system based on the hardness of the discrete logarithm problem. Benaloh [5] proposed a practical scheme for conducting secret-ballot elections in which the outcome of an election is verifiable by all participants and even by non-participating observers; his scheme is based on secret sharing homomorphisms [4] that allow computations on shared data.

A large number of studies utilized homomorphic encryption, as it enables voting aggregation in the ciphertext domain. For example, Cramer et al. [21] proposed a scheme in which each voter posts a single encrypted ballot; owing to the homomorphism of the cipher, the final tally is verifiable to any observer of the election. Damgård et al. [22] proposed a generalization of Paillier's probabilistic public-key system [41] and then showed how it can be used for efficient e-voting. While most homomorphic e-voting schemes are based on additive homomorphism, Peng et al. [43] proposed a scheme based on multiplicative homomorphism. In their scheme, the tallier recovers the product of the votes, instead of their sum, and then the product is factorized to recover the votes.

To the best of our knowledge, only two previous studies considered the question of private execution of the computation that the underlying voting rule dictates. Canard et al. [10] considered the Majority Judgment (MJ) voting rule [3], which does not fall under the score-based family of rules that we consider here. They first translate the complex control flow and branching instructions that the MJ rule entails into a branchless algorithm; then they devise a privacy-

preserving implementation of it using homomorphic encryption, distributed decryption schemes, distributed evaluation of Boolean gates, and distributed comparisons. Nair et al. [38] suggest to use secret sharing for the tallying process in Plurality voting. Their protocol provides anonymity but does not provide perfect secrecy as it reveals the final aggregated score of each candidate. In addition, their protocol is vulnerable to cheating attacks, as it does not include means for detecting illegal votes. In our study, which covers all score-based rules, we provide perfect privacy as well as means for preventing cheating by using a secret sharing-based secure multiparty computation (see Section 4.1).

3. Preliminaries

This section provides the required background on score based voting rules (Section 3.1) and secret sharing (Section 3.2).

3.1. Score-based voting rules

We consider a setting in which there are N voters, $\mathbf{V} = \{V_1, \dots, V_N\}$, that need to hold an election over M candidates, $\mathbf{C} = \{C_1, \dots, C_M\}$. The election determines a score $\mathbf{w}(m)$ for each candidate C_m , $m \in [M] := \{1, \dots, M\}$ in a manner that will be discussed below. Let $K \in [M]$ be some fixed integral parameter. Then the output of the voting algorithm is the subset of the K candidates with the highest \mathbf{w} -scores, where ties are broken either arbitrarily or by another rule that is agreed upfront. ($K = 1$ corresponds to the typical case of a single winner.) Our protocol can be easily extended to output also the ranking of the candidates or the final scores they received. As such extensions are straightforward, we focus here on the “lean” output consisting only of the identities of the K elected candidates.

In score-based voting rules, every voter V_n , $n \in [N] = \{1, \dots, N\}$, creates a ballot vector of the form $\mathbf{w}_n := (\mathbf{w}_n(1), \dots, \mathbf{w}_n(M))$, where all single votes, $\mathbf{w}_n(m)$, are nonnegative and uniformly bounded. Define

$$\mathbf{w} = (\mathbf{w}(1), \dots, \mathbf{w}(M)) := \sum_{n=1}^N \mathbf{w}_n. \quad (1)$$

Then $\mathbf{w}(m)$ is the aggregated score of the candidate C_m , $m \in [M]$.

We consider five types of voter inputs to be used in the above described rule template, which give rise to five well known voting rules [40]:

- **PLURALITY.** $\mathbf{w}_n \in \{\mathbf{e}_1, \dots, \mathbf{e}_M\}$ where $\mathbf{e}_m, m \in [M]$, is an M -dimensional binary vector of which the m th entry equals 1 and all other entries are 0. Namely, V_n casts a vote of 1 for exactly one candidate and a vote of 0 for all others, and the winner is the candidate who was the favorite of the maximal number of voters.

- **RANGE.** $\mathbf{w}_n \in \{0, 1, \dots, L\}^M$ for some publicly known L .² Here every voter gets to give a score, ranging from 0 to L , to each candidate.

- **APPROVAL.** $\mathbf{w}_n \in \{0, 1\}^M$. Every voter submits a binary vector in which (up to) K entries are 1, while the remaining entries are 0. Such a voting rule is used when it is needed to fill K equivalent positions; for example, if K members in the senate of a university retired, it is needed to select K new senate representatives from the faculty.

- **VETO.** $\mathbf{w}_n \in \{\hat{\mathbf{e}}_1, \dots, \hat{\mathbf{e}}_M\}$ where $\hat{\mathbf{e}}_m, m \in [M]$, is an M -dimensional binary vector of which the m th entry equals 0 and all other entries are 1. In this method every voter states his least preferred candidate. The winner is the candidate that got the minimal number of zero votes.

- **BORDA.** $\mathbf{w}_n \in \{(\pi(0), \dots, \pi(M-1)) : \pi \in \Pi_M\}$, where Π_M is the set of all permutations over the set $\{0, \dots, M-1\}$. Here, the input of each voter is his own ordering of the candidates, i.e., $\mathbf{w}_n(m)$ indicates the position of C_m in V_n 's order, where a position of 0 (resp. $M-1$) is reserved to V_n 's least (resp. most) favorite candidate.

3.2. Secret sharing

Secret sharing methods [45] enable distributing a secret among a group of participants. Each participant is given a random share of the secret so that: (a) the secret can be reconstructed only by combining the shares given to specific *authorized* subsets of participants, and (b) combinations of shares belonging to unauthorized subsets of participants reveal zero information on the underlying secret.

The notion of secret sharing was introduced, independently, by Shamir [45] and Blakley [7], for the case of threshold secret sharing. Assuming that there are D participants, $\mathbf{P} = \{P_1, \dots, P_D\}$, then the access structure in Shamir's and in Blakley's schemes consists of all subsets of \mathbf{P} of size at least D' , for some $D' \leq D$. Such secret sharing schemes are called D' -out-of- D .

²We note that RANGE is not commonly included in the family of score-based voting rules, but we include it in this family since it fits the same voting rule "template". RANGE is common in many applications, e.g. www.netflix.com and www.amazon.com, and it is often used in recommender systems [35].

Shamir's secret sharing scheme works as follows. Assume that p is a sufficiently large prime so that the domain of all possible secrets may be embedded in the finite field \mathbf{Z}_p . Denote the secret to be shared by x . The Shamir's scheme has the following two procedures: **Share** and **Reconstruct**:

- **Share** $_{D',D}(x)$. The procedure samples a uniformly random polynomial $g(\cdot)$ over \mathbf{Z}_p , of degree $D' - 1$, where the free coefficient is x . That is, $g(t) = x + \alpha_1 t + \alpha_2 t^2 + \dots + \alpha_{D'-1} t^{D'-1}$, where α_j , $1 \leq j \leq D' - 1$, are selected uniformly at random from \mathbf{Z}_p .³ The procedure outputs D values, x_1, \dots, x_D , where $x_d = g(d)$ is the share given to P_d , $d \in [D] = \{1, \dots, D\}$.

