Fully Tally-Hiding Verifiable E-Voting for Real-World Elections with Seat-Allocations

Carmen Wabartha[*2], Julian Liedtke[×2[0000−0002−8289−4970]], Nicolas Huber[×1,2[0000−0001−6905−3571]], Daniel Rausch[×2[0000−0002−1901−3659]], and Ralf Küsters[×2[0000−0002−9071−9312]]

[1] Corresponding Author
[2] University of Stuttgart
[*] st161329@stud.uni-stuttgart.de
[×] firstname.secondname@sec.uni-stuttgart.de

Abstract. Modern e-voting systems provide what is called verifiability, i.e., voters are able to check that their votes have actually been counted despite potentially malicious servers and voting authorities. Some of these systems, called tally-hiding systems, provide increased privacy by revealing only the actual election result, e.g., the winner of the election, but no further information that is supposed to be kept secret. However, due to these very strong privacy guarantees, supporting complex voting methods at a real-world scale has proven to be very challenging for tally-hiding systems.

A widespread class of elections, and at the same time, one of the most involved ones is parliamentary election with party-based seat-allocation. These elections are performed for millions of voters, dozens of parties, and hundreds of individual candidates competing for seats; they also use very sophisticated multi-step algorithms to compute the final assignment of seats to candidates based on, e.g., party lists, hundreds of electoral constituencies, possibly additional votes for individual candidates, overhang seats, and special exceptions for minorities. So far, it has not been investigated whether and in how far such elections can be performed in a verifiable tally-hiding manner.

In this work, we design and implement the first verifiable (fully) tally-hiding e-voting system for an election from this class, namely, for the German parliament (Bundestag). As part of this effort, we propose several new tally-hiding building blocks that are of independent interest. We perform benchmarks based on actual election data, which show, perhaps surprisingly, that our proposed system is practical even at a real-world scale. Our work thus serves as a foundational feasibility study for this class of elections.

1 Introduction

E-voting is of rising interest. In order to ensure secure and correct elections, modern e-voting systems are designed to be (end-to-end) verifiable [1–3,6–8,16–18,20,24], that is, voters should be able to check that their votes/ballots were

submitted correctly, and voters, election officials, and even external observers should be able to check that the election result corresponds to the votes that were cast. A stronger notion of verifiability is accountability, which states that, if the result turns out to be incorrect, then a misbehaving party causing this mistake can be identified and be held accountable. A very common method for election systems to achieve verifiability is by publishing the full tally, which consists of all (potentially aggregated) individual votes, along with additional evidence, such as zero-knowledge proofs (ZKPs), which proves that the tally was computed correctly. With the knowledge of the full tally, everyone is able to compute the actual election result, e.g., the winner of the election, and check whether this corresponds to the claimed election result.

More recently, verifiable tally-hiding e-voting systems (e.g. [4, 5, 9, 12, 14, 15, 19, 23, 27]), have been proposed that defer from revealing the full tally. They are rather designed to only publish the actual election result, e.g., the winning candidate(s) of an election, and as little further information as possible (ideally none), while the correctness of the election result can still be verified. Following the terminology of [15], tally-hiding systems can be divided into three classes: Fully tally-hiding systems (e.g., [5, 9, 14, 19]) are the strongest ones as they reveal only the election result. Publicly or partially tally-hiding systems (see, e.g., [15, 23]) are more relaxed in that they reveal some information beyond the election result, possibly only to certain parties. As discussed for example in [9, 14, 15, 19], tally-hiding systems offer several attractive features such as improved ballot privacy for voters, avoiding embarrassment or weakening of candidates, protection against a specific class of coercion attacks called Italian attacks [4, 13], and preventing Gerrymandering. So far, it has been shown that simple election schemes can be performed at a large scale, even in a fully tally-hiding manner. However, due to the strong privacy requirements, more complex voting methods have proven to be a challenge for all types of tally-hiding systems, with some types of elections even turning out to be practical only for very few candidates and/or voters (cf. Section 6).

