# Kryvos: Publicly Tally-Hiding Verifiable E-Voting

Nicolas Huber nicolas.huber@sec.uni-stuttgart.de Institute of Information Security and Center for Integrated Quantum Science and Technology (IQ<sup>ST</sup>) University of Stuttgart, Germany

Julian Liedtke julian.liedtke@sec.uni-stuttgart.de Institute of Information Security University of Stuttgart, Germany

**Ralf Küsters** 

ralf.kuesters@sec.uni-stuttgart.de Institute of Information Security and Center for Integrated Quantum Science and Technology (IQ<sup>ST</sup>) University of Stuttgart, Germany

Johannes Müller johannes.mueller@uni.lu Interdisciplinary Centre for Security, Reliability and Trust University of Luxembourg

Pascal Reisert pascal.reisert@sec.uni-stuttgart.de Institute of Information Security University of Stuttgart, Germany

# ABSTRACT

Elections are an important corner stone of democratic processes. In addition to publishing the final result (e.g., the overall winner), elections typically publish the full tally consisting of all (aggregated) individual votes. This causes several issues, including loss of privacy for both voters and election candidates as well as so-called Italian attacks that allow for easily coercing voters.

Several e-voting systems have been proposed to address these issues by hiding (parts of) the tally. This property is called *tally*hiding. Existing tally-hiding e-voting systems in the literature aim at hiding (part of) the tally from everyone, including voting authorities, while at the same time offering verifiability, an important and standard feature of modern e-voting systems which allows voters and external observers to check that the published election result indeed corresponds to how voters actually voted. In contrast, real elections often follow a different common practice for hiding the tally: the voting authorities internally compute (and learn) the full tally but publish only the final result (e.g., the winner). This practice, which we coin publicly tally-hiding, indeed solves the aforementioned issues for the public, but currently has to sacrifice verifiability due to a lack of practical systems.

In this paper, we close this gap. We formalize the common notion of publicly tally-hiding and propose the first provably secure verifiable e-voting system, called Kryvos, which directly targets publicly tally-hiding elections. We instantiate our system for a wide

CCS '22, November 7-11, 2022, Los Angeles, CA, USA.

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9450-5/22/11...\$15.00

https://doi.org/10.1145/3548606.3560701

range of both simple and complex voting methods and various result functions. We provide an extensive evaluation which shows that Kryvos is practical and able to handle a large number of candidates, complex voting methods and result functions. Altogether, Kryvos shows that the concept of publicly tally-hiding offers a new trade-off between privacy and efficiency that is different from all previous tally-hiding systems and which allows for a radically new protocol design resulting in a practical e-voting system.

# **CCS CONCEPTS**

Andreas Vogt andreas.vogt@fhnw.ch

University of Applied Sciences and

Arts Northwestern Switzerland Windisch, Switzerland

• Applied computing  $\rightarrow$  Voting / election technologies; • Se**curity and privacy**  $\rightarrow$  *Human and societal aspects of security and* privacy; Cryptography.

### **KEYWORDS**

e-voting; publicly tally-hiding; verifiability; zk-snark, homomorphic commitments

#### **ACM Reference Format:**

Nicolas Huber, Ralf Küsters, Toomas Krips, Julian Liedtke, Johannes Müller, Daniel Rausch, Pascal Reisert, and Andreas Vogt. 2022. Kryvos: Publicly Tally-Hiding Verifiable E-Voting. In Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (CCS '22), November 7-11, 2022, Los Angeles, CA, USA. ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3548606.3560701

#### **INTRODUCTION** 1

Elections are an import corner stone of democratic processes. Besides national or local political elections, elections also often are and have to be carried out in companies, organizations, associations, etc. There exists a multitude of different voting methods ranging from relatively simple voting methods, such as plurality/singlechoice voting, to more complicated ones, such as cumulative voting with multiple votes, and very complex ones, like preferential elections and multi-round votings. In practice, not just simple voting

**Toomas Krips** toomas.krips@ut.ee University of Tartu Tartu, Estonia

Daniel Rausch daniel.rausch@sec.uni-stuttgart.de Institute of Information Security University of Stuttgart, Germany

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

methods are used but often also more complex ones as they allow for capturing the voters' preferences more accurately. For example, Instant-Runoff Voting (IRV) [79] is an important preferential voting method that is used in many countries for municipal or national political elections, including Australia, India, the UK, and the US [44, 70, 73, 85]. Also, there are many different result functions used in elections. For example, one might only be interested in the winner of the election (e.g., for presidential elections) or the *n* best or worst candidates (ranked or not ranked), e.g., to fill positions or decide who moves on to a runoff election.

Verifiability. A fundamental security property of elections is verifiability [2, 8, 11, 22, 24, 26, 56, 59, 60, 63], i.e., voters should be able to verify that their votes were actually counted and every observer, including voters, election officials, and external observers, should be able to verify whether the final election outcome indeed corresponds to the votes submitted by the voters. This property is not just required for traditional paper-based elections, where votes are typically tallied by several (mutually distrusting) talliers and observers can monitor the tallying process, but it is particularly crucial for electronic voting (e-voting) systems: numerous problems with e-voting systems have been reported in many countries, where votes have been dropped or miscounted (see, e.g., [37, 46, 83, 87] and references therein). This does not come as a surprise since e-voting systems are among the most complex hardware and software systems. Hence, it is virtually impossible to avoid programming errors or subtle security vulnerabilities. Verifiability enables external and internal observers to detect and reject false election results, even when the underlying cause is an unknown programming error or vulnerability. Verifiability therefore has been studied intensively for e-voting systems (see, e.g., [27] for an overview), with Helios [3] being one of the most prominent verifiable e-voting systems used in practice.

**Publishing the Full Tally.** In purely paper-based elections and tallying processes, in order to obtain verifiability it appears unavoidable to publicly reveal the full tally consisting of all (aggregated) individual votes, rather than just the actual election result, such as the winner of the election or the *n* best candidates: arbitrary observers (both voters and external observers) should be able to verify the election outcome, and hence, they have to watch the tallying process and by this learn the full tally. While this is not strictly necessary for e-voting systems, currently most verifiable e-voting systems also have to publish the full tally by design, including, for example, the prominent Helios system. This, however, comes with several downsides:

- *Biased voters*: As mentioned above, some elections consist of several rounds. In particular, they might involve runoff elections. Revealing the intermediate tallies for the candidates during such an election introduces biases that are likely to influence voters' choices for the following rounds, which might be unintended.
- *Embarrassed candidates*: In some elections, for example, within companies or associations, it is unnecessarily embarrassing for the losing candidates to publish the (possibly low) number of votes they received.
- *Weak mandates*: If the winning margin is small, then revealing the full tally can undermine the power of the elected candidate. This might be undesirable.

- *Gerrymandering*: If the distribution of votes in different districts/groups is published, then this information can be used to adjust the specifications of those districts/groups for future elections so that one of the parties has an advantage.
- *Italian attacks*: In preferential voting (e.g., IRV), the individual choices of voters can be essentially unique and thus be regarded as "fingerprints", even if the number of candidates is moderate. This privacy loss can be exploited to perform so-called *Italian attacks* [13, 49]: even a passive adversary, who merely sees the full tally, is able to easily coerce any voter to vote in a specific (unlikely) way and then use the tally to check whether the voter complied.

These issues can be mitigated or resolved by hiding (part of) the tally, except for the election result. Voting systems with this property are called *tally-hiding*. We classify them as follows, depending on their specific objectives and approaches.

**Fully and Partially Tally-Hiding.** Several (verifiable) e-voting systems, see, e.g., [14, 21, 51, 62, 84, 86], have been proposed to address all of the aforementioned issues at once. They only publish the actual election result (e.g., only the name of the winner or the names of the *n* best candidates), and no party, neither internal nor external, gets to know any intermediate results. We call this strong notion of tally-hiding *fully tally-hiding*. They typically employ heavy-weight cryptography, such as universally verifiable Multi-Party Computation (MPC). As a result, existing fully tally-hiding e-voting systems are limited to rather simple voting methods and elections with only a few candidates. They quickly become inefficient for larger numbers of candidates or complex result functions and voting methods, such as ranked voting, say, via IRV (see Section 6 for more discussion).

There are also verifiable e-voting systems that focus on specific issues due to publishing the tally, most notably Italian attacks. They hide exactly those parts of the tally which cause the respective issues (see, e.g., [13, 25, 49, 55, 79]). We call this approach *partially tally-hiding*. Such systems can be more efficient than fully tally-hiding ones. For example, the e-voting system proposed in [79] can handle even complex IRV elections for real world data. But hiding the tally partially comes with the trade-off that some information is still revealed, such as the order of candidates, and hence, some of the problems of publishing the tally remain, such as biased voters and embarrassed candidates (see also Section 6). Thus, these systems are useful to solve certain aspects, but not others.

**Publicly Tally-Hiding.** In this work, we follow an approach that offers a trade-off between privacy and efficiency which is different from all previous fully and partially tally-hiding solutions. This approach is motivated by and enhances existing practices: There have been many cases where voting authorities chose to internally compute but not publish the full tally. They rather only published the final election result, e.g., the winner of the election or the *n* best candidates, as they wanted to mitigate some of the above issues, including privacy issues for voters (Italian attacks), for candidates (embarrassment, weak mandates), and/or manipulations (gerrymandering). Hence, while the talliers learn the full tally, the public learns only the election result. We call this approach *publicly tally-hiding* elections. Such publicly tally-hiding elections are carried out, among others, by ACM Special Interest Groups (SIG) [1],

the German Computer Science Association [42], CrossRef [32], the Society for Industrial and Applied Mathematics (SIAM) [82], or the German Research Fund [35], as confirmed by websites and/or personal communication. Civica Election Services (CES), a large e-voting provider, conducts several dozen elections per year where costumers demand to obtain only the actual election result from CES [78], according to CES, for example, to protect against weak mandates or gerrymandering issues. So while publicly tally-hiding elections are quite common, they, however, so far do not offer verifiability, i.e., it is impossible for a voter or external observer to verify that the election result was computed correctly.

In this work, we close this gap by proposing the first verifiable voting system that follows the common practice of publicly tallyhiding elections. More specifically, we propose a publicly tallyhiding e-voting system which follows a radically different approach than all previous tally-hiding ones. We allow each tallier to learn the homomorphically aggregated votes, while the public merely learns the final election result.<sup>1</sup> This allows for a completely different design of the e-voting system and the use of different cryptographic techniques compared to previous systems for tally hiding elections. For example, we can employ relatively lightweight zero-knowledge proofs as opposed to more heavy-weight universally verifiable MPC. Our publicly tally-hiding e-voting system achieves practical efficiency for dramatically larger numbers of candidates and more complex voting methods than all previous fully tally-hiding systems, while still hiding the full tally from the public, unlike partially tallyhiding systems. This of course comes with the trade-off that now the talliers learn the aggregated votes. Altogether, our paper opens a new line of research that allows for enhancing already existing practices, where voting authorities compute but do not publish the tally, with the fundamental and crucial property of verifiability.

#### **Contributions.**

- We formalize the notion of *publicly tally-hiding*.
- We propose Kryvos, the first provably secure verifiable e-voting system that explicitly targets the common practice of publicly tally-hiding elections, with the advantages mentioned before. Our system is designed in a completely different way than previous e-voting systems for (fully) tally-hiding. It builds on general purpose zero-knowledge proofs, more specifically, the highly efficient Groth16 SNARKs [45]. The core idea is that talliers prove the correctness of the election outcome according to the result function using these SNARKs. While the idea itself is simple, putting this idea to practice is non-trivial:
  - Kryvos follows the general design philosophy of Helios, one of the most prominent and practical verifiable e-voting systems which forms the basis of many other systems. Unfortunately, it is not possible to apply the above idea directly to Helios (cf. Section 2.2). Therefore, Kryvos is a fundamental redesign of Helios.
  - To be of practical use, also with performance that improves upon existing fully tally-hiding systems, the design of Kryvos

requires care and dealing with a number of pitfalls (see Section 2 and Section 4). We carefully evaluate various implementation and design options to come up with a suitable implementation.

 We provide instantiations of Kryvos for various voting methods, including plurality voting, various forms of cumulative elections with multiple votes, and ranked elections, such as IRV. These instantiations come with extensive evaluation and benchmarks, demonstrating the practicality of the system. Among others, we test our system with real-world data of complex IRV elections from Australia. Our benchmarks show that Kryvos can handle all of these voting methods more efficiently than existing fully tally-hiding solutions.