- **Reconstruct** $_{D'}(x_1, \dots, x_D)$. The procedure is given any selection of D' shares out of $\{x_1, \dots, x_D\}$, say $\{x_{j_1}, \dots, x_{j_{D'}}\}$ where $1 \leq j_1 < \dots < j_{D'} \leq D$, and it then interpolates a polynomial $g(\cdot)$ of degree at most $D' - 1$ such that $g(j_i) = x_{j_i}$ for all $i \in [D']$. The procedure then outputs $x = g(0)$. It is easy to see that any subset of $D' - 1$ (or less) shares reveals nothing about the secret x , whereas any subset of D' (or more) shares fully determines the polynomial g , and in particular, the secret $x = g(0)$.

We conclude this crash course on secret sharing with the observation that the secret sharing procedure is linear in the following sense. Let x and y be two secrets from \mathbf{Z}_p and $a, b \in \mathbf{Z}_p$ be two publicly known values. Assume that (x_1, \dots, x_D) and (y_1, \dots, y_D) are shares in a Shamir's D' -out-of- D secret sharing scheme in x and y , respectively. Then, as can be readily verified, $(ax_1 + by_1, \dots, ax_D + by_D)$ are shares in a Shamir's D' -out-of- D secret sharing scheme in $ax + by$. Indeed, if g_x and g_y are the share-generating polynomials of degree $D' - 1$ that were used to create the shares in x and y , respectively, then the set of shares $(ax_1 + by_1, \dots, ax_D + by_D)$ correspond to the share-generating polynomial $f := ag_x + bg_y$ which is a polynomial of degree $D' - 1$ for which $f(0) = ax + by$.

Our protocol involves a *distributed* third party, $\mathbf{T} = \{T_1, \dots, T_D\}$, called the tallier (\mathbf{T}) or talliers (T_d , $d \in [D]$). In the protocol, we use secret sharing for creating shares of the private ballots of the voters and distributing them among the D talliers. As those ballots are vectors (see Section 3.1), the secret sharing is carried out for each entry independently, so that each of the talliers receives a share vector in each ballot.

³Note that the actual degree of the polynomial could be less than $D' - 1$, if $\alpha_{D'-1} = 0$, but for simplicity we relate to such polynomials as having degree $D' - 1$.

4. The method: A secure protocol for score based rules

In this section we present our protocol. As indicated earlier, our protocol is mediated, in the sense that it assumes a set of talliers, $\mathbf{T} = \{T_1, \dots, T_D\}$, who assist in the computations, but are not allowed to learn any information on the private votes of the voters. The number of talliers, D , can be any integer $D > 1$. Higher values of D will imply higher computational and communication costs, but they will also imply greater security against coalitions of corrupted talliers.

A privacy-preserving implementation of score-based rules is described in Protocol 1. Before delving into that protocol, we make the following observation. For each of the five score-based rules, there is a known upper bound B on the entries of \mathbf{w} . $B = N$ in PLURALITY, APPROVAL, and VETO rules, $B = NM$ in BORDA, and $B = NL$ in RANGE. Let p be a fixed prime greater than B . Then all computations in Protocol 1 are carried out in the field \mathbf{Z}_p .

First, each voter V_n , $n \in [N]$, constructs his own ballot vector (Step 1), $\mathbf{w}_n \in (\mathbf{Z}_p)^M$. We assume that all voters know the index $m \in [M]$ of each candidate. For example, that index can be determined by the lexicographical ordering of the candidates according to their names.

In Step 2, V_n creates D random share vectors of his ballot vector, \mathbf{w}_n , using Shamir's secret sharing scheme with the threshold $D' = \lfloor (D+1)/2 \rfloor$. The reason for selecting this specific threshold will be clarified later on. The sharing is done on each entry of \mathbf{w}_n independently. Namely, for each $m \in [M]$, V_n generates a random polynomial $g_{n,m}$ of degree $D' - 1$ over \mathbf{Z}_p , where $g_{n,m}(0) = \mathbf{w}_n(m)$. Then, in Step 3, V_n sends the share vector $\mathbf{w}_{n,d} = (g_{n,1}(d), \dots, g_{n,M}(d))$ to T_d , for all $d \in [D]$.

After receiving the ballot shares from all voters, T_d computes the sum of all received share vectors $\hat{\mathbf{w}}_d = \sum_{n=1}^N \mathbf{w}_{n,d} \bmod p$ (Step 4). Each such vector $\hat{\mathbf{w}}_d$ on its own carries no information regarding the votes (since it is the sum of uniformly random and independent vectors). But in view of Eq. (1) and the linearity of the secret sharing operation (see Section 3.2), the set $\{\hat{\mathbf{w}}_d(m) : d \in [D]\}$ is a set of D shares in $\mathbf{w}(m)$ by a Shamir D' -out-of- D secret sharing scheme, for all $m \in [M]$, where \mathbf{w} is the aggregated vector of scores. In other words, for every $m \in [M]$, there exists a polynomial g_m of degree $D' - 1$ over \mathbf{Z}_p , where $g_m(0) = \mathbf{w}(m)$, and $g_m(d) = \hat{\mathbf{w}}_d(m)$ for all $d \in [D]$.

The heart of the protocol is in Step 5: here, the talliers engage in a secure multiparty computation (MPC) in order to find the indices of the K candidates with the highest aggregated scores. This is a non-trivial task since no one holds the vector \mathbf{w} . In Section 4.1 we explain the notion of MPC and describe the MPC

protocol that we use for the above described task. Once those indices are found, the voters proceed to find the identity of the K candidates behind those indices (Step 6).

Protocol 1 A basic protocol for secure score-based voting

Input: $\mathbf{w}_n, n \in [N]; K \in [M]$.

Output: The K candidates from \mathbf{C} with highest aggregated scores in $\mathbf{w} = \sum_{n=1}^N \mathbf{w}_n$.

- 1: Each voter $V_n, n \in [N]$, constructs his ballot vector \mathbf{w}_n according to the selected indexing and voting rule.
 - 2: Each voter $V_n, n \in [N]$, generates a random polynomial $g_{n,m}$ of degree $D' - 1$ over \mathbf{Z}_p , where $g_{n,m}(0) = \mathbf{w}_n(m), \forall m \in [M]$. Then, he creates the share vector $\mathbf{w}_{n,d} = (g_{n,1}(d), \dots, g_{n,M}(d))$.
 - 3: $V_n, \forall n \in [N]$, sends $\mathbf{w}_{n,d}$ to $T_d, \forall d \in [D]$.
 - 4: $T_d, \forall d \in [D]$, computes $\hat{\mathbf{w}}_d = \sum_{n=1}^N \mathbf{w}_{n,d} \pmod p$.
 - 5: T_1, \dots, T_D find the indices of the K candidates in \mathbf{C} with highest \mathbf{w} -scores and output them.
 - 6: The voters find the identities of the top K candidates.
-

4.1. Sorting shared vectors

The main challenge in Step 5 of Protocol 1 is to find the indices of the K largest entries in \mathbf{w} . Towards that end, the talliers can implement any sorting algorithm on \mathbf{w} until all K largest entries are found. However, the talliers must not reconstruct \mathbf{w} 's entries, nor even learn any piece of information about them. They must perform oblivious comparisons using only the shares that they hold in \mathbf{w} 's entries.