A very important class of elections in practice is *parliamentary election with party-based seat allocation* as carried out by many countries around the world. These are among the most complex types of elections: They usually involve millions of voters, dozens of parties, hundreds of individual candidates, and hundreds of electoral constituencies. In some cases, voters have not just one but multiple votes that they can distribute among parties and possibly also individual candidates. Sophisticated multi-step algorithms are used to compute the election result, i.e., the assignment of seats to individual candidates. An important component for this process is a so-called *seat allocation method*, which takes as input a number of available seats and a set of parties with their total number of votes and then computes the number of seats assigned to each party. While a crucial part this seat allocation method is only a small step in the computation of the actual election result. Additional steps are taken, e.g., to combine the results of different constituencies to distribute seats that are directly allocated to individual candidates instead of just parties, to take into account minimum vote

counts for parties before they are assigned any seats, and to include special exceptions for minorities. Furthermore, the seats assigned to each party need to be mapped to individual candidates, typically according to party candidate lists for each constituency and weighted by how many votes a party has obtained in the respective constituency. In some cases, even the size of the parliament is modified while computing the election result, possibly only after the seat allocation method has already been computed to more closely reflect the vote distribution.

Perhaps due to this intimidating complexity, so far, it has not been investigated *whether and in how far this class of elections can be performed in a tally-hiding manner, and whether this is possible even at the same scale in terms of voters, parties, candidates, and constituencies as needed in real-world elections*. There are only a few existing works that propose tally-hiding algorithms for computing certain seat allocation methods, namely, the d'Hondt method [9] and the Hare-Niemeyer method [14]. As explained above, while seat allocation methods are important components, they constitute just a small portion of the entire election scheme, and hence, these prior works do not answer the above question. In this work, we therefore, for the first time, investigate this open research question.

Contributions. More specifically, we design, implement, and benchmark the *first verifiable (and even accountable) fully tally-hiding e-voting system for a major real-world party-based parliamentary election*, namely, the election of the German parliament (Bundestag). Perhaps surprisingly, and as our main insight, with this system, we are able to *show that such a parliamentary election scheme with party-based seat allocation can actually be performed in a verifiable, fully tally-hiding manner at a real-world scale*. Our system supports the strongest level of tally-hiding, namely full tally-hiding. That is, if desired, one can reveal only the allocation of individual candidates to seats and the number of voters who cast a vote and nothing else to anybody. But one can also easily relax the kind of information that is revealed, e.g., by additionally publishing the winners of individual constituencies.

On a technical level, to obtain our voting system, we follow and slightly modify a generic approach for constructing verifiable fully tally-hiding systems, namely the Ordinos framework [19]. The Ordinos framework provides a general blueprint for the structure of such systems. Some components in this blueprint are unspecified and have to be filled in by protocol designers on a case-by-case basis to obtain a concrete instantiation of Ordinos that can perform an election for a specific voting method. It has been shown in [19] that, as long as those components meet specific requirements, the overall system/instantiation is a secure verifiable, fully tally-hiding e-voting system. The main challenge lies in constructing those components for a specific voting method in such a way that they provide all expected security properties while achieving practical performance.

The most important and also most difficult to design component is a publicly verifiable secure multi-party computation (MPC) protocol that computes the election result for the German parliament from the set of (encrypted) ballots. Due to the inherent complexity and scale of this election, this requires special

care to obtain not just a theoretically secure but also a practically efficient system. Specifically, we first propose several MPC building blocks, including the first MPC subroutine for computing the Sainte-Laguë seat allocation method used for parliamentary elections, not just in Germany but also in, e.g., Indonesia, New Zealand, Nepal, Sweden, Norway, and Kosovo. Based on these building blocks, we then construct an efficient MPC protocol that performs the entire election evaluation for the German parliament. Along the way, we evaluate different options for designing our algorithms and propose several novel optimizations to improve the overall efficiency. We note that many of our ideas and building blocks, such as our MPC protocol for the Sainte-Laguë method, are of interest also for other parliamentary election schemes since such elections often use similar concepts and components.

The overall practicality of e-voting systems following the Ordinos approach is determined essentially only by the performance of the MPC component. Hence, to evaluate the performance and identify potential limitations of our system, we have implemented our full MPC protocol for electing the German parliament and performed extensive benchmarks based on actual real-world election data consisting of the votes of all respective constituencies. Our solution needs about a day to compute the election result, which is within the usual time frame expected for this election, thus demonstrating that our MPC protocol is practical even for such a complex large-scale political election.

Altogether, our results serve as a foundational feasibility study for (fully) tally-hiding elections for the important class of parliamentary elections with party-based seat allocation. Of course, as can be seen in countries already using or aiming at online elections, establishing and actually deploying a full-fledged ready-to-use system in the real world requires a huge effort beyond studying feasibility. Future deployments can build on our results by considering further aspects of parliamentary elections that are out of scope of this work, such as deciding which parties run the election in a distributed fashion, tackling the risk of voter coercion, establishing procedures for handling voter complaints, etc.