**Structure.** In Section 2, we sketch our design rationale, discuss various design choices, and derive efficient realizations of core building blocks. We then, in Section 3, introduce Kryvos, with its implementation and detailed evaluation presented in Section 4. In Section 5, security and privacy of Kryvos is formally stated and proven. Related work is discussed in Section 6. We conclude in Section 7. Full details and proofs are given in our technical report [54], with the implementation provided in [53].

# 2 DESIGN RATIONALE

Many verifiable e-voting systems follow the concept of the prominent Helios system [3]. Helios is based on a  $(t,n_t)$ -threshold IND-CPA-secure additively homomorphic public-key encryption scheme, such as exponential ElGamal. Essentially, given the encrypted votes of the voters, these encrypted votes are homomorphically aggregated to compute the publicly known encrypted tally Enc(T) consisting of a sequence of ciphertexts containing the total number of votes for each candidate/choice, i.e., there is one ciphertext per candidate/choice. By aggregating votes, Helios breaks the link between voters and their votes, thereby achieving ballot privacy, and additionally hides individual votes within the tally T. The talliers then decrypt  $Enc(\mathbb{T})$ , i.e., all ciphertexts therein, in a distributed way and publish T along with proofs of correct decryption, which reveals the full tally  $\mathbb{T}$  and allows everyone to publicly compute the result of the election  $f_{res}(\mathbb{T})$ , e.g., the winner of the election. In systems which aim for full tally-hiding, i.e., where only  $f_{res}(\mathbb{T})$ , but not  $\mathbb{T}$  is revealed, talliers essentially perform certain forms of MPC on  $Enc(\mathbb{T})$  in order to compute  $f_{res}(\mathbb{T})$ . As already mentioned and further discussed in Section 6, this has several shortcomings, which is why in this work we, for the first time, directly follow the common approach of public tally-hiding, i.e., while talliers may learn  $\mathbb{T}$ , everybody else should only learn  $f_{res}(\mathbb{T})$ .

Intuitively, our idea for Kryvos is to build a system similar to Helios where encrypted votes can still be publicly aggregated and the talliers can then compute  $\mathbb{T}$  from  $Enc(\mathbb{T})$ . The aggregation guarantees that talliers do not learn more about individual votes than talliers in a regular privacy-preserving (non tally-hiding) system, such as Helios. However, instead of publishing  $\mathbb{T}$ , talliers rather internally compute and then publish only the result  $f_{res}(\mathbb{T})$ . To still obtain verifiability in this situation, we employ non-interactive zero-knowledge proofs (NIZKPs) to let (one of) the talliers show that  $f_{res}(\mathbb{T})$  was computed correctly from  $Enc(\mathbb{T})$ . Since NIZKPs are rather lightweight compared to MPC, the hope and rationale is that

<sup>&</sup>lt;sup>1</sup>This aggregated tally might still reveal individual votes to the talliers, depending on how ballots are encoded and aggregations are performed. For simple voting methods, the aggregated tally effectively hides individual votes from the talliers, but it might not for complex voting methods, such as IRV.

this should allow for constructing more efficient e-voting systems that support more kinds of elections in a (publicly) tally-hiding way than existing (full) tally-hiding systems. To support a multitude of existing election types with different result functions, we make use of recent advances in the field of *general purpose zero knowledge proof systems* (cf. Section 2.1), which allow for proving statements with respect to essentially arbitrary functions. While this general idea might seem conceptually simple, it requires careful design and optimization to achieve not only a secure but also a practical solution as outlined in this section and Section 4.

#### 2.1 General Purpose Proof Systems

A general purpose proof system takes as input an arbitrary indicator function  $f_R : \{0,1\}^* \times \{0,1\}^* \rightarrow \{0,1\}$  for some binary relation *R*, such that  $f_R(x,w) = 1$  iff  $(x,w) \in R$  for public statement x and secret knowledge/witness w. It then allows for computing a zero-knowledge proof that proves knowledge of w such that  $f_R(x,w) = 1$ . In our setting, a tallier can use  $x = (Enc(\mathbb{T}), y)$ as public input along with a suitable witness (which, among others, contains some tally  $\mathbb{T}_w$ ) and the function  $f_R$ :  $f_R(x,w) = 1$ iff  $\mathbb{T}_w$  corresponds to the plaintext in  $\text{Enc}(\mathbb{T})$  and  $y = f_{\text{res}}(\mathbb{T}_w)$ . The chosen proof system should be able to handle complex result functions even for a large number of candidates/choices, which, among others, requires reasonably low proof creation and verification times. Of the various different general purpose proof systems [4, 9, 23, 40, 43, 45, 71], zero-knowledge succinct non-interactive arguments of knowledge (SNARKs, for short, where we implicitly assume the zero-knowledge property) currently fit these requirements best. More specifically, we use the highly efficient state-ofthe-art Groth16 SNARK [45] as, offers constant proof size with constant verification time (independently of the function  $f_R$ ), while achieving quite fast polynomial proving time and thus scaling well even for highly complex functions. We briefly discuss our choice of the Groth16 SNARK in comparison with alternative proof systems in Section 6. We note that our construction is based on arithmetic circuits (see next paragraph) and hence is not limited to the Groth16 SNARK. In fact, every proof system which is based on arithmetic circuits (such as [4, 10, 18]) can be used to instantiate Kryvos.

A Groth16 SNARK uses the language of quadratic arithmetic programs (QAPs) to specify the underlying relation *R* respectively indicator function  $f_R$ . Typically, in order to obtain a QAP,  $f_R$  is first expressed as an arithmetic circuit which is then converted to a set of constraints that is finally compiled into a QAP instance. A constraint over *n* variables is an equation  $\sum_{i=1}^{n} a_i u_i \cdot \sum_{i=1}^{n} a_i v_i = \sum_{i=1}^{n} a_i w_i$ , where  $u_i, v_i$  and  $w_i$  are constants that define the constraint. The overall performance (runtime, bandwidth, memory overhead) of a Groth16 SNARK is directly determined by the number of constraints and hence the size of the arithmetic circuit.

A bit simplified, the Groth16 SNARK consists of three algorithms: (Setup,Prove,Verify). The Setup( $f_R$ ) algorithm generates two common reference strings, CRS<sub>EK</sub> (*vealuation key* CRS, needed to create proofs) and CRS<sub>VK</sub> (*verification key* CRS, needed to verify proofs) that depend on  $f_R$ . This creates an instance of Groth16 that is specific to the function  $f_R$ . The CRS<sub>EK</sub> can then be used by anyone to create a proof  $\pi \stackrel{\leq}{\leftarrow}$  Prove(CRS<sub>EK</sub>, x, w) with properties as described above. One can use Verify(CRS<sub>VK</sub>,  $x, \pi$ ) to verify the proof, which requires only a very small CRS<sub>VK</sub>. We note that Groth16 SNARKs are based on bilinear groups whose order depends on the size of the input values. The currently most efficient implementation [81], which we use, uses bilinear groups of size up to 256 bits and supports operations in  $\mathbb{F}_p^*$  with *p* of size up to 255 bits.

# 2.2 Public Tally-Hiding for Systems With Homomorphic Encryption

Our approach of obtaining a verifiable publicly tally-hiding e-voting system by having talliers prove in ZK that  $f_{res}(\mathbb{T})$  was computed correctly from  $Enc(\mathbb{T})$  does not directly apply to all existing evoting systems, which often use homomorphic encryption for tallying votes. For example, simply modifying Helios by letting talliers publish only  $f_{res}(\mathbb{T})$  plus a ZKP showing correctness of the result, instead of publishing T, does not yield a secure system. Such a ZKP would have to be computed locally by one of the talliers, who would require as witnesses all private key shares in order to prove knowledge of  $\mathbb{T}_w$  such that  $\mathbb{T}_w = \text{Dec}(\text{Enc}(\mathbb{T}))$ . Revealing those key shares would in turn allow that tallier to decrypt individual votes and hence break ballot privacy entirely. One interesting option to tackle this problem might be to use the very recent concept of distributed SNARKs which permit several parties, each holding a share of a witness, to compute a SNARK in a distributed and privacy preserving fashion [57, 75]. However, to protect the secrecy of witness shares, these constructions heavily rely on MPC components, so it is unclear whether and in how far this approach would yield voting systems that are more efficient than fully tally-hiding ones. Indeed, benchmarks by [75] indicate that currently available distributed SNARKs are still too slow for our purposes.

In this work, we therefore propose a different approach, namely, constructing an e-voting system based on an additively homomorphic commitment scheme. This follows a line of commitment-based e-voting systems initiated by [30]. In Section 3 we show that using this approach it is possible to build a publicly tally-hiding system such that a single tallier is able to locally and thus efficiently compute a SNARK showing correctness of the result.

#### 2.3 Efficiently Proving Knowledge of the Tally

A SNARK can in principle be used to prove statements for arbitrary NP-relations, however, the resulting performance (in terms of runtime, memory overhead, bandwidth, etc.) quickly deteriorates and becomes impractical for large circuits such as, e.g., cryptographic algorithms (see, e.g., [39, 61]). Therefore, a main challenge in using SNARKs consists of carefully constructing suitable circuits with minimal numbers of constrains for the intended relations such that the resulting SNARKs are practical. While we discuss the overall circuits for our SNARKs in Section 4, we need as a central building block for all of our SNARKs an efficient circuit for verifying decommitments, which we construct in the following.

**Commitment Scheme.** We use Pedersen commitments over a group  $(\mathcal{G}, \cdot)$  of order q with generator g as introduced in [77]. Recall that to compute a Pedersen commitment, one first takes an auxiliary value h defined by  $h = g^a$  for a uniform  $a \stackrel{\$}{\leftarrow} \mathbb{F}_q$  and then commits to a value  $v \in \mathbb{F}_q$  with  $\operatorname{Com}(v,r) = g^v h^r$  using a uniform  $r \stackrel{\$}{\leftarrow} \mathbb{F}_q \setminus \{0\}$ .

Pedersen commitments are a standard solution that have a simple structure and hence also rather small computational complexity.

The choice of the group G significantly affects efficiency of the corresponding SNARK, with some groups being more favorable than others. We use a Montgomery elliptic curve, specifically Curve25591, which is well known to provide good synergy with Groth16 SNARKs. Specifically, coordinates of points on this curve use only 255 bit as provided by [81]. Furthermore, the Montgomery ladder algorithm allows for relatively cheap exponentiation within the SNARK as discussed next.<sup>2</sup>

**Designing the Circuit.** A Pedersen commitment  $g^v \cdot h^r$  requires exponentiations and multiplications, with exponentiations being the most expensive operation. We leverage results from the C0C0framework [61], which observes that the highly efficient Montgomery ladder algorithm for exponentiations can be implemented with relatively few constraints. However, a major limitation of this approach is that the resulting values cannot directly be multiplied.

We therefore extend this approach to also support efficient multiplication of such values. Among others, we achieve this by designing small circuits for the fast *y*-coordinate recovery algorithm by Okeya and Sakurai [74], conversion between projective and affine coordinates, and point multiplication. We construct the overall circuit for computing a Pedersen commitment by combining these components, where combinations are chosen in such a way that they require the least number of constraints. We provide full details of our circuit, including a discussion of choices, in our technical report [54].

**Optimizing Constraints.** A crucial and labor intensive step for obtaining an efficient SNARK is to optimize the implementation via constraints (see e.g., [61, 76]). While one can in principle, specify an arithmetic circuit defined in a suitable language and then automatically compute a representation via constraints, such a generic conversion results in needlessly large numbers of constraints and can lead to impractical SNARKs. Obtaining efficient SNARK proofs is crucial for our voting system. We therefore manually implemented and optimized our circuit via constraints. This includes individual gates and operations such as various comparison and branching operations, elliptic curve point conversions, and multiplications as well as combinations of those gates within specific algorithms, like the Montgomery ladder and the Okea-Sakurai algorithm, and within the overall commitment circuit.

As an example of the many optimizations that we have performed — all of which, even if small, add up quickly — consider a division gate. For divisions, one has to handle the special case of a divisor of value zero. Such cases can occur while evaluating a path not taken by the control flow during branching. Using multiple such divisions is costly. However, it is possible to optimize the number of constraints for the divisions by reusing intermediate results for specific values for multiple divisions. With this, one only needs to perform the error handling once. The same technique can be applied for various other error handling cases, for example for dealing with the point at infinity. Using knowledge of the nesting of the gates allows for even more efficient error handling. For example, if the

<sup>2</sup>To stay compatible with the usual notation for Pedersen commitments, we denote the group operation on an elliptic curve with · instead of the usual +. The wording *addition* and *multiplication* is consequentially changed to *multiplication* and *exponentiation*.