Assume that the two entries that need to be compared are $u = \mathbf{w}(m)$ and $v = \mathbf{w}(m')$ for some $m, m' \in [M]$. Each tallier $T_d, d \in [D]$, holds random shares $u_d, v_d \in \mathbf{Z}_p$ in u and v , respectively, in a Shamir's D' -out-of- D secret sharing scheme. The talliers wish to find whether $u < v$, but without revealing any information beyond that on u and v .

To privately verify such questions, we use a secure multiparty computation (MPC) protocol [49]. An MPC protocol allows T_1, \dots, T_D to compute any function f over private inputs x_1, \dots, x_D that they hold, so that at the end of the

protocol everyone learns $f(x_1, \dots, x_D)$ *but nothing else*.⁴ A common approach towards designing efficient MPC protocols is to represent the function f by an arithmetic circuit C such that for every set of inputs, x_1, \dots, x_D , the output of the circuit, $C(x_1, \dots, x_D)$, equals $f(x_1, \dots, x_D)$.

The circuit is composed of input and output gates such that each input gate is fed with a single secret value by one of the parties, and an output gate determines a single value that is revealed to all parties. Additionally, between the input and output gates there are multiple layers of arithmetic gates that connect them. An arithmetic gate can be either addition or multiplication. Each gate is given exactly two inputs, and it produces one output such that the output of a gate at layer ℓ can be given as input to multiple gates in layer $\ell + 1$. (All input gates constitute the first layer of the circuit.) Only the value that the output gate issues is revealed to the parties; all intermediate values that pass from one gate to another remain secret from everyone. Specifically, a secure protocol allows the parties to maintain the invariant that the actual value output from each gate is secret-shared, as described in Section 3.2. When reaching the output gate, each party broadcasts the corresponding share that it holds, so that everyone can reconstruct the output.

The computational and communication costs of computing such circuits depend mainly on the number of multiplication gates and on the number of layers in the circuit, as we proceed to explain. To compute a multiplication gate of two secrets, s_1 and s_2 , the parties have to interact; i.e., each party needs to send some information to the other parties. However, to compute a multiplication gate of one secret s and a public value c , the parties do not need to interact (such a gate requires only local computation). The same holds for an addition gate. Therefore, for the efficiency of secure computation, circuit designers are mostly concerned with the number of multiplication gates in the circuit and with the *depth* of the circuit, i.e., the number of interactions that have to be performed sequentially (since they depend on each other and cannot be performed in parallel).

Specifically, in this work we use a design of an arithmetic circuit by Nishide & Ohta [39], which performs an MPC comparison of two secret values $u, v \in \mathbf{Z}_p$. It is assumed that each of the interacting parties, T_d , $d \in [D]$, holds shares u_d and v_d in u and v , respectively, in a Shamir's D' -out-of- D secret sharing scheme, where $D' \leq \lfloor (D+1)/2 \rfloor$. The circuit outputs the bit that indicates whether $u < v$.

⁴Of course, some information may be inferred from the desired output $f(x_1, \dots, x_D)$, but this is inevitable and allowed. For example, if the desired output is a median of x_1, \dots, x_D , then at the completion of the protocol, every tallier whose input is smaller than the median can infer that there are at least $\frac{D}{2}$ talliers that hold greater values than his own.

The circuit has a *constant depth* (15 to be concrete). This is advantageous as the depth does not depend neither on the number of parties nor on the field size p , so changing those parameters does not have a significant effect on performance (see Table 1 in Section 5).

4.2. The protocol's security

Here we discuss the security of the whole protocol. An important goal of secure voting is to provide anonymity; namely, it should be impossible to connect a ballot to the voter who cast it. Protocol 1 achieves that goal since each cast ballot is distributed into random shares and then each share is sent to a different tallier. Each such share carries zero information on the underlying ballot. Even subsets of $D' - 1 = \lfloor (D - 1)/2 \rfloor$ shares reveal no information on the secret ballot. Our working assumption is that the set of talliers has an *honest majority*; that assumption means that even if some of the talliers are dishonest and try to collude in order to extract sensitive information on the ballots, the number of colluding (dishonest) talliers is smaller than the number of the honest talliers. Since each ballot is shared by a D' -out-of- D secret sharing scheme, with $D' = \lfloor (D + 1)/2 \rfloor$, then the number of talliers that have to collude in order to recover the private ballots is at least $D' \geq D/2$, and that scenario is impossible under the honest majority assumption. Hence, under that assumption the talliers cannot recover the ballots, nor can they infer even partial information on them.

As for the MPC computation that the talliers carry out in Step 5 of Protocol 1, its security is proven in [39]. By utilizing that protocol, the talliers may find the indices of the K winning candidates without learning any information beyond the order that the aggregated scores induce on the candidates⁵. Also the security of that computation (in similarity to the security of the secret sharing of the individual ballots) is guaranteed under the assumption of an honest majority.

Hence, to summarize, the voters' privacy is perfectly preserved by our protocol, unless at least $D' = \lfloor (D + 1)/2 \rfloor$ talliers betray the trust vested in them. For example, with $D = 3$ talliers, at least two talliers would need to collude in order to recover the ballots; similarly, if $D = 5$ at least three talliers would need to collude for that purpose. If such a collusion does not occur, as implied in settings with

⁵The circuit that we use in order to verify inequalities may be modified in order to hide intermediate comparison results and output only the K indices of candidates with highest scores. Such a version, which we do not describe herein, will output only the K winning candidates without disclosing their order.

an honest majority, the talliers will be able to compute the final election results without learning *anything* beyond those computed results.

A collusion scenario that threatens the security of our protocol is highly improbable, and its probability decreases as D increases. Ideally, the talliers would be parties that enjoy high level of trust within the organization or state in which the elections take place, and whose business is based on such trust. Betraying that trust may incur devastating consequences for the talliers. Hence, even if D is set to a low value such as $D = 5$, in which case at least 3 talliers need to collude in order to recover the personal ballots, the probability of such a breakdown of trust in any conceivable application scenario (with a proper selection of the talliers) would be negligible.

Another possible attack scenario is as follows: a voter V_j can eavesdrop on the communication link between another voter V_n and each of the talliers, and intercept the messages that V_n sends to the talliers (in Protocol 1's Step 3) in order to recover \mathbf{w}_n from them; additionally, V_j may replace V_n 's original messages that carry shares of \mathbf{w}_n with other messages (say, ones that carry shares of \mathbf{w}_j , or any other desired fake ballot). Such an attack can be easily thwarted by requiring each party (a voter or a tallier) to have a certified public key, encrypt each message that he sends out using the receiver's public key and then sign it using his own private key; also, when receiving messages, each party must first verify them using the public key of the sender and then send a suitable message of confirmation to the sender. Namely, each message that a voter V_n sends to a tallier T_d in Step 3 of Protocol 1 should be signed with V_n 's private key and then encrypted by T_d 's public key; and T_d must acknowledge its receipt and verification.

In view of the above discussion, the tradeoff in setting the number of talliers D is clear: higher values of D provide higher security since more talliers would need to be corrupted in order to breach the system's security. However, increasing D has its costs: more independent and reputable talliers are needed, and the communication and computational costs of our protocol increase, albeit modestly (see Section 5).