Structure. We recall the Ordinos framework in Section 2. In Section 3, we present novel building blocks that we have constructed to realize the voting methods in this work. In Section 4, we present the Sainte-Laguë method, including a novel variant to compute the Sainte-Laguë seat allocation and different tally-hiding algorithms to compute the Sainte-Laguë method. We present our voting system for the German Bundestag in Section 5. We discuss related work and conclude in Section 6. Our implementation is available at [26]. A full version of this paper with complete details for all of our results is available at [25].

2 Preliminaries

Notation. We write $[n]$ to denote the set $\{0, \ldots, n-1\}$. Let n_{cand} be the number of candidates/parties/choices, and let n_{votes} be the (maximal) number of votes. We will use n_{seats} to denote the number of seats that are being distributed among n_{parties} parties. With n^j_{votes} we denote the number of votes, and with

n^j_{seats} the number of seats that party j has received. The format of a plain, i.e., unencrypted, ballot is defined via a finite *choice space* $C \subseteq \mathbb{N}^{n_{\mathsf{cand}}}$, i.e., a ballot assigns each candidate a number subject to constraints defined by C. For example, a single vote election where a plain ballot contains one vote for a single candidate/party/choice can be modeled via the choice space $C_{\mathsf{single}} := \{(b_0, \ldots, b_{n_{\mathsf{cand}}-1}) \in \{0,1\}^{n_{\mathsf{cand}}} \mid \sum_i b_i = 1\}$. For voter j we denote her plain ballot by $v^j := (v^j_i)_{i \in [n_{\mathsf{cand}}]} \in C$.

The Ordinos Framework. The Ordinos framework was introduced in [19] as a general blueprint for constructing verifiable, fully tally-hiding e-voting systems. Systems following the Ordinos approach use a voting authority, an arbitrary number of voters, n_t trustees, an authentication server, and an append-only bulletin board (B) and roughly work as follows. In an initial *setup phase*, parameters of the election are generated and published on B, including a public key and corresponding secret key shares for an additively homomorphic t-out-of-n_t threshold public key encryption scheme $\mathcal{E} = (E, D)$. Each trustee has one secret key share and publishes a non-interactive zero-knowledge proof of knowledge (NIZKP) $\pi^{\mathsf{KeyShareGen}}$ to prove knowledge of their key share. The choice space C and the result function f_{res} of the election are published on B as well, where f_{res} takes as input a tally and outputs the corresponding election result, e.g., the candidate with the most votes. In the following *voting phase*, the voters first encrypt their ballots and then publish them on B, authenticating themselves as eligible voters with the help of the authentication server and the authentication server adds a signature to the ballot. An encrypted ballot of voter j has the form $(E_{\mathsf{pk}}(v^j_i))_{i \in [n_{\mathsf{cand}}]}$, i.e., each component of the plain ballot is encrypted separately. The encrypted ballot comes with a NIZKP π^{Enc} that proves validity of the plain ballot, i.e., $v^j = (v^j_i)_{i \in [n_{\mathsf{cand}}]} \in C$. The published encrypted ballots can be homomorphically (and publicly) aggregated to obtain an encryption of the aggregated full tally, i.e., one obtains one ciphertext on each $v_i := \sum_{j \in [n_{\text{votes}}]} v^j_i$, where v_i is the total number of votes/points that candidate/choice i obtained in the election. In the *tallying phase*, the trustees run a publicly verifiable MPC protocol $\mathsf{P}_{\mathsf{MPC}}$ to compute f_{res}. This protocol takes as (secret) inputs the secret key shares of the trustees and the (public) encrypted aggregated tally and outputs the election result $res = f_{\mathsf{res}}(v_0, \ldots, v_{n_{\mathsf{cand}}-1})$. This result, along with any material that is needed to allow external parties to verify the MPC computation, is published by the trustees on B. Finally, in the *verification phase*, voters can check that their ballots appear on B, and everyone can verify that the election result *res* was computed correctly from the encrypted ballots by re-computing the homomorphic aggregation, checking all NIZKPs, and checking the MPC computation (which typically involves additional NIZKP verifications).

Many of the above components are not fixed by the Ordinos framework because they strongly depend on the specific election method that is to be implemented. Specifically, the following parameters and components have to be specified or constructed by a protocol designer to create an instantiation of Ordinos for a concrete election method: *(i)* the choice space C and election result function f_{res}, *(ii)* a threshold encryption scheme $\mathcal{E}$, *(iii)* NIZKPs $\pi^{\mathsf{KeyShareGen}}$

and π^{Enc}, *(iv)* a EUF-CMA-secure signature scheme $\mathcal{S}$, and *(v)* an MPC protocol $\mathsf{P}_{\mathsf{MPC}}$ for computing the election result function f_{res}.