#### CCS '22, November 7-11, 2022, Los Angeles, CA, USA.

	N Pedersen Commitments				1 Pedersen Vector Commitment			
Choices/	Constraints	Prove	CRS <sub>EK</sub>	CRSVK	Constraints	Prove	CRS <sub>EK</sub>	CRSVK
N		[s]	[GB]	[MB]		[s]	[GB]	[MB]
1	6,549	0.903	0.01	0.98	6,549	0.903	0.01	0.98
5	32,745	3.472	0.05	4.9	10,077	1.249	0.01	1.5
10	65,490	6.684	0.09	9.8	14,487	1.681	0.02	2.17
25	163,725	13.107	0.18	24	27,717	2.546	0.03	4.15
50	327,450	32.376	0.46	49	49,767	5.142	0.07	7.45
100	654,900	64.491	0.91	98	93,867	9.467	0.13	14
150	982,350	96.606	1.37	147	137,967	13.792	0.19	20
200	1,309,800	128.722	1.82	196	182,067	18.117	0.25	27
250	1,637,250	160.837	2.28	245	226,167	22.442	0.31	34

Table 1: Comparison of Pedersen commitments and Pedersen vector commitments in a single Groth16 SNARK. We consider inputs v of at most 32 bits.

output of a gate is an elliptic curve point that cannot be the point at infinity, independently of the input values, no such checks have to be performed for the following gate.

**Further Optimizing The Commitment Circuit.** Even with the Montgomery ladder and an optimized QAP representation, the exponentiation step is still very expensive. In our and similar settings, we can further improve this step via two additional optimizations:

Firstly, we can upper bound the length of the input v of the commitment scheme with 32 bits (instead of the full 255 bits), which is sufficient for our application. We discuss the effects of this optimization in our technical report [54].

Secondly, as part of our system, we will need to commit to multiple input values  $v_i$ , say, N many. If one implements this in a straightforward manner via standard Pedersen commitments, then there will be N separate commitments  $com(v_i, r_i)$ , one per  $v_i$ . Each of these commitments introduces another exponentiation for some randomness  $r_i$  to obtain  $h^{r_i}$ , which becomes very costly for large N (cf. left side of Table 1). We solve this issue by using Pedersen vector commitments [16]. Essentially, these commitments allow for committing to a vector  $v = (v_1, \ldots, v_N)$  with entries in  $\mathbb{F}_q$  by computing  $\operatorname{Com}(\boldsymbol{v},r) = g_1^{v_1} \cdot \ldots \cdot g_N^{v_N} \cdot h^r$  for generators  $g_1, \ldots, g_N$  of  $\mathcal{G}$  chosen uniformly at random. In this case, we say that N is the *slot size* of the commitment. One observes that Pedersen vector commitments are also additively homomorphic, perfectly hiding, and computationally binding under the discrete logarithm assumption [5]. This construction requires just one exponentiation with randomness of 255 bits to commit to all N inputs. This in turn drastically improves both the computational cost but also the sizes of the CRSs of the SNARK as shown in Table 1. We will further discuss the optimal choice of N in the context of our voting scheme in Section 4, and our technical report [54] illustrates further possible tradeoffs.

# 3 THE Kryvos SYSTEM

In this section, we present the Kryvos e-voting system, which is the first verifiable e-voting system which directly follows the publicly tally-hiding paradigm. Kryvos is a generic framework that supports many different voting methods and result functions. We present its implementation and benchmarks for specific voting methods and result functions in Section 4.

#### 3.1 System description

We now describe the Kryvos e-voting system. Our description focuses on those parts of Kryvos that are at the core of the public tally-hiding property. Well-known and standard enhancements to security, in particular distributed implementations of bulletin boards (e.g., [33, 52, 58]) and mechanisms to mitigate trust on the voting devices (e.g., [12, 41]), are orthogonal to and fully compatible with Kryvos.

**Participants.** The Kryvos protocol is run among a voting authority Auth, voters  $V_1, \ldots, V_{n_v}$ , talliers  $T_1, \ldots, T_{n_t}$ , and a bulletin board (BB). We assume that for each party there exists a mutually authenticated channel to the BB.

**Voting method.** Kryvos is parameterized by  $n_{\text{choices}} \in \mathbb{N}$  denoting the number of choices of an election and a set of valid voting options  $C \subseteq (\mathbb{F}_q)^{n_{\text{choices}}}$ , the *choice space* over the prime field  $\mathbb{F}_q$ . Both parameters are instantiated depending on the specific voting method that should be captured. For example, for a single-vote election, where each voter can give a single vote to one out of three candidates, we have  $n_{\text{choices}} = 3$  and the choice space  $C_{\text{single}} = \{(a,b,c) \mid a,b,c \in \{0,1\}, a+b+c=1\}$ .

Kryvos is further parameterized by a deterministic polynomial time function  $f_{\text{res}}: \{0,1\}^* \rightarrow \{0,1\}^*$  which computes the *final result* of the election based on a set of aggregated votes, i.e., the tally of the election. For example, if we are interested only in the winner/winners of the election, then  $f_{\text{res}}$  returns the index/indices of the highest entry/entries. The tuple  $(C, f_{\text{res}})$  defines the abstract voting method for which we describe Kryvos; specific instantiations are given in Section 4.

Setup phase. In this phase, all election parameters are fixed and posted on the BB by Auth: the list id of eligible voters, opening and closing times, the election ID  $\mathsf{id}_{\mathsf{election}},$  and the voting method  $(C, f_{res})$ . Additionally, Auth publishes a decomposition  $n_{choices} =$  $n_{\text{tuples}} \cdot N$  that induces a decomposition on the choice space  $C \subseteq$  $(\mathbb{F}_q)^{n_{\text{choices}}} = (\mathbb{F}_q)^N \times \ldots \times (\mathbb{F}_q)^N$  yielding a splitting of a vote  $v \in C$ into  $n_{tuples}$  tuples, each of size N. The reason for this splitting lies in the usage of homomorphic vector commitments as introduced in Section 2. Splitting allows for decomposing a commitment to a vote v into  $n_{\text{tuples}}$  vector commitments with N slots. We will argue in Section 4 that this decomposition can be necessary for very large values of  $n_{\text{choices}}$ . A vote  $v = (v_1, \ldots, v_{n_{\text{tuples}}})$  then contains in  $v_1$  the respective entries for the first N choices, in  $v_2$  the next N choices and so on.<sup>3</sup> This assignment is published by Auth on B. Furthermore, Auth creates and publishes the CRSs required for the Groth16 SNARK proofs on B. Each tallier  $T_k$  runs the key generation algorithm KeyGen & of an IND-CCA2-secure public-key encryption scheme  $\mathcal{E} = (\text{KeyGen}_{\mathcal{E}}, \text{Enc}, \text{Dec})$  to generate its public/private (encryption/decryption) key pair  $(pk_k, sk_k)$  and posts  $(k, pk_k)$  on B.

**Voting phase.** Let  $m^i = (m^{i,1}, \ldots, m^{i,n_{\text{choices}}}) \in C$  be the vote of voter  $V_i$ . The voter creates full-threshold secret sharings of every component  $m^{i,j}$  of her vote to share  $m^{i,j}$  among the  $n_t$  talliers:  $(m_1^{i,j}, \ldots, m_{n_t}^{i,j})$  with  $\sum_{k=1}^{n_t} m_k^{i,j} \mod q = m^{i,j}$ .

Then, for each tallier  $T_k$ , the voter decomposes  $(m_k^{i,1}, \ldots, m_k^{i,n_{choices}})$ into  $(t_k^{i,1}, \ldots, t_k^{i,n_{tuples}})$ , where  $n_{choices} = n_{tuples} \cdot N$  is as defined by Auth in the setup phase (see above). Now the voter creates a commitment to each tuple:  $c_k^{i,l} \leftarrow \text{Com}(t_k^{i,l}; r_k^{i,l})$ . Observe that since the commitment scheme is additively homomorphic, the commitment  $c^{i,l} := \sum_{k=1}^{n_t} c_k^{i,l}$  opens to the tuple  $(m^{i,1+(l-1)\cdot N}, \ldots, m^{i,l\cdot N})$ , i.e., the votes for the choices belonging to the *l*-th factor in the decomposition of  $(\mathbb{F}_q)^{n_{\text{choices}}}$  specified by Auth, using opening value  $r^{i,l} := \sum_{k=1}^{n_t} r_k^{i,l}$ . That is, one can obtain commitments on the original vote of  $V_i$  by combining all of the commitments on the full-threshold shares.

To guarantee the well-formedness of all commitments and the vote contained therein, the voter creates a Groth16 SNARK proof  $\Pi^{i}_{\text{ballot}}$  proving that  $(c^{i,1}, \dots, c^{i,n_{\text{tuples}}})$  commits to a vote  $m^{i} \in C$ ; we provide more information, including formal definitions of the relations shown by the proof, in our technical report [54].

For each tallier  $T_k$ , the voter uses  $T_k$ 's public key  $pk_k$  to securely send the opening values for  $(c_k^{i,1}, \ldots, c_k^{i,n_{tuples}})$  to  $T_k$ :

$$\boldsymbol{e}_k^i \gets \mathsf{Enc}(\mathsf{pk}_k, ((\boldsymbol{t}_k^{i,1}, \dots, \boldsymbol{t}_k^{i,n_{\mathsf{choices}}}), (\boldsymbol{r}_k^{i,1}, \dots, \boldsymbol{r}_k^{i,n_{\mathsf{tuples}}}))).^4$$

To complete the voting process,  $V_i$  submits her ballot  $b^i = (i, (c_k^{i,l})_{l,k}, \Pi_{\text{ballot}}^i, (e_k^i)_k)$  to the bulletin board. B then verifies that i) the voter is eligible to vote, ii) has not submitted a valid ballot before, iii) the voter's SNARK proof is valid, and iv) no voter has previously submitted a vector containing any of the ciphertexts in  $(e_k^i)_k$ .<sup>5</sup> If all checks succeed, then the BB adds  $b^i$  to the list of ballots  $\vec{b}$  and publicly updates  $\vec{b}$ .

Now, the list  $\hat{\mathbf{b}}$  needs to be prepared by the talliers for the tallying phase. Each tallier  $\mathsf{T}_k$  decrypts every  $e_k^i$  posted on B:

$$((\tilde{t}_k^{i,1},\ldots,\tilde{t}_k^{i,n_{\text{tuples}}}),(\tilde{r}_k^{i,1},\ldots,\tilde{r}_k^{i,n_{\text{tuples}}})) \leftarrow \mathsf{Dec}(\mathsf{sk}_k,e_k^i).$$

Then  $T_k$  checks whether each pair  $(\tilde{t}_k^{i,l}, \tilde{r}_k^{i,l})$  is a valid opening for  $c_k^{i,l}$ . If this is not the case, then  $T_k$  publishes a NIZKP of correct decryption of  $e_k^i$  on B so that one can verify that  $e_k^i$  is invalid; as a consequence, the corresponding ballot will not be counted.<sup>6</sup> All of the following steps, including the tallying phase, are performed only for those ballots that have passed this check; for simplicity of presentation, we therefore assume that all voters have submitted a valid ballot.

**Tallying phase.** Everyone can homomorphically aggregate the public commitments on B by computing  $c^{\perp,l} \leftarrow \sum_{k=1}^{n_l} \sum_{i=1}^{n_o} c_k^{i,l}$  for each  $1 \leq l \leq n_{\text{tuples}}$ . In parallel, the talliers homomorpically aggregate the corresponding opening values. First, each tallier  $T_k$  internally computes for each  $1 \leq l \leq n_{\text{tuples}}$ :  $t_k^{\perp,l} \leftarrow \sum_{i=1}^{n_o} t_k^{i,l}$ ;  $r_k^{\perp,l} \leftarrow \sum_{i=1}^{n_o} r_k^{i,l}$ , where the  $t_k^{i,l}$ 's and  $r_k^{i,l}$ 's are the opening values decrypted above. Next, each tallier  $T_k$  shares  $(t_k^{\perp,l})_l$  and  $(r_k^{\perp,l})_l$  with the other

<sup>&</sup>lt;sup>3</sup>If  $n_{\text{choices}}$  is not divisible by N, we can artificially grow the choice space so that  $n_{\text{choices}}$  is a multiple of N by appending zeros to a vote.

 $<sup>^4</sup>$  We assume that all plaintexts have fixed length. This can easily be guaranteed here. This is necessary to provide privacy for the votes.

<sup>&</sup>lt;sup>5</sup>This so-called *ballot weeding* process prevents malicious voters from submitting ballots that are related to the ballots from honest voters, which would break privacy. In particular, due to the CCA2-secure encryption, which must contain the correct randomness for the commitments, the only way to submit a valid related ballot is by creating a new ballot that re-uses some of the (unmodified) ciphertexts from the honest ballots. This is caught by the ballot weeding procedure if at least one tallier is honest (which is always assumed for privacy).