A fundamental assumption in all secure voting systems that rely on fully trusted talliers (that is, talliers who receive the actual ballots from the voters) is that the talliers do not misuse the ballot information and that they keep it secret. In contrast, our protocol significantly reduces the trust vested in the talliers as it denies the talliers access to the actual ballots. Even in scenarios where some (a minority) of the talliers betray that trust, privacy is ensured. Such a reduction of trust in the talliers is essential in order to increase the confidence of the voters in the voting system so that they would be further motivated to exercise their right to

vote and moreover, vote according to their true preferences, without fearing that their private vote would be disclosed to anyone.

4.3. Validating the legality of the cast ballots

Protocol 1 is designed for honest voters, namely, voters who cast legal votes. However, voters may attempt cheating by submitting illegal ballots in order to help their candidate of choice. For example, assume that V_n 's favorite candidate is C_m and the voting rule is PLURALITY. Then a honest V_n would cast the ballot $\mathbf{w}_n = \mathbf{e}_m$ (see Section 3.1). A dishonest V_n , on the other hand, could cast the ballot $\mathbf{w}_n = N\mathbf{e}_m$. Such an illegal ballot would boost C_m 's chances of winning, or, if V_n is the only dishonest voter, it would even ensure C_m 's win. Similar options of cheating exist also with the other voting rules. Since the talliers do not see the actual ballots, if a voter can pull such a cheat, it might remain undetected.

In real-world voting scenarios, where voters typically cast their ballots on certified computers in voting centers, the chances of hacking such computers and tampering with the software that they run are small. However, for full-proof security, we proceed to describe an MPC solution that enables the talliers to validate the legality of each ballot, even though those ballots remain hidden from them. In case a ballot is found to be illegal, the talliers may recover it (by adding up all shares) and use the recovered ballot as a proof of the voter's dishonesty.

Let us start by examining the PLURALITY rule. A ballot \mathbf{w}_n is legal in this case iff $\mathbf{w}_n(m) \leq 1$, i.e. if $\mathbf{w}_n(m) \in \{0, 1\}$ for all $m \in [M]$, and $\sum_{m \in [M]} \mathbf{w}_n(m) = 1 \pmod p$ (assuming $p > M$). Each of the above m inequalities can be verified by an MPC sub-protocol that computes an arithmetic circuit that outputs the product $\mathbf{w}_n(m) \cdot (\mathbf{w}_n(m) - 1)$; a suitable MPC protocol that we may adopt in our context is described in Chida et al. [17]. The talliers accept the vote $\mathbf{w}_n(m)$ as legal (namely, being either 0 or 1) iff the result is 0. Finally, verifying that there exists exactly one entry in the vector \mathbf{w}_n with the value 1 is done by computing the sum $\sum_{m \in [M]} \mathbf{w}_n(m)$ and verifying that it equals 1. The fact that $N < p$ ensures that there will not be a wrap around.

To validate ballots \mathbf{w}_n in the case of VETO or APPROVAL, we also need to check that each entry, $\mathbf{w}_n(m)$, $m \in [M]$, is either 0 or 1, as described above for PLURALITY. If all entries were validated, the talliers need to proceed and check an aggregated condition on the ballot's entries. The aggregated condition in VETO can be checked by computing the sum $\sum_{m \in [M]} \mathbf{w}_n(m)$ and verifying that it equals $M - 1$. The aggregated condition in APPROVAL requires that the sum $S_n := \sum_{m \in [M]} \mathbf{w}_n(m)$ is at most K . That condition can be checked by a circuit that outputs $S_n \cdot (S_n - 1) \cdots (S_n - K)$. The talliers will accept \mathbf{w}_n as legal iff the

latter product equals 0. Indeed, if that product equals 0 then the talliers can deduce that V_n had voted for at most K candidates, without knowing the exact number of candidates for whom V_n had voted. Any result other than 0 will testify that V_n had cheated; in such cases the talliers can reject his vote and consider further consequences.

In the case of RANGE, a ballot \mathbf{w}_n is legal iff $\mathbf{w}_n(m) \leq L$ for all $m \in [M]$. Each of these m inequalities can be verified by a circuit that outputs $\mathbf{w}_n(m) \cdot (\mathbf{w}_n(m) - 1) \cdots (\mathbf{w}_n(m) - L)$. The talliers will accept the ballot \mathbf{w}_n iff the result of that circuit will be 0 for each $m \in [M]$.

The BORDA rule is slightly different. A ballot is legal under this rule if it consists of some permutation of the M values in $\{0, 1, \dots, M - 1\}$. Hence, a ballot is legal iff it satisfies the following two conditions:

$$\mathbf{w}_n(m) \leq M - 1 \quad \forall m \in [M]; \quad \text{and} \quad (2)$$

$$\mathbf{w}_n(m) \neq \mathbf{w}_n(m') \quad \forall m > m' \in [M]. \quad (3)$$

The entry-wise conditions in Eq. (2) can be verified as described earlier. The global condition in Eq. (3) can be verified by verifying that each of the $\binom{M}{2}$ differences $\mathbf{w}_n(m) - \mathbf{w}_n(m')$, $m > m' \in [M]$, is nonzero. However, for privacy reasons, the talliers must not recover those differences, since they would reveal the entire ballot. To prevent such leakage of information and still allow the verification, the talliers can generate $\binom{M}{2}$ random secret elements in the field \mathbf{Z}_p , denoted by $\rho_{m,m'}$, $m > m' \in [M]$, and then compute $\zeta_{m,m'} := \rho_{m,m'} \cdot (\mathbf{w}_n(m) - \mathbf{w}_n(m'))$. If the underlying field \mathbf{Z}_p is large, then with high probability (of $1/p$) the selected $\rho_{m,m'}$ is nonzero. In such cases, $\zeta_{m,m'} \neq 0$ iff $\mathbf{w}_n(m) - \mathbf{w}_n(m') \neq 0$, and $\zeta_{m,m'}$ reveals no information at all on $\mathbf{w}_n(m) - \mathbf{w}_n(m')$ (since $\rho_{m,m'}$ can be any nonzero element in the underlying field).

Generating shares of a random secret multiplier is a simple task. In fact, such a protocol is executed by the talliers anyway as part of the secure computation protocol of Chida et al. [17]. Specifically, for securely computing a multiplication gate in that protocol, the talliers generate two random secret values in the field. Therefore, in our computational cost analysis (Section 5), we upper bound the cost of generating shares in a random secret by the cost of evaluating a multiplication gate (where in fact the latter cost is strictly higher than the former).

Lastly, we consider the case of false negatives. In the field \mathbf{Z}_p there is a probability of $1/p$ that the resulting random element would be zero. In such a case, $\zeta_{m,m'}$ would be zero even though $\mathbf{w}_n(m) - \mathbf{w}_n(m') \neq 0$. Hence, the talliers would get

a false alarm regarding a ballot of some voter, as if it contains two equal entries, when in fact all of the ballot’s entries are distinct, as required. But the probability of such a false alarm is $1/p$. In Section 5 we present runtimes for the field \mathbf{Z}_p with $p = 2^{61} - 1$; in such large fields the probability $1/p$ is negligible. Furthermore, even if for some difference $\mathbf{w}_n(m) - \mathbf{w}_n(m')$ the talliers get an indication that it equals zero, they can repeat the test with another random and independent multiplier $\rho_{m,m'}$. If even that additional test yields $\zeta_{m,m'} = 0$, the talliers can withhold that ballot until its validity is verified (say, by performing additional independent tests until the probability of an error reduces to below some given threshold, or by revealing the value of the difference $\mathbf{w}_n(m) - \mathbf{w}_n(m')$).