<sup>&</sup>lt;sup>6</sup>For example, one can combine the IND-CCA2-secure PKE by Cramer-Shoup [31] with the NIZKP by Camenisch-Shoup [19]. See our technical report [54] for the precise relation to be proven.

talliers so that for each  $1 \leq l \leq n_{\text{tuples}}$ , the trustees can internally compute  $t^{\perp,l} \leftarrow \sum_{k=1}^{n_t} t_k^{\perp,l}$ ;  $r^{\perp,l} \leftarrow \sum_{k=1}^{n_t} r_k^{\perp,l}$ . After this step, all talliers can compute the (aggregated) tally of all votes  $\mathbb{T} := (m^{\perp,1}, \ldots, m^{\perp,n_{\text{choices}}})$ , where  $m^{\perp,j}$  is the total number of votes for choice j.

The final step is then performed by a designated tallier. First, this tallier computes the election result res  $\leftarrow f_{\text{res}}(\mathbb{T})$ . She then computes a Groth16 SNARK on public inputs  $c^{\perp,1}, \ldots, c^{\perp,n_{\text{tuples}}}$  and res that proves knowledge of  $\mathbb{T}$  (as well as knowledge of randomness) such that  $c^{\perp,1}, \ldots, c^{\perp,n_{\text{tuples}}}$  is a list of commitments on  $\mathbb{T}$  and res is the output of  $f_{\text{res}}(\mathbb{T})$  (cf. our technical report [54] for the formal definition of the relation). The result res and the SNARK proof are then published on the BB.

**Public verification phase.** In this phase, every participant, including the voters or external observers, can verify the correctness of the tallying procedure, in particular, the correctness of all ZKPs. Note in particular that homomorphic aggregations of commitments can be re-computed by external observers without knowing any openings.

#### 3.2 Discussion

**Using SNARKs to prove ballot validity.** Using a general Groth16 SNARK to prove ballot validity has the advantage that Kryvos can support essentially arbitrary choice spaces and hence voting methods, including very complex ones such as ranked voting methods and methods that assign points from a specific set and according to certain rules to each candidate. As we show in Section 4.2, this approach is indeed practical. This support for arbitrary ballot formats is also of interest beyond the area of (publicly) tally-hiding systems. For example, as far as we are aware Kryvos is the first e-voting system that supports ZKPs for Borda tournament style ballots. For Condorcet ballots, Kryvos improves over previous efficient ZKPs that even required talliers to perform part of the proof [48].

Kryvos can in principle also be combined with other existing ZKPs for ballot correctness that are designed for a specific choice space, where such ZKPs are available (e.g., [29]).

**Generating CRSs.** Note that our system requires honestly generated CRSs for the Groth16 SNARKs. There are several well-known mechanisms that are practical with the parameters we need, such as distributed generation of the CRS by multiple parties. We briefly discuss such mechanisms in our technical report [54] for completeness. For simplicity and since this is an orthogonal issue, we here compute CRSs on a local computer.

# 4 IMPLEMENTATION OF Kryvos AND EVALUATION

We provide a proof of concept implementation of Kryvos that we have instantiated for a wide range of both simple and complex voting methods as well as various result functions; our implementation is available at [53]. We make use of the libsnark library [81] for the Groth16 SNARK. While in Section 3 we gave a high-level description of Kryvos, there are various ways in how Kryvos can be implemented, involving among others fixing a slot size *N*, a suitable choice space, designing and optimizing circuits for both ballot and tallying SNARKs. Coming up with an efficient implementation needs care. We follow various implementation paths, carefully evaluate them, and by this develop a practical implementation.

All of the following benchmarks, were obtained using Ubuntu 20.04 with 16 GB of RAM and eight cores. All benchmarks are performed without leveraging parallelism, i.e., we use just a single core to make the results independent of the specific core count, thereby making it easier to compare with other benchmarks. Of course, in practice one would parallelize, e.g., if multiple SNARK proofs need to be computed, thereby reducing the overall runtime depending on the number of cores available.

#### 4.1 The Optimal Slot Size

Recall from Section 3.1 that Kryvos uses Pedersen vector commitments, where the number of inputs (slot size) N can be chosen arbitrarily and in particular independently of the number of choices  $n_{\text{choices}}$ . Before we present our concrete instantiations, we first discuss and evaluate the choice of N in relationship to  $n_{\text{choices}}$ .

The slot size *N* heavily influences the circuit size for proving a correct opening of the commitments, which is also the main determining factor for the size and hence the performance of the overall ballot and tallying SNARKs (cf. in Section 4.2). Hence, the runtime and bandwidth of ballot/tallying SNARKs for N = 1 and  $N = n_{\text{choices}}$  roughly corresponds to the left and right columns of Table 1, with each row indicating a possible value for  $n_{\text{choices}}$ . We have evaluated the exact performance of showing an opening of an entire ballot/tally (possibly consisting of several separate commitments) via a SNARK not just for these extremes but also for possible intermediate choices, e.g., having two commitments with  $N = \frac{n_{\text{choices}}}{2}$ , and provide full results with an in-depth discussion in our technical report [54].

To summarize our findings, the combined runtime and bandwidth used for ballot and tallying SNARKs for both proving correctness of a ballot and correctness of SNARKs is minimal for  $N = n_{\text{choices}}$ . As per Table 1, this yields practical performance up to  $n_{\text{choices}} = 250$  choices, which is sufficient for most of our use cases and will therefore be used in what follows. The only exception are IRV elections, where  $n_{\text{choices}} \gg 250$  (e.g.,  $n_{\text{choices}} = 1,957$ for 6 candidates). The limiting factor in this setting is the ballot SNARK: For  $N = n_{choices}$ , computing the SNARK would require downloading a CRS of size 2.4 GB and computing 170 seconds, which is typically not considered acceptable for a voter. By setting N to instead be a large fraction of  $n_{\text{choices}}$ , say,  $N = \frac{n_{\text{choices}}}{8}$ , the overall runtime for proving knowledge of a ballot will be slightly larger but one can compute this statement in parallel using multiple much smaller SNARKs. This leads to a smaller CRS and reduces overall runtime to an acceptable level for voters with multi core CPUs. We discuss this in detail in Section 4.3.

# 4.2 Implementation and Evaluation of Kryvos for various voting methods/result functions

We have instantiated and evaluated Kryvos for a variety of voting systems and result functions. As it turns out, there are essentially two main requirements for a voting method to be supported by Kryvos in practice: Firstly, the voting method must support a ballot format that allows for homomorphic aggregation of ballots. This is necessary to protect the privacy of voters against the talliers. Secondly, it must be possible to define a choice space such that the resulting ballots, including the corresponding ballot SNARKS, can be created efficiently.

We note that a voting method essentially defines a choice space which is in turn required for proving ballot validity. More precisely, proving ballot validity includes proving that a vote belongs to the respective choice space. An election result function, however, does not (only) depend on the voting method and its corresponding choice space. Some result functions only make sense for certain choice spaces though. For systems like Single-Vote and Multi-Vote, we consider various result functions that e.g. compute only the candidate with the most votes or all candidates that gained at least a certain threshold of votes. The chosen result function only influences the tallying SNARK but not the ballot SNARK.

Concretely, we have implemented the following popular and widely used voting methods and result functions for Kryvos with full details provided in our technical report [54]: (i) Single-Vote and Multi-Vote with the result functions Most Votes. Vote Threshold and Best n (unordered), (ii) Borda Count (as used in the Eurovision Song Contest [38] and for parliamentary elections in the Republic of Nauru [80]) with the same result functions as Single-Vote and Multi-Vote, (iii) Majority Judgement [6] (used for example in political research polling in the US, France and Germany [7]) with result functions to compute a Median Grade and complete Majority Judgement Evaluation, (iv) Condorcet Methods (as used for internal elections by the Debian Project [34]) with a result function for the computation of the Smith Set, (v) and Instant Runoff-Voting (IRV) (as used used for example in political elections in Australia [73], India [44], the UK [85], and the US [70]) where we implemented the result function used for the New South Wales Legislative Assembly election [73]. In what follows, we give an overview of our benchmarking results. The IRV election method is then discussed in detail in Section 4.3. Full details of the remaining voting methods and result functions, including descriptions of those methods and formal definitions of the corresponding choice spaces, are available in our technical report [54].

**Tallying for various result functions.** The tallying phase mainly consists of aggregating the ballots and computing the final Groth16 SNARK proof that shows correctness of the election result. Aggregation of Pedersen (vector) commitments is very fast and can mostly be performed already during the previous voting phase. Hence, the driving factor and focus of our benchmarks lies in the final Groth16 SNARK creation. For all result functions, the tallying SNARK essentially takes the aggregated commitment(s) on the tally and the election result as public input, the commitment opening(s) as secret input, and then proves that the opening corresponds to the commitment and that the election result was obtained by applying  $f_{\rm res}$  to the tally contained in the opening. In our technical report [54], we provide formal definitions of all relations shown by the SNARKs and NIZKPs (for the encryptions of the ballots) used for our instantiations.

The runtime for computing the resulting tallying SNARKS for all of our result functions, including IRV, is shown in Figure 1, where the runtime scales linearly in the number of choices  $n_{choices}$ . Our benchmarks show that SNARK computation and hence the tallying

Comparison Between Ordinos and Kryvos



Figure 1: Benchmarks of all implemented election result functions in comparison to Ordinos. Note that the x-axis uses the number of choices instead of the number of candidates. phase of Kryvos is very fast even for large numbers of choices, e.g., almost all voting methods and result functions (green line) require only about 100s for  $n_{choices} = 1000$  choices. Furthermore, we can see that in most cases the result function has only a negligible impact on the overall runtime, which is rather dominated by proving knowledge of a decommitment and hence practically identical in almost all cases. The only exception is the majority judgment voting method with 6 grades and full result evaluation returning a single winner. The algorithm for computing the result is so complex that it has a small but noticeable effect on the overall runtime, as can between seen in Figure 1. Figure 1 also shows the runtime of the tallying phase of two instantiations of Ordinos, the most current fully implemented and provably secure verifiable and fully tallyhiding voting system, which illustrates that Kryvos indeed achieves its goal of much better performance over fully tally-hiding systems. We give a more detailed comparison of Kryvos with Ordinos and the system by Canard et al. [21] for Majority Judgement in Section 6.

While in many voting methods, such as single and multi vote elections, a single choice directly corresponds to one candidate, this is not always the case. E.g., the in case of majority judgement with r grades, we need r choices to represent each candidate (essentially, the voter sets one out of r choices to 1 with all other choices being 0 to indicate a grade in such a way that it can be aggregated). In the case of Condorcet voting, which is a ranked voting method where every candidate is compared with every other candidate, the number of choices is even quadratic in the number of candidates. Therefore, Figure 2 translates the runtime results of Figure 1 from numbers of choices n<sub>choices</sub> to the corresponding number of supported candidates. Additionally, Figure 2 also shows the sizes of the CRSs. As can be seen, in most cases even elections with more than 80 candidates can still be tallied in under 3 minutes. This candidate size should already cover most elections from practice. But our system also scales quite well for much larger numbers. The only exceptions are IRV (not shown in this figure but discussed separately in Section 4.3) and Condorcet (with the Smith set evaluation function, but the same should also apply for other result functions), which scales worse due to the quadratic translation from choices to candidates. We note, however, that Condorcet elections are typically performed only for small to moderate numbers of candidates (less than a dozen) since voters have to fully rank all candidates on their ballot, which beyond already 20 candidates becomes very tedious and impractical. So for practical purposes the performance of Kryvos when applied to Condorcet appears to be sufficient.

**Ballot validity.** To show validity of a ballot (for all voting methods except IRV, which is discussed in Section 4.3), a voter computes



Figure 2: Benchmarks of Different Election Result Functions.



### Figure 3: Benchmarks of SNARKs for proving Ballot validity. Note that the line for Single-Vote coincides with the line for Multi-Vote.

a single Groth16 SNARK that takes the single Pedersen vector commitment of slot size  $N = n_{choices}$  obtained by aggregating the commitments on all vote shares of a voter as public input, the opening as secret input, and then proves that the opening is from the correct choice space and corresponds to the commitment. Just as for the tallying SNARK, the runtime for computing a ballot SNARK is dominated by proving knowledge of a decommitment. Hence, runtime is linear in the number of choices  $n_{choices}$  and essentially independent of the choice space. In Figure 3, we show the runtime translated to the number of candidates as well as the sizes of the CRSs. As this figure shows, voters can efficiently prove the validity of their ballots for all voting methods. Even in the worst case of Condorcet in an election with, say, 20 candidates (which is at the very high end for Condorcet), a voter can prove ballot validity in less than 40 seconds using a CRS<sub>EK</sub> of less than 0.5 GB.

**Public Verification Phase.** Each Groth16 SNARK proof can be verified in ~15 ms using the smaller CRS<sub>VK</sub>, where a single proof is of size ~ 5 KB. As an example, consider a single vote election ( $C_{single}$ ) with an arbitrary number of talliers and up to  $n_{choices} = 1,000$  candidates and  $n_v = 100,000$  voters. Such an election requires 100,000 Groth16 SNARK proofs for showing well-formedness of ballots and a single Groth16 SNARK proof for showing correctness of the election result. In total, the size of all proofs is ~ 500 MB with a total sequential verification time of ~ 25 minutes. Verification requires two CRS<sub>VK</sub>s, both of them having size less than 150 MB. We note that this verification can be highly parallelized and also already performed while other phases, such as the voting phase, are still running.

# 4.3 Instantiation of Instant-Runoff Voting

Instant-runoff-voting (IRV) is a quite complex ranked voting method used for example in political elections in Australia [73], India [44],

the UK [85], and the US [70]. In IRV, the ballot of a voter contains a (full or partial) ranking of all of the candidates according to her preferences. Tallying of these ballots is then done in multiple rounds: in each round, if a candidate is ranked first by an absolute majority of voters, this candidate immediately wins the election. Otherwise, the candidate who is ranked first by the least number of voters is eliminated. As a result of this elimination, all ballots are edited by removing the eliminated candidate, allowing all of the following candidates to move up in the ranking. This process is then iterated. To handle possible ties, IRV versions include different tie-breaker mechanisms to decide how the round proceeds. One of the more complex versions of IRV is used for the New South Wales Legislative Assembly election [73] (called NSW-IRV in the following), which not only supports partial rankings of candidates but also includes a sophisticated tie-breaking mechanism (cf. our technical report [54]). We therefore use NSW-IRV for our benchmarks.

For an IRV instantiation of Kryvos, we need to be able to aggregate the tally. This ensures a minimal level of vote privacy towards the talliers and that runtime of the tallying phase is independent of the number of voters. We achieve this by using the choice space  $C_{\text{single}}$  from the single vote method but interpret each choice as a (full or partial) ranking of candidates. For example, for  $n_{\text{cand}} = 5$ , we have  $n_{\text{choices}} = 120(= 5!)$  choices for fully ranked IRV, where each choice represents a permutation of the full set of candidates, and  $n_{\text{choices}} = 326$  for partially ranked IRV (as in NSW-IRV), where each choice represents a permutation of a partial set of candidates. The result function  $f_{\text{res}}$  of Kryvos is set to capture the multiple tallying rounds, including modifications to the ballots after each round as well as tie breaker mechanisms of the specific IRV version.

However, this choice space grows exponentially in the number of candidates and thus quickly leads to situations where voters cannot prove validity of their ballots in reasonable time (for the slot size  $N = n_{choices}$ ). E.g., while the case of  $n_{cand} = 5$  and hence  $n_{choices} = 326$  is still manageable for voters (the ballot SNARK requires only 30 seconds to compute), for  $n_{cand} = 6$  and hence  $n_{choices} = 1,957$  computing the ballot SNARK already requires 3 minutes, which is typically not considered acceptable. One option to try and improve this situation would be to replace the ballot SNARKs with traditional specialized ZKPs that are optimized for the choice space  $C_{single}$ . However, as we show in our technical report [54], this approach only improves ballot ZKP runtime for very small slot sizes N close to 1, in which case the runtime of the tallying SNARK becomes entirely impractical.

However, we can actually improve the runtime of the ballot SNARK by splitting the ballot into a small number of parts and then showing multiple SNARKs in parallel. Specifically, encoding a ballot via  $n_{\text{tuples}} > 1$  separate commitments, each having slot size  $N \approx n_{\text{choices}}/n_{\text{tuples}}$ , one can show the well-formedness of a ballot belonging to  $C_{\text{single}} := \{(x_1, \ldots, x_{n_{\text{choices}}}) \mid x_i \in \{0,1\}, \sum_{i=1}^{n_{\text{choices}}} x_i = 1\}^7$  via multiple sub-statements, where each statement is concerned only with one commitment, namely:

 (a) For each commitment c<sup>i,j</sup>, 1 ≤ j < n<sub>tuples</sub>: all slots x<sub>(j-1)·N+1</sub>,..., x<sub>j·N</sub> in c<sup>i,j</sup> are either 0 or 1.

<sup>&</sup>lt;sup>7</sup>This approach is also applicable to certain other choice spaces.

CCS '22, November 7-11, 2022, Los Angeles, CA, USA.

# Commitments (= $n_{tuples}$ )	1	2	3	5	10
Slot Size N	1,957	979	653	392	196
Ballot SNARK: # Constraints	1,741,532	874,046	584,884	353,377	179,525
Ballot SNARK: CRS <sub>EK</sub> [GB]	2.42	1.22	0.81	0.49	0.25
Ballot SNARK: CRS <sub>VK</sub> [MB]	261	131	88	53	27
Ballot SNARK: Prove [s]	171	85	57	34	17
Ballot SNARK: Total Time	171	257	230	209	196
Tallying SNARK: # Constraints	1,748,926	1,755,475	1,762,024	1,774,240	1,802,575
Tallying SNARK: CRS <sub>EK</sub> [GB]	2.43	2.44	2.45	2.47	2.5
Tallying SNARK: CRS <sub>VK</sub> [MB]	262	263	264	266	270
Tallying SNARK: Prove [s]	171	172	173	174	177

**Table 2: Comparison of various slot sizes for NSW-IRV with**  $n_{\text{cand}} = 6$  and thus  $n_{\text{choices}} = 1,957$ .

- (b) For  $c^{i,n_{tuples}}$ : the first slots  $x_{(n_{tuples}-1)\cdot N+1}, \ldots, x_{n_{choices}}$  are either 0 or 1, and all remaining slots  $x_{n_{choices}+1}, \ldots, x_{n_{tuples}\cdot N}$  are 0.
- (c) For c<sup>i</sup> := ∑<sub>j=1</sub><sup>n<sub>tuples</sub> c<sup>i,j</sup>: The commitment contains a single slot with value 1 and all other slots are 0.<sup>8</sup>
  </sup>

In our technical report [54] we describe a mechanism that allows for using single CRS for showing all of the above statements, instead of having three separate CRS each of them with roughly the same size. Using this method, we have to determine a suitable slot size N (and hence corresponding  $n_{tuples}$ ). A smaller N allows for higher parallelism up to the number of CPU cores available and hence improves runtime for the ballot SNARK. At the same time, a smaller N increases the combined runtime of all ballot sub-SNARKs (as shown in Section 4.1 but also since (c) introduces an additional decommitment over a new aggregated commitment, which is not necessary if all statements are shown in a single SNARK) and of the single tallying SNARK, which cannot be parallelized. Table 2 shows the benchmarks for various possible slot sizes for an IRV election with 6 candidates. The line "Ballot SNARK: Prove" indicates the runtime for computing a single substatement and hence also indicates the overall runtime if all statements (1 statement for  $n_{tuples} = 1$  and  $n_{\text{tuples}}$  + 1 statements otherwise due to the addition of (c)) can be computed in parallel. The line "Ballot SNARK: Total Time" gives the combined sequential runtime. As the table shows, even if we use only  $n_{\text{tuples}} = 3$  and hence  $N \approx n_{\text{choices}}/3$ , then a voter with a commonly available quad core CPU can already construct a ballot in less than a minute, which is manageable. The performance progressively improves as a larger number of cores is available, while the tallying SNARK remains almost as efficient as for  $n_{tuples} = 1$ , i.e.,  $N = n_{choices}$ . Hence, using the above construction Kryvos can also handle IRV elections with 6 candidates.

Following the case study of [79], we illustrate the overall performance of Kryvos by using real election data from the 2015 New South Wales state election for the Legislative Assembly [36]. More specifically, we consider the electoral districts of Albury (five candidates) and Auburn (six candidates), which were also target elections for a partially tally-hiding election system proposed in [79]. Our results are given in Table 3, showing that such (complex) elections Nicolas Huber et al.

District	Albury		Auburn	
Number of Candidates	5		6	
Number of Voters	46347		43783	
	Kryvos	[79]	Kryvos	[79]
Tallying	30 s	2h	177 s	15 h
Tallying: SNARK Proof [s]	30		177	
Tallying: SNARK CRS <sub>EK</sub> Size [MB]	450		2,500	
Tallying: SNARK CRS <sub>VK</sub> Size [MB]	47		270	

Table 3: Benchmarks of NSW-IRV. We note that the benchmarks of [79] are taken directly from [79] which uses a setup comparable to the one we used for Kryvos, cf. our technical report [54]. For Kryvos, we used a slot size of 196.

can be performed in a publicly tally-hiding way using Kryvos. considered.

# **5 SECURITY**

In this section, we formally show that Kryvos enjoys verifiability and privacy properties as desired. For this purpose, we, for the first time, define the notion of publicly tally-hiding e-voting (Section 5.3).

#### 5.1 Computational Model of Kryvos

We start by formally modeling Kryvos using a general and established computational framework (see, e.g., [27, 65, 66]) that we can use both for analyzing verifiability and privacy of Kryvos.

The computational model introduces the notion of a process which can be used to model protocols (we recall some details in our technical report [54].). Essentially, a process  $\hat{\pi}_{\rm P}$  modeling some protocol P is a set of interacting ppt Turing machines which capture the honest behavior of protocol participants. The protocol P runs alongside an adversary A, modeled via another process  $\pi_A$ , which controls the network and may corrupt protocol participants; here we assume static corruption. We write  $\pi = (\hat{\pi}_{\rm P} || \pi_{\rm A})$  for the combined process. Kryvos can be modeled in a straightforward way as a protocol  $P_{Kryvos}(n_v, n_t, C, f_{res}, \mu)$  in the above sense. Recall from Section 3 that we denote the number of voters by  $n_v$ , the number of talliers by  $n_t$ , and the voting method by  $(C, f_{res})$ . By  $\mu$  we denote a probability distribution over C according to which each honest voter makes her choice. (Note that by this we model that the adversary knows this distribution.) This choice is called the actual choice of the voter. Besides the parties mentioned above, Kryvos contains a BB. In our model of Kryvos, the voting authority Auth is part of an additional agent, the scheduler S. Besides playing the role of the authority, S schedules all other agents in a run according to the protocol phases. The trust assumptions are as explained and motivated at the beginning of Section 3.1 and the end of Section 3.2. In particular, we assume that the voting authority Auth (more generally, the scheduler S) and the BB are honest, i.e., are not corrupted by the adversary.

#### 5.2 Verifiability

In this section, we establish the level of verifiability provided by Kryvos. To this end, we use the generic and widely used verifiability definition proposed in [65] which we briefly recall first. This definition has already been applied to numerous e-voting protocols and building blocks thereof [17, 47, 62–68].

**Verifiability Framework.** The verifiability definition [65] assumes a "virtual" entity, called the *judge* J, whose role is to either

<sup>&</sup>lt;sup>8</sup>While (a), (b), and (c) in conjunction prove the well-formedness of a ballot in ZK, they actually do not prove knowledge of a witness. Intuitively, this is because one can re-arrange the components of a ballot with valid proofs to create a new ballot with valid proofs, even without knowing the contents of the commitments. However, Kryvos actually only requires a zero knowledge proof of correctness in order to be secure (cf. Section 5). Alternatively, one could add the position of each commitment as additional public input.

accept or reject a protocol run. In a real election, the program of the judge can be executed by any party, including external observers and even voters themselves. The judge takes as input solely public information (e.g., the zero-knowledge proofs in Kryvos published on the BB to perform certain checks). In the context of e-voting, for verifiability to hold, the judge should only accept a run if "the announced election result corresponds to the actual choices of the voters". This statement is formalized via the notion of a *goal*  $\gamma$  of a protocol P. A goal  $\gamma$  is simply a set of protocol runs for which the mentioned statement is true, where the description of a run contains the description of the protocol, the adversary with which the protocol runs, and the random coins used by these entities.

Now, following [65] (see also our technical report [54]), we say that a goal  $\gamma$  is *verifiable* by the judge J in a protocol P, if the probability that J accepts a run *r* of P even though the goal  $\gamma$  is violated (i.e.,  $r \notin \gamma$ ) is negligible in the security parameter.

**Analysis.** We prove the verifiability result for Kryvos under the following assumptions:

(V1) The encryption scheme is correct, the commitment scheme is homomorphic and computationally binding, and all NIZKPs are (computationally) sound.