5. Evaluation: Computational and communication costs

We analyze herein the computational and communication costs for the voters (Section 5.1), and for the talliers, where the latter discussion is separated to two parts — the cost for computing the final election results (Section 5.2) and the cost for validating the legality of the cast ballots (Section 5.3).

5.1. Costs for the voters

The costs for each voter are negligible, as a voter needs only to generate $M(D - 1) \log_2 p$ random bits, perform $M(D - 1)$ additions in \mathbf{Z}_p , and then send D messages of $M \log_2 p$ bits each (Protocol 1, Steps 2-3).

5.2. The talliers: the cost of computing the final election results

The cost in Step 4 of Protocol 1 is negligible ($(N - 1)M$ additions in \mathbf{Z}_p), but the determination of the winners (Step 5) is more costly as it invokes a protocol for secure comparison. The number of multiplication gates in the comparison circuit is $279 \cdot \log p + 5$ in a circuit of depth 15. A secure evaluation of a multiplication gate incurs a communication of $12 \log p$ bits per tallier (according to Chida et al. [17, Table 2]). Hence, the overall communication per tallier for a single comparison is roughly $3348 \cdot \log^2 p$ bits (or < 1.5 megabytes when $p = 2^{61} - 1$).

In order to evaluate the runtime of performing such a comparison, we ran it on Amazon AWS `m5.4xlarge` machines at N. Virginia over a network with bandwidth 9.6Gbps. We performed our evaluation with $D \in \{3, 5, 7, 9\}$ talliers. As for the bound B on aggregated scores, which affects the runtime, we examined two cases: $B < p_1$ and $B < p_2$, where $p_1 := 2^{13} - 1$ and $p_2 := 2^{31} - 1$ are two Mersenne primes. Namely, in cases where there are few voters, or ballots’ entries are small, and as a result $B < p_1$, we used the bound p_1 ; otherwise, we used

the bound p_2 , which seems to suffice for all conceivable application scenarios. Using Mersenne primes is advantageous in the context of secure computation, since multiplication of two field elements can be done without performing an expensive division. The results are presented in the first two rows of Table 1.

Note that those runtimes are for a single comparison. In order to determine the identity of the K winners, it is necessary to perform up to KM comparisons. We can see that even in large election scenarios that require choosing an underlying field of size p_2 , with $D = 9$ talliers (for enhanced security), with large numbers of candidates, M , and selected winners, K , the election results can be determined within only few seconds.

	$D = 3$	$D = 5$	$D = 7$	$D = 9$
$B < p_1 = 2^{13} - 1$	2.83	4.3	6.6	12.81
$B < p_2 = 2^{31} - 1$	9.07	9.54	9.64	15.0
Validating $5 \cdot 10^4$ PLURALITY ballot entries	41.3	42.2	52.9	65.55
Validating $5 \cdot 10^4$ RANGE ballot entries with $L = 20$	826	844	1058	1311
Validating $5 \cdot 10^4$ RANGE ballot entries with $L = 100$	4210	4945	5770	7050
Validating $5 \cdot 10^4$ BORDA ballots with $M = 5$	1652	1688	2116	2622
Validating $5 \cdot 10^4$ BORDA ballots with $M = 10$	7434	7596	9522	11799

Table 1: Rows 1-2 show runtimes (milliseconds) for a secure comparison protocol with a varying number of talliers, D , and two field sizes, p_1 and p_2 . Rows 3-7 show runtimes for a batch validation of $5 \cdot 10^4$ ballot entries or full ballots for various voting rules.

5.3. The talliers: the cost of validating ballots

For ballot validation, we consider three rules — PLURALITY, RANGE, and BORDA. (The validation of ballots in the two remaining rules, VETO and APPROVAL, is similar to that in PLURALITY.)

We extrapolated the runtime for executing the validation procedure for the various rules from the experimental results reported by Chida et al. [17]. They experimented on a similar network setting, but used a larger field with $p = 2^{61} - 1$. They experimented with a circuit that consists of one million multiplication gates that are evenly spread over $\{20, 100, 1000\}$ layers; hence, in each layer there are $\{5 \cdot 10^4, 10^4, 10^3\}$ multiplication gates, respectively. The reported runtimes as a function of D , the number of parties (talliers in our case), are shown in Table 2.

#layers	#multiplication gates per layer	$D = 3$	$D = 5$	$D = 7$	$D = 9$
20	50000	826	844	1058	1311
100	10000	842	989	1154	1410
1000	1000	1340	1704	1851	2243

Table 2: Runtimes (milliseconds) for computing 10^6 multiplication gates, spread evenly over 20, 100, and 1000 layers, as a function of the number D of talliers. The first two columns show the number of layers and the number of multiplication gates per layer in each setting.

5.3.1. The PLURALITY rule

As described in Section 4.3, a PLURALITY ballot validation depends on a circuit that checks whether $\mathbf{w}_n(m) \leq 1$. As explained there, such a circuit requires only one multiplication gate and has depth 1. In addition, to verify that there is only one ballot entry with the value 1, we need only to perform summation, which requires no multiplication gates. Overall, a validation circuit of q ballots requires exactly qM multiplication gates in *one layer*. As explained earlier, the runtimes in row 1 of Table 2 are for a circuit with 10^6 multiplication gates that were spread evenly over 20 layers. Hence, the runtime of executing $10^6/20 = 5 \cdot 10^4$ multiplication gates in a single layer is obtained by dividing those runtimes by 20, as shown in row 3 of Table 1. Those are the runtimes for Validating $5 \cdot 10^4$ PLURALITY ballot entries (from several different voters). In fact, those runtimes constitute an upper bound on the actual runtimes for our application, since they were obtained with a field size of $p = 2^{61} - 1$ that is larger than p_1 or p_2 that suffice for our needs. Those numbers indicate that validation is an extremely lightweight task in the case of the PLURALITY rule (as well as VETO and APPROVAL). Validating 50 million ballot entries can be done in roughly one minute.

5.3.2. The RANGE rule

Here we evaluate the cost of validating a ballot in the RANGE rule, as a function of the maximum score, L , that a voter can give to a candidate. The validation of an entry in a RANGE ballot requires a circuit with L multiplications in a row (when implemented naïvely, without optimizations). Thus, the experiment in Chida et al. [17] with a circuit that consists of a million multiplication gates spread over 20 layers (the runtimes of which are reported in row 1 of Table 2) captures exactly $5 \cdot 10^4$ RANGE ballots entries, in the case $L = 20$. We report those numbers in row 4 of Table 1. When L equals 100, we rely on the runtimes that are reported in Chida et al. [17] for a circuit of a million multiplication gates spread over 100

layers, as appear in row 2 of Table 2. In terms of validating RANGE ballot entries, those are the runtimes for validating 10000 RANGE ballot entries with $L = 100$. We multiply those runtimes by 5, in order to get the same batch size as in the case of $L = 20$, and report the resulting runtimes in row 5 row of Table 1.

We infer that even though validating RANGE ballot entries, especially for large values of L , is more costly than validating PLURALITY ballot entries, the validation task remains a very practical one. For example, validating 50 million RANGE ballot entries, even for $L = 100$ and $D = 9$, can be done in under two hours. (Recall that the election process usually spans a long period, say one day. Hence, the validation process can be spread along the entire election period, by validating each time the batch of ballots that were cast since the last validation.)

5.3.3. The BORDA rule

Finally, we turn to evaluate the cost of validating a ballot in the BORDA rule, as a function of the number of candidates, M . As described in Section 4.3, validating a BORDA ballot consists of two stages. In the first stage we check that each ballot entry is in the range $[0, M - 1]$. This can be done by a circuit of depth $M - 1$ with an overall number of $M(M - 1)$ multiplication gates. In the second stage we check that the ballot is a permutation of $\{0, \dots, M - 1\}$. This verification consists of generating $\binom{M}{2}$ secret random sharings and a secure computation of $\binom{M}{2}$ multiplication gates (see Section 4.3) in a circuit with a single layer. The computation and communication required for generating shares of a secret random value are strictly smaller than those required for a multiplication gate, since in the protocol of Chida et al. [17], evaluating a multiplication gate involves the generation of two secret random sharings. Thus, the cost of the second stage in the validation of a BORDA ballot can be bounded by the cost of performing $2 \cdot \binom{M}{2} = M(M - 1)$ multiplications, in two consecutive layers. In total, the overall cost of both stages in validating a single BORDA ballot can be bounded by performing $2M(M - 1)$ multiplications over $M - 1$ layers.⁶

When $M = 5$, the overall cost of validating a single ballot can be bounded, as stated above, by performing $2M(M - 1) = 40$ multiplications spread over 4 layers. Hence, to validate a batch of $2.5 \cdot 10^4$ BORDA ballots in this case we need to perform one million multiplications spread over 4 layers. Those runtimes, for

⁶Note that by taking a spread of the needed workload over more layers than what is actually needed, one gets higher runtimes. We perform that relaxation for the sake of simplicity, since we are only interested in upper bounds on the validation cost that will demonstrate its lightweight nature.

the various values of D , can be bounded by the runtimes in row 1 of Table 2, since the latter runtimes are for the same number of multiplications, only spread over more layers (20 instead of 4). Finally, we conclude that the overall runtime for validating a batch of $5 \cdot 10^4$ BORDA ballots is given by doubling the runtimes in row 1 of Table 2, as shown in row 6 of Table 1.

Bounding the runtimes for $M = 10$ is done in a similar fashion. The overall cost of validating a single ballot is smaller than that of performing $2M(M - 1) = 180$ multiplications spread over 9 layers. Hence, to validate a batch of $10^6/180$ BORDA ballots in this case we need to perform one million multiplications spread over 9 layers. Those runtimes, for the various values of D , can be bounded by the runtimes in row 1 of Table 2. Finally, we conclude that the overall runtime for validating a batch of $5 \cdot 10^4$ BORDA ballots is given by multiplying the runtimes in row 1 of Table 2 by 9, as shown in row 7 of Table 1.

In cases where $M \geq 22$, we can use the runtimes reported in rows 2 or 3 of Table 2, since in such cases the number of layers (in the first stage), which is $M - 1$, would be greater than 20. However, as the BORDA rule is impractical for such values of M , because it is hard to expect voters to have a full ranking over such a large number of candidates, we leave it for the interested reader to compute the resulting runtime bounds for such values of M . But the numbers for $M = 5$ and $M = 10$ indicate that the validation of BORDA ballots is a very practical task, despite the higher complexity of that task in comparison to that in PLURALITY or RANGE. For example, the validation of one million ballots when there are $M = 10$ candidates and $D = 9$ talliers would take less than 4 minutes.

6. Discussion

Electronic voting schemes should, ideally, comply with some essential requirements. Below we list those requirements, as defined in Chang & Lee [12]. We discuss, in a theoretical manner, the compliance of our proposed protocol with each of the requirements. (Such a discussion of compliance with essential requirements is common in papers dealing with secure electronic voting schemes, see e.g. Chung & Wu [18], Li et al. [33], Liaw [34], Wu et al. [47].)

Anonymity/Privacy. *The votes should remain anonymous throughout the entire process.* Our protocol achieves anonymity, as explained in Section 4.2. In particular, our protocol outputs only the desired election results, without revealing any information on the ballots.

Fairness. *A fair mechanism does not reveal intermediate results.* Our protocol is fair since all intermediate results (such as aggregated scores of candidates) are

kept secret-shared; the protocol is designed to output just the identities of the winner(s) and nothing beyond that.

Convenience. *No special equipment is required and the voters do not need to learn any specialized technique.* Our protocol relies on basic and general-purpose cryptographic functions that can be found in many free libraries. In fact, as our protocol includes a mechanism for validating the cast ballots (Section 4.3), voters may vote from anywhere (and not necessarily from voting centers). They only need to download a simple software package that implements our protocol vis-a-vis the talliers.

Robustness. *No malicious intruder should be able to interrupt the procedure.* This requirement is addressed by the standard security mechanisms that should be implemented on top of our protocol, as discussed in Section 4.2: signing ballots before sending, verifying received ballots, and confirming the receipt and verification of each ballot.

Mobility. *The mechanism can be implemented to run on the World Wide Web.* Our protocol allows mobility since it requires no special equipment and it relies on common cryptographic toolkit.

Uniqueness. *Each voter is allowed to vote only once.* If a voter attempts to vote twice (Step 3 in Protocol 1), only his first vote will be processed, while the second one will be ignored. Also, one voter cannot shoot down a message from another voter (in order to prevent the latter voter's ballot to reach the talliers) since the voter expects to receive from each tallier a confirmation of receiving his signed and authenticated message. In case such a confirmation is not received, the voter can resend his message.

Completeness. *Only eligible voters are allowed to vote.* This requirement is achieved by the usage of certificated public keys and signatures.

Uncoercibility. *A voter must not be able to prove to a third party how he had voted, in order to prevent bribery.* The only way to prove to a third party W how a voter V_n had voted is if all D talliers send to W the shares of V_n 's ballot \mathbf{w}_n . Hence, under our assumption that the tallier has an honest majority, our protocol offers uncoercibility.

Correctness. *Each ballot must be counted correctly.* Correctness follows directly from the talliers' semi-honesty (namely, that they follow the prescribed computations correctly). To ensure the talliers' semi-honesty in practical deployments, the talliers' software must be verified and authenticated, and it would be best to run it on a dedicated and tamper-proof machine (namely, a computer that is physically protected from any hacking attempts).

Efficiency. *The computational load of the whole process is required to be such that the result is obtained within a reasonable amount of time.* Our protocol complies with this requirement, as is evident from our cost analysis that shows that it is extremely lightweight (see Section 5).

The right to abstain. *A voter that wishes to abstain should be able to do so without revealing that fact to any other party [29].* In all considered rules, each ballot is a vector and the final result is determined by the sum of the private vectors. Hence, a voter V_n who wishes to abstain can use $\mathbf{w}_n = \mathbf{0}$. Such an action is equivalent to not participating. Even the talliers cannot tell that a message received from some voter contains an all-zero vote, because of the secret sharing. Hence, our protocol allows confidential abstinence.