(V2) The scheduler S, the judge J (see our technical report [54]), and the BB are honest:  $\varphi = hon(S) \wedge hon(J) \wedge hon(B)$ .<sup>9</sup>

Note that an arbitrary number of voters and talliers may be controlled by the adversary.

We instantiate the verifiability definition with the goal  $\gamma(\varphi)$  proposed in [27], which captures the intuition of  $\gamma$  given before and with  $\varphi$  as above (see our technical report [54]).

THEOREM 5.1 (VERIFIABILITY). Under the assumptions (V1) and (V2), the goal  $\gamma(\varphi)$  is verifiable by the judge J in the protocol  $P_{Kryvos}(n_v, n_t, C, f_{res}, \mu)$ .

Our formal proof of the verifiability theorem is provided in our technical report [54]. This result mainly uses the soundness of the NIZKPs/SNARKs employed in Kryvos. The underlying relations of these proofs (see our technical report [54]) yield a "global" relation which effectively ensures the goal  $\gamma(\varphi)$ .

### 5.3 Privacy

Our privacy analysis of Kryvos makes use of the privacy definition for e-voting protocols proposed in [66]. This definition is particularly useful for our purposes as it allows us to *measure* the level of privacy a protocol provides, especially compared to protocols that publish the full tally, unlike other definitions (see, e.g., [15]). We first briefly recall the privacy definition from [66] (Definition 5.2). We then formally introduce the novel notion of publicly tally-hiding e-voting (Definition 5.3) and show that Kryvos meets this notion.

**Privacy Framework.** The definition proposed in [66] formalizes privacy of an e-voting protocol as the inability of an adversary to distinguish whether some voter  $V_{obs}$  (the *voter under observation*) who runs her honest program voted for choice  $m_0$  or choice  $m_1$ .

To define this notion formally, we first introduce the following notation for an arbitrary e-voting protocol P for voting method  $(C, f_{res})$ . Given a voter  $V_{obs}$  and  $m \in C$ , we consider instances of P of

the form  $(\hat{\pi}_{V_{obs}}(m) \| \pi^* \| \pi_A)$  where  $\hat{\pi}_{V_{obs}}(m)$  is the honest program of the voter  $V_{obs}$  under observation who takes *m* as her choice,  $\pi^*$  is the composition of programs of the remaining parties in P, and  $\pi_A$  is the program of the adversary. In the case of Kryvos,  $\pi^*$ includes the scheduler, the BB, all other voters, and all talliers.

Let  $\Pr[(\hat{\pi}_{V_{obs}}(m) \| \pi^* \| \pi_A)^{(\ell)} \mapsto 1]$  denote the probability that the adversary writes the output 1 on some dedicated tape in a run of  $(\hat{\pi}_{V_{obs}}(m) \| \pi^* \| \pi_A)$  with security parameter  $\ell$  and some  $m \in C$ , where the probability is taken over the random coins used by the parties in  $(\hat{\pi}_{V_{obs}}(m) \| \pi^* \| \pi_A)$ .

Now, vote privacy is defined as follows, where for Kryvos we quantify over all adversaries  $\pi_A$  which neither corrupt the BB nor the scheduler S.

Definition 5.2 (Privacy). Let P be a voting protocol,  $V_{obs}$  be the voter under observation, and  $\delta \in [0,1]$ . Then, P achieves  $\delta$ -privacy, if for all choices  $m_0, m_1 \in C$  and all adversaries  $\pi_A$  the difference

 $\Pr[(\hat{\pi}_{\mathsf{V}_{\mathsf{obs}}}(m_0) \| \pi^* \| \pi_{\mathsf{A}})^{(\ell)} \mapsto 1] - \Pr[(\hat{\pi}_{\mathsf{V}_{\mathsf{obs}}}(m_1) \| \pi^* \| \pi_{\mathsf{A}})^{(\ell)} \mapsto 1]$ 

is  $\delta\text{-bounded}$  as a function of the security parameter  $1^{\ell}.^{10}$ 

In other words, the level  $\delta$  is an upper bound of an arbitrary adversary's advantage to "break" vote privacy. Therefore,  $\delta$  should be as small as possible. Note, however, that even for an ideal e-voting protocol with a completely passive adversary,  $\delta$  might not be 0 (depending on the result function): for example, if the full tally is published as part of the result, then there might be a non-negligible chance that all honest voters, including the voter under observation, voted for the same candidate, in which case the adversary can easily derive from the final tally how the voter under observation voted.

**Ideal Privacy.** Often, and this will also be the case for Kryvos, formal privacy results are formulated w.r.t. the privacy level  $\delta^{\text{ideal}}_{(n_v, n_v^h, \mu)}$ 

 $(C, f_{res})$  an ideal voting protocol  $I_{voting}(n_v, n_v^h, C, f_{res}, \mu)$  for voting method  $(C, f_{res})$  provides. In this protocol (cf. our technical report [54]), honest voters pick their choices according to the distribution  $\mu$ . In every run, there are  $n_v^h$  many honest voters and  $n_v$ voters overall. The ideal protocol collects the votes of the honest voters and the dishonest ones (where the latter ones are independent of the votes of the honest voters) and outputs the result according to the function  $f_{res}$ .

A generic formula of  $\delta_{(n_v,n_v^h,\mu)}^{\text{ideal}}(C, f_{\text{res}})$  for arbitrary voting methods  $(C, f_{\text{res}})$  was published in [62]. In this work, for the first time, we instantiate  $\delta_{(n_v,n_v^h,\mu)}^{\text{ideal}}(C, f_{\text{res}})$  for complex ranked-choice voting methods, such as IRV. We depict concrete values of the ideal privacy level for the IRV function  $f_{\text{IRV}}$  for two different distributions in Figure 4. While one distribution is uniform, the other is modeled more realistically: we sorted the candidates into a political spectrum and assumed that if a voter chooses a candidate as her first preference, she will most likely rank a candidate with a similar political opinion as rank 2 (see our technical report [54] for details). As the figure shows, revealing the full tally leads to a very low level of privacy, as expected, causing severe privacy issues, like for example Italian attacks. Comparing these levels with the ones revealing only the winner demonstrates that voting systems which hide the tally for

<sup>&</sup>lt;sup>9</sup>See our technical report [54] for more details on  $\varphi$ .

<sup>&</sup>lt;sup>10</sup>A function *f* is  $\delta$ -bounded if, for every c > 0, there exists  $\ell_0$  such that  $f(\ell) \leq \delta + \ell^{-c}$  for all  $\ell > \ell_0$ . Also see e.g. [72] for an explicit formula for  $\delta$ .

CCS '22, November 7-11, 2022, Los Angeles, CA, USA.



Figure 4: Level of privacy ( $\delta$ ) for the ideal protocol with 5 candidates, IRV voting, and no dishonest voters.  $f_{IRVcomplete}$  reveals the complete tally and  $f_{IRV}$  only reveals the winner of the election.

ranked-choice voting methods provide dramatically better vote privacy than those voting systems which always reveal the complete tally. Even more: we see that hiding the tally is in fact *necessary* to achieve a reasonable privacy level for IRV elections that include more than just a few candidates. For simple voting methods, privacy levels comparing tally-hiding and non-tally-hiding systems can be found in [62].

**Publicly Tally-Hiding.** We now introduce the novel notion of publicly tally-hiding e-voting. Intuitively, an e-voting protocol P (like Kryvos) is publicly tally-hiding for some voting method (C,  $f_{res}$ ) if the following two conditions hold true:

- (1) *Public privacy:* Under the assumption that all talliers are honest, P provides the same level of privacy as the ideal voting protocol  $I_{\text{voting}}$  for voting method (*C*, *f*<sub>res</sub>), which reveals nothing but the actual election result by definition. That is, no one, except for the talliers, learns anything beyond the published election result.
- (2) *Internal privacy*: Under the assumption that less talliers than a certain threshold *t* are dishonest, P provides the same level of privacy as the ideal voting protocol  $I_{\text{voting}}$  for voting method (*C*,  $f_{\text{complete}}$ ), where  $f_{\text{complete}}$  is the function that returns the *full tally* (e.g., the number of votes for all choices/candidates). In other words, the talliers do not learn more than they would for a non-tally-hiding secure e-voting protocol.

We now define this notion formally. We assume some set  $\mathcal{T}$  of talliers and some threshold t; for Kryvos we have  $\mathcal{T} = \{\mathsf{T}, \ldots, \mathsf{T}_{n_t}\}$  and  $t = n_t$ .

Definition 5.3 (Publicly Tally-Hiding). Let P be a voting protocol with a set of talliers  $\mathcal{T}$  and  $t \leq |\mathcal{T}|$ . We say that P is  $(\delta_p, \delta_i)$ -publicly tally-hiding w.r.t.  $(\mathcal{T}, t)$  iff the following two conditions hold true:

**Public Privacy:** If all parties  $T \in \mathcal{T}$  are honest, then P achieves  $\delta_p$ -privacy.

**Internal Privacy:** If at most t - 1 parties  $T \in \mathcal{T}$  are dishonest, then P achieves  $\delta_i$ -privacy.

We call  $\delta_p$  the public and  $\delta_i$  the internal privacy level.

As explained, the public and internal privacy levels of a secure protocol should correspond to privacy levels of the ideal protocols mentioned above. However, we do not explicitly require this in the definition in order to be able to use this definition to measure the privacy level of an e-voting system. Moreover, we note that for all "reasonable" publicly tally-hiding voting protocols, the public privacy level  $\delta_p$  should be much better (closer to 0) than the internal privacy level  $\delta_i$ .

**Analysis.** To analyze the publicly tally-hiding property of Kryvos, we make the following assumptions about the primitives we use (see also Section 3):

(P1) The public-key encryption scheme is IND-CCA2-secure, the SNARKs and the NIZKPs are perfectly zero-knowledge, and the commitment scheme is perfectly hiding.

(P2) An adversary  $\pi_A$  does neither corrupt the scheduler S nor the BB, and at least  $n_n^h$  voters are honest.

THEOREM 5.4 (PUBLICLY TALLY-HIDING). Let  $\mathcal{T} = \{T, ..., T_{n_t}\}$ and  $t = n_t$ . Then, assuming (P1) and (P2) hold, the voting protocol  $P_{Kryvos}(n_v, n_t, C, f_{res}, \mu)$  is  $(\delta_{(n_v, n_v^h, \mu)}^{ideal}(C, f_{res}), \delta_{(n_v, n_v^h, \mu)}^{ideal}(C, f_{complete}))$ -publicly tally-hiding w.r.t.  $(\mathcal{T}, n_t)$ .

Theorem 5.4 essentially states that the public privacy level  $\delta_p$  of Kryvos is the ideal one for  $(C, f_{\text{res}})$ , and its internal privacy level  $\delta_i$  is the ideal one for  $(C, f_{\text{complete}})$  where  $f_{\text{complete}}$  returns the number of votes for each choice/candidate. Figure 4 illustrates that for IRV Kryvos dramatically improves public privacy compared to non-tally-hiding systems (for which  $\delta_i = \delta_p \geq \delta^{\text{ideal}}(C, f_{\text{complete}})$ ).

The formal proof of Theorem 5.4 is provided in our technical report [54].

# 6 RELATED WORK

In this section, we discuss and compare Kryvos with related e-voting systems that were designed to be tally-hiding [14, 21, 28, 50, 51, 62, 84, 86] or to protect against Italian attacks [13, 25, 49, 55, 79, 86], respectively. We also briefly discuss our choice of the Groth16 SNARK in comparison with alternative proof systems.

Fully tally-hiding e-voting. The idea of fully tally-hiding evoting and the first such system was proposed by Benaloh [14]. Hevia and Kiwi [51] published a fully tally-hiding e-voting system specifically tailored to jury votings. Wen and Buckland were the first to show how the tally can be fully hidden in IRV elections [86]. However, unlike Kryvos, the aforementioned protocols [14, 51, 86] were neither formally proven secure nor were they implemented to show their practicality (e.g., the computational complexity of [86] suggests that it is actually not truly practical). Another fully tallyhiding e-voting protocol has been proposed by Szepieniec and Preneel [84], but unfortunately, their protocol is insecure. The authors discuss some mitigations but do not solve the problem (see [84], Appendix A, for details). Cortier et al. [28] have proposed fully tally-hiding MPC components for the tallying phase of a multitude of result functions and studied their asymptotic complexity but do not provide an implementation for showing their practicality. We therefore concentrate our discussion regarding fully tally-hiding systems on [21, 62] in what follows. The design of these systems is quite different to Kryvos as they are based on MPC.