In view of the above, our proposed scheme is suitable for electronic elections of any kind and scale, from small-scale elections similar to the examples given in Section 1, to national elections in populations of any size, provided that the underlying voting rule is score-based. In national elections, it is possible to implement our electronic voting system in a manner that resembles existing voting systems. Citizens will arrive at the voting centers and identify using their standard identification card. They will then obtain access to a computer into which they will privately enter their selection. The software running the protocol will translate their selection into a ballot vector, as described in Section 3.1; subsequently, shares of that ballot will be distributed to the D talliers' servers, who will then proceed to process them as described in Section 4. As opposed to existing tallying systems, in our system, the *final and true* results will be determined with utmost certainty and security within seconds after the voting period ends.

7. Conclusion

We begin this section with a birds-eye summary of our study. Then, we identify the limitations of our method and provide an outlook on future research directions.

7.1. Summary

We considered a setting in which a group of voters wishes to elect K candidates out of a given list of candidates. We considered score-based rules and showed how to securely compute the winners using a secure multiparty protocol. Our protocol offers perfect ballot secrecy: apart from the desired output (the identity of the elected K candidates), all other information, such as the actual ballots

or the aggregated scores that each of the candidates received, is kept secret from all parties – voters and talliers alike. (As indicated in Section 4.2, the protocol herein allows the talliers to deduce the ranking among the candidates, but a small modification of the underlying circuit may hide even that information.)

Such level of privacy may be essential in some scenarios. For example, when a prize committee needs to select K recipients out of M candidates, it is desirable to determine only the identity of the K prize recipients without revealing their internal ranking or their aggregated scores, as such pieces of information might expose undesired information about the ballots.

Such perfect secrecy may increase the confidence of the voters in the voting system, so that they would be encouraged to participate in the elections and vote truthfully without fearing that their private vote would be disclosed to anyone else. Furthermore, our technique can be used during iterative voting [1, 32, 36, 23]. If, during iterative voting, the candidates’ scores are kept secret from all parties, then voters would not obtain from the voting process information that could have been used for strategic voting. (Of course, the perfect secrecy of our protocol cannot stop strategic voting altogether, since voters may still base strategic voting on information from polls, rumors and other communication channels.)

The protocol complies with conventional security desiderata, as discussed in Section 6. An analysis of its computational and commutation costs shows that it is practical as it is extremely lightweight. The protocol is based solely on existing cryptographic arsenal. This is a prominent advantage of our protocol; indeed, protocols that can be implemented on top of existing libraries are advantageous over protocols that require the development, scrutiny and assimilation of new cryptographic components and, therefore, might be unattractive to practitioners.

7.2. Limitations and future work

Our method has several limitations, which suggest corresponding future research endeavours to resolve them.

- (a) **Convincing the public.** As with each new technology, our model faces the difficulty of “selling” it to the public and the legislators, and convincing them of its safety and other desired advantages. As a first step in this direction, a wide-range user study should be carried out.
- (b) **Extension to order-based rules.** Herein we focused on score-based rules. Another important family of rules consists of rules where the voter submits an ordered list of preferences. Such rules are called *order-based* or

pairwise-comparison voting rules (see Brandt & Sandholm [9]). Two prominent rules in this family are – COPELAND [19] and MAXIMIN (a.k.a KRAMER-SIMPSON) [31, 46]. We intend to devise a protocol for securely computing the winners in voting systems that are based on such rules.

- (c) **Extension to multi-winner elections.** Our current protocol can select the top K candidates, for any $K < M$; namely, it can output the K candidates that got the highest aggregated scores. However, there are multi-winner election protocols that are designed specifically for selecting the K candidates that would satisfy the voters the most [27, 25], in the sense that they also comply with additional social conditions (e.g., that the selected winners include a minimal number of representatives of specific gender, race, region etc.). This problem has unique features and it therefore requires its own secure protocol. Examples for voting rules that are designed for such a purpose are CHAMBERLIN-COURANT [11] and MONROE [37].
- (d) **A hierarchical tallier model.** We assumed a “flat” tallier model, where all talliers are operating on all ballots. However, in large voting systems, a hierarchical tallier may be more suitable. For example, in the US, it may be more suitable to use a hierarchy by county (first level), state (second level), and national (third and highest level). A modification of our protocol for such settings is in order.
- (e) **Restraining malicious talliers.** Our protocol assumes that the talliers are semi-honest, i.e., they follow the prescribed computations correctly. As discussed in Section 6, the semi-honesty of the talliers can be ensured in practice by securing the software and hardware of the talliers. However, one can provide even a stronger protection shield by cryptographic means. Namely, it is possible to design an MPC protocol that would be immune even to malicious talliers that may attempt to infer secret ballot information or intermediate voting results, or may even try to affect the final voting results. MPC protocols that are designed to be immune to malicious parties are usually significantly costlier and more complex than the corresponding MPC protocols for semi-honest parties. However, since in common voting systems, the final results are published usually hours and even days after voting had ended, the implementation of costlier MPC protocols that are designed to resist malicious talliers is not expected to be problematic, and it will enhance even further the security of the system and the trust of voters in the preservation of their privacy.

To conclude, the secure voting protocol that we presented here is of relevance to any setting in which voting should be implemented in a manner that ensures perfect privacy. The protocol promotes truthful voting and the final results are obtained within seconds. Even though the presented protocol is relevant and adequate as is to a wide range of practical voting scenarios, it opens up several interesting research directions that could widen its application scope and further strengthen its security.

Funding

Lihy Dery was supported by the Ariel Cyber Innovation Center in conjunction with the Israel National Cyber directorate of the Prime Ministers Office. The authors hereby declare that this support did not influence the work reported in this paper and that the funding source had no involvement in the project.

References

- [1] Airiau, S., & Endriss, U. (2009). Iterated majority voting. In *International Conference on Algorithmic Decision Theory* (pp. 38–49).
- [2] Alwen, J., Shelat, A., & Visconti, I. (2008). Collusion-free protocols in the mediated model. In *CRYPTO* (pp. 497–514).
- [3] Balinski, M., & Laraki, R. (2007). A theory of measuring, electing, and ranking. *Proceedings of the National Academy of Sciences*, 104, 8720–8725.
- [4] Benaloh, J. (1986). Secret sharing homomorphisms: Keeping shares of a secret secret. In *CRYPTO* (pp. 251–260).
- [5] Benaloh, J. (1987). *Verifiable Secret-Ballot Elections*. Technical Report MSR-TR-1987-1.
- [6] Benkaouz, Y., & Erradi, M. (2015). A distributed protocol for privacy preserving aggregation with non-permanent participants. *Computing*, 97, 893–912.
- [7] Blakley, G. (1979). Safeguarding cryptographic keys. In *International Workshop on Managing Requirements Knowledge* (pp. 48:313–317).
- [8] Brandt, F., Conitzer, V., Endriss, U., Lang, J., & Procaccia, A. D. (2016). *Handbook of computational social choice*. Cambridge University Press.