Canard et al. [21] proposed a fully tally-hiding e-voting system specifically for Majority Judgement. However, there is a non-negligible chance that the system does not output a result.<sup>11</sup> They

<sup>&</sup>lt;sup>11</sup>The reason for this is due to the underlying evaluation algorithm that does not provide a result for every possible tally. Kryvos uses a different algorithm that always outputs the correct result.

implemented their system and provide benchmarks which demonstrate that their system can handle large numbers of voters, just like Kryvos. While their system needs almost 20 minutes to tally 5 candidates with 5 possible grades and up to  $2^{20} - 1$  voters, Kryvos can handle all practically relevant numbers of candidates and grades with up to  $2^{32} - 1$  voters, as shown in Figure 2. For example, 7 candidates and 6 grades are evaluated in 5.678 seconds. Hence, Kryvos shows for the first time that publicly tally-hiding systems can not only realize majority judgement but also, as initially hoped, indeed achieve much better efficiency than a fully tally hiding system (by providing weaker privacy towards talliers). We further note that, in contrast to Kryvos, the security of [21] was not formally analyzed and the implementation was not tested in a distributed network but on a single computer only. Due to the online complexity of the underlying MPC protocol employed in [21], it is not clear how [21] performs in real-world distributed tallying scenarios.

Ordinos [50, 62] is the first provably secure verifiable fully tallyhiding e-voting system. It has been implemented for a wide range of voting methods and result functions with detailed benchmarks provided in [50, 62]. The main difference between Kryvos and Ordinos is their balance between efficiency and the tally-hiding property they provide: Ordinos provides the stronger notion of *fully* tallyhiding and is practical for rather simple voting methods (single and certain multi votes), simple result functions, and only a limited number of choices (yet large numbers of voters), whereas Kryvos provides the relaxed notion of *public* tally-hiding and is practical even for very complex voting methods (e.g., Borda or Condorcet methods) and result functions (e.g., IRV) and large numbers of possible choices. This is also illustrated in Figure 1 where we compare the runtime of various result functions of Ordinos and Kryvos. Unlike for Kryvos, the result function has a strong impact on performance in Ordinos. The benchmarks of the vote threshold result function of Ordinos provide a lower bound for all other result functions that Ordinos supports. We also include benchmarks of the Condorcet Smith set result function from the extension of Ordinos [50] to illustrate that this is indeed the case.

We again note that both publicly tally-hiding and fully tallyhiding protocols offer the same level of privacy w.r.t. the public. However, fully tally-hiding protocols offer a better level of privacy w.r.t. talliers, whereas publicly tally-hiding systems can provide better efficiency, as demonstrated by Kryvos.

**Partially tally-hiding e-voting.** A number of partially tallyhiding protocols have been proposed that solve the long-standing issue of Italian attacks for complex voting methods, such as Condorcet, Borda, IRV, and STV (e.g., [13, 25, 49, 55, 79]). In fact, as far as we are aware, all existing partially tally-hiding systems focus on mitigating Italian attacks. Some of these systems are quite efficient and able to handle real world elections, such as [79], which is one of the most efficient ones and has been shown to be practical for the 2015 New South Wales IRV elections.

As already described in Section 1, partially and publicly tallyhiding systems offer different trade-offs between privacy and efficiency compared to fully tally-hiding ones. These trade-offs lead to incomparable privacy properties. Specifically, while publicly tally hiding protocols hide the full tally from the public, partially tallyhiding systems still reveal some intermediate information about the tally to the public. For example, in the voting protocol for IRV elections by Ramchen et al. [79], the public learns, among others, the order of the weakest candidates, which might embarrass those candidates. On the other hand, partially tally-hiding protocols hide parts of the tally even from internal parties, whereas publicly tally-hiding protocols reveal the full aggregated tally to the talliers. Hence, Kryvos does not protect against Italian attacks performed by (one of) the talliers. This is in contrast to the partially and fully tally-hiding systems mentioned above, which protect against Italian attacks also in this situation.

In terms of efficiency, Table 3 illustrates that Kryvos is well able to handle the same real world IRV elections as [79]. Hence, both protocols are practical solutions for IRV elections that provide different incomparable balances in terms of privacy.

**Regarding the choice of the Groth16 SNARK.** There are many other possibilities for instantiating the zero-knowledge proofs for Kryvos. However, in comparison with Groth16, many existing constructions have significantly worse benchmarks in terms of proof size and/or verification time [10, 18]. Furthermore, constructing a secure e-voting system requires care and not every proof system is directly applicable in this context. For example, [69] tries to construct a secure e-voting system based on Commit-and-Prove-SNARKs [20] but faces several challenges in breaking the link between ballot and voter. However, we emphasize that the (results for the) efficient combination of algorithms for computing Pedersen commitments will also carry over to other SNARKs as long as the base field of the elliptic curve for the Pedersen commitments is compatible with the base field of the SNARK. This is the case for Groth16, but might require some work for other SNARKs.

# 7 CONCLUSION

With Kryvos, we proposed the first provably secure verifiable evoting system which directly follows the common practice of publicly tally-hiding elections. This enabled us to come up with a radically different design which offers a new, practically relevant balance between privacy and efficiency compared to all previous approaches.

Along the way, we presented and carefully evaluated various design and implementation options for publicly tally-hiding systems, which does motivate and justify the current system.

We finally note that Kryvos is of interest also beyond its tallyhiding property since it is the first verifiable homomorphic e-voting system that supports very complex ballot formats, along with complex well-formedness conditions. By this, Kryvos also opens new options for homomorphic e-voting systems in general.

# ACKNOWLEDGEMENTS

This research was supported by the DFG through grant KU 1434/11-1, by the CRYPTECS project, which has received funding from the German BMBF through grant 16KIS1441 and from the French ANR through grant ANR-20-CYAL-0006, by the European Social Fund via IT Academy programme, by the Carl Zeiss Foundation, and by the Centre for Integrated Quantum Science and Technology (IQ<sup>ST</sup>). Johannes Müller was supported by the Luxembourg National Research Fund (FNR), under the CORE Junior project FP2 (C20/IS/14698166/FP2/Mueller).

CCS '22, November 7-11, 2022, Los Angeles, CA, USA.

## REFERENCES

- [1] ACM. 2020. Q&A on ACM's Internet Voting. https://www.acm.org/binaries/ content/assets/acmelections/acminternetvoting-1.pdf.
- [2] B. Adida. 2008. Helios: Web-based Open-Audit Voting. In USENIX 2008. 335-348.
- Ben Adida, Olivier de Marneffe, Olivier Pereira, and Jean-Jaques Quisquater. 2009. Electing a University President Using Open-Audit Voting: Analysis of Real-World Use of Helios. In USENIX/ACCURATE Electronic Voting Technology (EVT 2009).
- [4] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkitasubramaniam. 2017. Ligero: Lightweight Sublinear Arguments Without a Trusted Setup. In Proceedings of the 2017 ACM CCS. 2087–2104.
- [5] Thomas Attema, Ignacio Cascudo, Ronald Cramer, Ivan Bjerre Damgård, and Daniel Escudero. 2022. Vector Commitments over Rings and Compressed  $\Sigma$ -Protocols. Cryptology ePrint Archive (2022).
- [6] Michel Balinski and Rida Laraki. 2007. A Theory of Measuring, Electing, and Ranking. Proceedings of the National Academy of Sciences 104, 21 (2007), 8720-8725.
- [7] Michel Balinski and Rida Laraki. 2014. Judge: Don't Vote! Operations Research 62.3 (2014). 483-511.
- [8] Susan Bell, Josh Benaloh, Mike Byrne, Dana DeBeauvoir, Bryce Eakin, Gail Fischer, Philip Kortum, Neal McBurnett, Julian Montoya, Michelle Parker, Olivier Pereira, Philip Stark, Dan Wallach, and Michael Winn. 2013. STAR-Vote: A Secure, Transparent, Auditable, and Reliable Voting System. USENIX Journal of Election Technology and Systems (JETS) 1 (August 2013), 18-37.
- [9] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. 2018. Scalable, Transparent, and Post-Quantum Secure Computational Integrity. IACR Cryptology ePrint Archive 2018 (2018), 46.
- [10] Eli Ben-Sasson, Alessandro Chiesa, Michael Riabzev, Nicholas Spooner, Madars Virza, and Nicholas P. Ward. 2019. Aurora: Transparent Succinct Arguments for R1CS. In Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 11476). Springer, 103-128.
- [11] J.G. Benaloh. 1987. Verifiable Secret Ballot Elections. Ph.D. Dissertation. Yale University.
- [12] Josh Benaloh. 2007. Ballot Casting Assurance via Voter-Initiated Poll Station Auditing. In 2007 USENIX/ACCURATE Electronic Voting Technology Workshop, EVT'07, Boston, MA, USA, August 6, 2007.
- [13] Josh Benaloh, Tal Moran, Lee Naish, Kim Ramchen, and Vanessa Teague. 2009. Shuffle-Sum: Coercion-Resistant Verifiable Tallying for STV Voting. IEEE Trans. Information Forensics and Security 4, 4 (2009), 685-698.
- [14] J. D. Benaloh. 1986. Improving Privacy in Cryptographic Elections (Technical Report). Technical Report. Technical report.
- [15] D. Bernhard, V. Cortier, D. Galindo, O. Pereira, and B. Warinschi. 2015. SoK: A Comprehensive Analysis of Game-Based Ballot Privacy Definitions. In IEEE S&P 2015 499-516
- [16] Jonathan Bootle and Jens Groth. 2018. Efficient Batch Zero-Knowledge Arguments for Low Degree Polynomials. In Public-Key Cryptography - PKC 2018 - 21st IACR International Conference on Practice and Theory of Public-Key Cryptography, Rio de Janeiro, Brazil, March 25-29, 2018, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 10770). Springer, 561-588.
- [17] Xavier Boyen, Thomas Haines, and Johannes Müller. 2020. A Verifiable and Practical Lattice-Based Decryption Mix Net with External Auditing. In Computer Security - 25th European Symposium on Research in Computer Security, ESORICS 2020.
- [18] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Gregory Maxwell. 2018. Bulletproofs: Short Proofs for Confidential Transactions and More. In 2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA. IEEE Computer Society, 315-334.
- [19] Jan Camenisch and Victor Shoup. 2003. Practical Verifiable Encryption and Decryption of Discrete Logarithms. In CRYPTO 2003, Proceedings (LNCS, Vol. 2729). Springer, 126-144.
- [20] Matteo Campanelli, Dario Fiore, and Anaïs Querol. 2019. LegoSNARK: Modular Design and Composition of Succinct Zero-Knowledge Proofs. In Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS 2019, London, UK, November 11-15, 2019. ACM, 2075-2092.
- [21] Sébastien Canard, David Pointcheval, Quentin Santos, and Jacques Traoré. 2018. Practical Strategy-Resistant Privacy-Preserving Elections. In ESORICS 2018. 331-
- [22] D. Chaum, R. Carback, J. Clark, A. Essex, S. Popoveniuc, R. L. Rivest, P. Y. A. Ryan, E. Shen, and A. T. Sherman. 2008. Scantegrity II: End-to-End Verifiability for Optical Scan Election Systems using Invisible Ink Confirmation Codes. In USENIX/ACCURATE Electronic Voting Technology (EVT 2008). USENIX Association.
- [23] Alessandro Chiesa, Dev Ojha, and Nicholas Spooner. 2020. Fractal: Post-quantum and Transparent Recursive Proofs from Holography. In EUROCRYPT 2020, Proceedings, Part I (LNCS, Vol. 12105). Springer, 769–793. M. R. Clarkson, S. Chong, and A. C. Myers. 2008. Civitas: Toward a Secure Voting
- [24] System. In S&P 2008. 354-368.