- [9] Brandt, F., & Sandholm, T. (2005). Decentralized voting with unconditional privacy. In *AAMAS* (pp. 357–364).
- [10] Canard, S., Pointcheval, D., Santos, Q., & Traoré, J. (2018). Practical strategy-resistant privacy-preserving elections. In *European Symposium on Research in Computer Security* (pp. 331–349).
- [11] Chamberlin, J. R., & Courant, P. N. (1983). Representative deliberations and representative decisions: Proportional representation and the borda rule. *The American Political Science Review*, (pp. 718–733).
- [12] Chang, C.-C., & Lee, J.-S. (2006). An anonymous voting mechanism based on the key exchange protocol. *Computers & Security*, 25, 307–314.
- [13] Chaum, D. (1982). Blind signatures for untraceable payments. In *CRYPTO* (pp. 199–203).
- [14] Chaum, D. (1988). Elections with unconditionally-secret ballots and disruption equivalent to breaking RSA. In *EUROCRYPT* (pp. 177–182).
- [15] Chen, C.-L., Chen, Y.-Y., Jan, J.-K., Chen, C.-C. et al. (2014). A secure anonymous e-voting system based on discrete logarithm problem. *Applied Mathematics & Information Sciences*, 8, 2571–2578.
- [16] Chen, L., Xu, L., Xu, S., Gao, Z., & Shi, W. (2019). Election with bribe-effect uncertainty: A dichotomy result. In *IJCAI* (pp. 158–164).
- [17] Chida, K., Genkin, D., Hamada, K., Ikarashi, D., Kikuchi, R., Lindell, Y., & Nof, A. (2018). Fast large-scale honest-majority MPC for malicious adversaries. In *CRYPTO* (pp. 34–64).
- [18] Chung, Y.-F., & Wu, Z.-Y. (2009). Approach to designing bribery-free and coercion-free electronic voting scheme. *Journal of Systems and Software*, 82, 2081–2090.
- [19] Copeland, A. H. (1951). A reasonable social welfare function. In *Mimeographed notes from a Seminar on Applications of Mathematics to the Social Sciences, University of Michigan*.
- [20] Cor, F., Cruciani, E., D’Angelo, G., & Ponziani, S. (2019). Exploiting social influence to control elections based on scoring rules. In *IJCAI* (pp. 201–207).

- [21] Cramer, R., Gennaro, R., & Schoenmakers, B. (1997). A secure and optimally efficient multi-authority election scheme. In *EUROCRYPT* (pp. 103–118).
- [22] Damgård, I., Jurik, M., & Nielsen, J. B. (2010). A generalization of Paillier’s public-key system with applications to electronic voting. *International Journal of Information Security*, 9, 371–385.
- [23] Dery, L., Obraztsova, S., Rabinovich, Z., & Kalech, M. (2019). Lie on the fly: Strategic voting in an iterative preference elicitation process. *Group Decision and Negotiation*, 28, 1077–1107.
- [24] Dey, P., Misra, N., Nath, S., & Shakya, G. (2019). A parameterized perspective on protecting elections. In *IJCAI* (pp. 238–244).
- [25] Elkind, E., Faliszewski, P., Laslier, J.-F., Skowron, P., Slinko, A., & Talmon, N. (2017). What do multiwinner voting rules do? an experiment over the two-dimensional euclidean domain. In *AAAI*.
- [26] Elkind, E., Gan, J., Obraztsova, S., Rabinovich, Z., & Voudouris, A. A. (2019). Protecting elections by recounting ballots. In *IJCAI* (pp. 259–265).
- [27] Faliszewski, P., Skowron, P., Slinko, A., & Talmon, N. (2017). Multiwinner voting: A new challenge for social choice theory. *Trends in computational social choice*, 74, 27–47.
- [28] Fujioka, A., Okamoto, T., & Ohta, K. (1992). A practical secret voting scheme for large scale elections. In *AUSCRYPT* (pp. 244–251).
- [29] Gritzalis, D. A. (2002). Principles and requirements for a secure e-voting system. *Computers & Security*, 21, 539–556.
- [30] Ibrahim, S., Kamat, M., Salleh, M., & Aziz, S. R. A. (2003). Secure e-voting with blind signature. In *NCTT* (pp. 193–197).
- [31] Kramer, G. H. (1977). A dynamical model of political equilibrium. *Journal of Economic Theory*, 16, 310–334.
- [32] Lev, O., & Rosenschein, J. S. (2016). Convergence of iterative scoring rules. *Journal of Artificial Intelligence Research*, 57, 573–591.

- [33] Li, C.-T., Hwang, M.-S., & Liu, C.-Y. (2008). An electronic voting protocol with deniable authentication for mobile ad hoc networks. *Computer Communications*, 31, 2534–2540.
- [34] Liaw, H.-T. (2004). A secure electronic voting protocol for general elections. *Computers & Security*, 23, 107–119.
- [35] Masthoff, J. (2011). Group recommender systems: Combining individual models. In *Recommender systems handbook* (pp. 677–702). Springer.
- [36] Meir, R., Polukarov, M., Rosenschein, J. S., & Jennings, N. R. (2010). Convergence to equilibria in plurality voting. In *AAAI*.
- [37] Monroe, B. L. (1995). Fully proportional representation. *American Political Science Review*, (pp. 925–940).
- [38] Nair, D. G., Binu, V. P., & Kumar, G. S. (2015). An improved e-voting scheme using secret sharing based secure multi-party computation. *CoRR*, abs/1502.07469.
- [39] Nishide, T., & Ohta, K. (2007). Multiparty computation for interval, equality, and comparison without bit-decomposition protocol. In *PKC* (pp. 343–360).
- [40] Nurmi, H. (2012). *Comparing voting systems* volume 3. Springer Science & Business Media.
- [41] Paillier, P. (1999). Public-key cryptosystems based on composite degree residuosity classes. In *EUROCRYPT* (pp. 223–238).
- [42] Park, C., Itoh, K., & Kurosawa, K. (1993). Efficient anonymous channel and all/nothing election scheme. In *EUROCRYPT* (pp. 248–259).
- [43] Peng, K., Aditya, R., Boyd, C., Dawson, E., & Lee, B. (2004). Multiplicative homomorphic e-voting. In *INDOCRYPT* (pp. 61–72).
- [44] Sako, K., & Kilian, J. (1995). Receipt-free mix-type voting scheme - A practical solution to the implementation of a voting booth. In *EUROCRYPT* (pp. 393–403).
- [45] Shamir, A. (1979). How to share a secret. *Commun. ACM*, 22, 612–613.

- [46] Simpson, P. B. (1969). On defining areas of voter choice: Professor tullock on stable voting. *The Quarterly Journal of Economics*, (pp. 478–490).
- [47] Wu, Z.-Y., Wu, J.-C., Lin, S.-C., & Wang, C. (2014). An electronic voting mechanism for fighting bribery and coercion. *Journal of Network and Computer Applications*, 40, 139–150.
- [48] Yang, Y. (2019). Complexity of manipulating and controlling approval-based multiwinner voting. In *IJCAI* (pp. 637–643).
- [49] Yao, A. (1982). Protocols for secure computation. In *FOCS* (pp. 160–164).
- [50] Zagórski, F., Carback, R. T., Chaum, D., Clark, J., Essex, A., & Vora, P. L. (2013). Remotegrity: Design and use of an end-to-end verifiable remote voting system. In *ACNS* (pp. 441–457).