- [25] Michael R. Clarkson and Andrew C. Myers. 2005. Coercion-Resistant Remote Voting Using Decryption Mixes. In In Frontiers in Electronic Elections (FEE 2005).
- [26] V. Cortier, D. Galindo, S. Glondu, and M. Izabachène. 2014. Election Verifiability for Helios under Weaker Trust Assumptions. In ESORICS 2014. 327-344
- V. Cortier, D. Galindo, R. Küsters, J. Müller, and T. Truderung. 2016. SoK: Verifia-[27] bility Notions for E-Voting Protocols. In S&P 2016. 779-798.
- [28] Véronique Cortier, Pierrick Gaudry, and Quentin Yang. 2021. A toolbox for verifiable tally-hiding e-voting systems. IACR Cryptol. ePrint Arch. 2021 (2021), 491.
- [29] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. 1994. Proofs of Partial Knowledge and Simplified Design of Witness Hiding Protocols. In Advances in Cryptology - CRYPTO '94, Proceedings. 174-187.
- Ronald Cramer, Matthew K. Franklin, Berry Schoenmakers, and Moti Yung. 1996. [30] Multi-Authority Secret-Ballot Elections with Linear Work. In EUROCRYPT 1996. 72-83.
- [31] Ronald Cramer and Victor Shoup. 1998. A Practical Public Key Cryptosystem Provably Secure Against Adaptive Chosen Ciphertext Attack. In Advances in Cryptology - CRYPTO '98, 18th Annual International Cryptology Conference, Proceedings (Lecture Notes in Computer Science, Vol. 1462), Hugo Krawczyk (Ed.). Springer, 13-25.
- [32] CrossRef. 2019. Election process and results. https://www.crossref.org/boardand-governance/elections/
- [33] Chris Culnane and Steve A. Schneider. 2014. A Peered Bulletin Board for Robust Use in Verifiable Voting Systems. In IEEE CSF 2014. 169-183.
- [34] Debian Project. 2012. Ubuntu IRC Council Position. https://www.debian.org/ vote/.
- [35] Deutsche Forschungsgemeinschaft (DFG). 2019. DFG Fachkollegienwahl 2019. https://www.dfg.de/download/pdf/dfg\_im\_profil/gremien/fachkollegien/ fk-wahl2019/fkwahl\_2019\_wahlergebnis\_endgueltig\_200218.pdf. Electoral Commission NSW. 2015. NSW State Electoin
- NSW State Electoin Results 2015. [36] https://pastvtr.elections.nsw.gov.au/SGE2015/la-home.htm.
- [37] Jeremy Epstein. 2015. Weakness in Depth: A Voting Machine's Demise. IEEE Secur. Priv. 13, 3 (2015), 55-58.
- European Broadcasting Union. 2020. Eurovision Song Contest How it works. [38] https://eurovision.tv/about/how-it-works
- Boyuan Feng, Lianke Qin, Zhenfei Zhang, Yufei Ding, and Shumo Chu. 2021. ZEN: [39] Efficient Zero-Knowledge Proofs for Neural Networks. Technical Report 2021/87. Cryptology ePrint Archive.
- [40] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. 2019. PLONK: Permutations over Lagrange-bases for Oecumenical Noninteractive arguments of Knowledge. IACR Cryptol. ePrint Arch. 2019 (2019), 953.
- [41] David Galindo, Sandra Guasch, and Jordi Puiggali. 2015. 2015 Neuchâtel's Castas-Intended Verification Mechanism. In VoteID 2015, Proceedings. 3-18.
- Gesellschaft für Informatik (GI). 2019. GI Wahlen. https://gi.de/wahlen. [42]
- [43] Irene Giacomelli, Jesper Madsen, and Claudio Orlandi. 2016. ZKBoo: Faster Zero-Knowledge for Boolean Circuits. In 25th USENIX Security Symposium, 2016. USENIX Association, 1069-1083.
- Government of India. 2020. Constitution of India. https://www.india.gov.in/my-[44] government/constitution-india
- Jens Groth. 2016. On the Size of Pairing-Based Non-interactive Arguments. In [45] EUROCRYPT 2016, Proceedings, Part II (LNCS, Vol. 9666). Springer, 305-326.
- [46] Thomas Haines, Sarah Jamie Lewis, Olivier Pereira, and Vanessa Teague. 2020. How Not to Prove Your Election Outcome. In 2020 IEEE SP 2020. IEEE, 644-660.
- [47] Thomas Haines and Johannes Müller. 2020. SoK: Techniques for Verifiable Mix Nets. In IEEE 33rd Computer Security Foundations Symposium, CSF, 2020. IEEE Computer Society.
- [48] Thomas Haines, Dirk Pattinson, and Mukesh Tiwari. 2019. Verifiable Homomorphic Tallying for the Schulze Vote Counting Scheme. In Verified Software. Theories, Tools, and Experiments - 11th International Conference, VSTTE 2019, New York City, NY, USA, July 13-14, 2019, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 12031). Springer, 36-53.
- [49] James Heather. 2007. Implementing STV securely in Prêt à Voter. In IEEE CSF 2007. 157-169
- [50] Fabian Hertel, Nicolas Huber, Jonas Kittelberger, Ralf Küsters, Julian Liedtke, and Daniel Rausch. 2021. Extending the Tally-Hiding Ordinos System: Implementations for Borda, Hare-Niemeyer, Condorcet, and Instant-Runoff Voting. Technical Report 2021/1420. Cryptology ePrint Archive.
- [51] Alejandro Hevia and Marcos A. Kiwi. 2002. Electronic Jury Voting Protocols. In LATIN 2002, Proceedings. 415-429.
- [52] Lucca Hirschi, Lara Schmid, and David A. Basin. 2020. Fixing the Achilles Heel of E-Voting: The Bulletin Board. IACR Cryptol. ePrint Arch. 2020 (2020), 109.
- [53] Nicolas Huber, Ralf Küsters, Toomas Krips, Julian Liedtke, Johannes Müller, Daniel Rausch, Pascal Reisert, and Andreas Vogt. 2022. Implementation of Kryvos. https://github.com/JulianLiedtke/kryvos.
- [54] Nicolas Huber, Ralf Küsters, Toomas Krips, Julian Liedtke, Johannes Müller, Daniel Rausch, Pascal Reisert, and Andreas Vogt. 2022. Kryvos: Publicly Tally-Hiding Verifiable E-Voting. Cryptology ePrint Archive 2022/1132 (2022).

Kryvos: Publicly Tally-Hiding Verifiable E-Voting

- [55] Wojciech Jamroga, Peter B. Rønne, Peter Y. A. Ryan, and Philip B. Stark. 2019. Risk-Limiting Tallies. In E-Vote-ID 2019, Proceedings (LNCS, Vol. 11759). Springer, 183–199.
- [56] A. Juels, D. Catalano, and M. Jakobsson. 2005. Coercion-Resistant Electronic Elections. In Proceedings of Workshop on Privacy in the Eletronic Society (WPES 2005). ACM Press, 61–70.
- [57] Sanket Kanjalkar, Ye Zhang, Shreyas Gandlur, and Andrew Miller. 2021. Publicly Auditable MPC-as-a-Service with succinct verification and universal setup. In IEEE European Symposium on Security and Privacy Workshops, EuroS&P 2021, Vienna, Austria, September 6-10, 2021. IEEE, 386–411.
- [58] Aggelos Kiayias, Annabell Kuldmaa, Helger Lipmaa, Janno Siim, and Thomas Zacharias. 2018. On the Security Properties of e-Voting Bulletin Boards. In SCN 2018, Proceedings. 505–523.
- [59] Aggelos Kiayias, Thomas Zacharias, and Bingsheng Zhang. 2015. DEMOS-2: Scalable E2E Verifiable Elections without Random Oracles. In CCS 2015. 352–363.
- [60] Aggelos Kiayias, Thomas Zacharias, and Bingsheng Zhang. 2015. End-to-End Verifiable Elections in the Standard Model. In EUROCRYPT 2015. 468–498.
- [61] Ahmed Kosba, Zhichao Zhao, Andrew Miller, Yi Qian, Hubert Chan, Charalampos Papamanthou, Rafael Pass, Elaine Shi, et al. 2015. CØCØ: A Framework for Building Composable Zero-Knowledge Proofs. Cryptology ePrint Archive (2015).
- [62] Ralf Küsters, Julian Liedtke, Johannes Müller, Daniel Rausch, and Andreas Vogt. 2020. Ordinos: A Verifiable Tally-Hiding Remote E-Voting System. In IEEE EuroS&P 2020.
- [63] R. Küsters, J. Müller, E. Scapin, and T. Truderung. 2016. sElect: A Lightweight Verifiable Remote Voting System. In CSF 2016. 341–354.
- [64] R. Küsters and T. Truderung. 2016. Security Analysis of Re-Encryption RPC Mix Nets. In IEEE EuroS&P 2016. IEEE Computer Society, 227–242.
- [65] R. Küsters, T. Truderung, and A. Vogt. 2010. Accountability: Definition and Relationship to Verifiability. In ACM CCS 2010. 526–535.
- [66] R. Küsters, T. Truderung, and A. Vogt. 2011. Verifiability, Privacy, and Coercion-Resistance: New Insights from a Case Study. In IEEE S&P 2011. 538–553.
- [67] R. Küsters, T. Truderung, and A. Vogt. 2012. Clash Attacks on the Verifiability of E-Voting Systems. In IEEE S&P 2012. 395–409.
- [68] R. Küsters, T. Truderung, and A. Vogt. 2014. Formal Analysis of Chaumian Mix Nets with Randomized Partial Checking. In S&P 2014. 343–358.
- [69] Jiwon Lee, Jaekyoung Choi, Jihye Kim, and Hyunok Oh. 2019. SAVER: Snarkfriendly, Additively-homomorphic, and Verifiable Encryption and decryption with Rerandomization. *IACR Cryptol. ePrint Arch.* 2019 (2019), 1270.
- [70] Maine State Legislature. 2020. Ranked Choice Voting in Maine. http://legislature. maine.gov/lawlibrary/ranked-choice-voting-in-maine/9509.
- [71] Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. 2019. Sonic: Zero-Knowledge SNARKs from Linear-Size Universal and Updatable Structured Reference Strings. In Proceedings of the 2019 ACM CCS. 2111–2128.

- [72] David Mesten, Johannes Müller, and Pascal Reisert. 2022. To appear. How Efficient are Replay Attacks against Vote Privacy? A Formal Quantitative Analysis. In IEEE 35rd Computer Security Foundations Symposium, CSF, 2022.
- [73] NSW Government. 2020. Constitution Act No 32. https://legislation.nsw.gov.au/ ~/view/act/1902/32.
- [74] Katsuyuki Okeya and Kouichi Sakurai. 2001. Efficient Elliptic Curve Cryptosystems from a Scalar Multiplication Algorithm with Recovery of the y-Coordinate on a Montgomery-Form Elliptic Curve. In Cryptographic Hardware and Embedded Systems - CHES 2001, Third International Workshop, Paris, France, May 14-16, 2001, Proceedings (Lecture Notes in Computer Science, Vol. 2162). Springer, 126–141.
- [75] Alex Ozdemir and Dan Boneh. 2021. Experimenting with Collaborative zk-SNARKs: Zero-Knowledge Proofs for Distributed Secrets. IACR Cryptol. ePrint Arch. (2021), 1530.
- [76] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. 2013. Pinocchio: Nearly Practical Verifiable Computation. In 2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013. IEEE Computer Society, 238–252.
- [77] Torben P. Pedersen. 1991. Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing. In Proceedings of the 11th Annual International Cryptology Conference (CRYPTO 1991) (Lecture Notes in Computer Science, Vol. 576). Springer, 129–140.
- [78] Personal communication (email) with Philip Wright, Technical Director of CES. 2020.
- [79] Kim Ramchen, Chris Culnane, Olivier Pereira, and Vanessa Teague. 2019. Universally Verifiable MPC and IRV Ballot Counting. In Financial Cryptography and Data Security - FC 2019, Revised Selected Papers (LNCS, Vol. 11598). Springer, 301–319.
- [80] Republic of Nauru. 2016. Electoral Act No. 15. http://ronlaw.gov.nr/nauru\_lpms/ files/acts/d83250a1ebdc56c1701fa7aa245af5b1.pdf.
- [81] scipr-lab. 2017. libsnark. https://github.com/scipr-lab/libsnark.
- [82] Society for Industrial and Applied Mathematics (SIAM). 2019. SIAM Announces New 2020 Leadership. https://sinews.siam.org/Details-Page/siam-announcesnew-2020-leadership-1.
- [83] D. Springall, T. Finkenauer, Z. Durumeric, J. Kitcat, H. Hursti, M. MacAlpine, and J. A. Halderman. 2014. Security Analysis of the Estonian Internet Voting System. In Proceedings of the 2014 ACM CCS. 703–715.
- [84] Alan Szepieniec and Bart Preneel. 2015. New Techniques for Electronic Voting. USENIX Journal of Election Technology and Systems (JETS) 3, 2 (2015), 46 – 69.
- [85] The National Archives. 2011. Greater London Authority Act 1999. https://www. legislation.gov.uk/ukpga/1999/29/contents.
- [86] Roland Wen and Richard Buckland. 2009. Minimum Disclosure Counting for the Alternative Vote. In VoteID 2009. Proceedings (LNCS, Vol. 5767). Springer, 122–140.
   [87] S. Wolchok, E. Wustrow, I. A. Halderman, H. K. Prasad, A. Kankipati, S. K.
- [87] S. Wolchok, E. Wustrow, J. A. Halderman, H. K. Prasad, A. Kankipati, S. K. Sakhamuri, V. Yagati, and R. Gonggrijp. 2010. Security Analysis of India's electronic Voting Machines. In *Proceedings of the 17th ACM CCS*. 1–14.