Voting systems following the Ordinos approach are intended to provide verifiability and full tally-hiding. As already mentioned, verifiability intuitively means that everyone can check whether the election result returned by the voting system corresponds to the actual votes. Full tally-hiding intuitively means that no one, including attackers, learns anything besides the number of submitted ballots and the final election result; this property, therefore, also implies the security notion of ballot privacy. We refer interested readers to [19] for formal definitions of both verifiability and full tally-hiding. Küsters et al. [19] have shown that if the above components defined by protocol designers meet certain properties, then the resulting Ordinos instance, indeed achieves both security notions:

Theorem 1 (Verifiability and Full Tally Hiding [19], informal). *Let $\mathcal{E}$ be an additively homomorphic threshold public-key encryption scheme $\mathcal{E}$, $\pi^{\mathsf{KeyShareGen}}$ and π^{Enc} be secure NIZKPs for $\mathcal{E}$, $\mathcal{S}$ be an EUF-CMA-secure signature scheme, and $\mathsf{P}_{\mathsf{MPC}}$ be a publicly verifiable secure MPC protocol for computing f_{res}, i.e., if the result does not correspond to the input, then this can be detected, and at least one misbehaving trustee can be identified; this must hold even if all trustees running the MPC protocol are malicious. Then, the resulting instance of Ordinos is verifiable and fully tally-hiding.*[3]

Existing building blocks. In this work, we will use a threshold variant of the Paillier encryption scheme [10] to implement $\mathcal{E}$. Given a public key pk for this encryption scheme, there exist publicly verifiable MPC building blocks [10,14,22] that allow the owners of the corresponding secret key shares to compute the following basic operations for $a, b, c \in \mathbb{Z}_n$ (all operations are $\mod n$ where n is determined by pk):

- $E_{\mathsf{pk}}(c) = f_{\mathsf{add}}(E_{\mathsf{pk}}(a), E_{\mathsf{pk}}(b))$ s.t. $c = a + b$; for brevity, we denote this operation by $\oplus$.
- $E_{\mathsf{pk}}(c) = f_{\mathsf{mul}}(E_{\mathsf{pk}}(a), E_{\mathsf{pk}}(b))$ s.t. $c = a \cdot b$, for brevity, we denote this operation by $\odot$.
- $E_{\mathsf{pk}}(c) = f_{\mathsf{gt}}(E_{\mathsf{pk}}(a), E_{\mathsf{pk}}(b))$ s.t. $c = 1$ iff $a \geq b$ and 0 otherwise.
- $E_{\mathsf{pk}}(c) = f_{\mathsf{eq}}(E_{\mathsf{pk}}(a), E_{\mathsf{pk}}(b))$ s.t. $c = 1$ iff $a = b$ and 0 otherwise.
- $(E_{\mathsf{pk}}(s_i))_{i=1}^n = f_{\mathsf{max}}((E_{\mathsf{pk}}(v_i))_{i=1}^n)$ s.t. $s_i \in \{0, 1\}$ and $s_i = 1$ means that $v_i = \max_{j \in \{1,\ldots,n\}} v_j \wedge \forall j \in \{i+1, \ldots, n\} : v_j < v_i$.
- $c = f_{\mathsf{dec}}(E_{\mathsf{pk}}(c))$ s.t. $E_{\mathsf{pk}}(c)$ is an encryption of c.

The MPC building blocks for computing the above operations have a useful property, namely, encrypted outputs from one building block can be used as inputs for another building block such that the resulting combined protocol is still a secure, publicly verifiable MPC protocol [22]. In other words, they allow for building more complex protocols such as $\mathsf{P}_{\mathsf{MPC}}$ for Ordinos that meet the

[3] Full tally-hiding requires that at most $t - 1$ trustees are malicious, verifiability does not require any honest trustees at all. This theorem uses further standard e-voting assumptions, such as honesty of B. We refer interested readers to [19] for full details.

requirements of Theorem 1. We further note that the MPC building blocks for computing $f_{\mathsf{gt}}()$ and $f_{\mathsf{eq}}()$ proposed by [22] offer sublinear runtime as long as an upper bound $< n$ for both input values a and b is known; hence, performance drastically increases as long as a, b are known to remain small. This, in turn, also improves performance of MPC protocols based on these two building blocks, including the MPC building block for computing $f_{\mathsf{max}}()$ [14].

3 New MPC Building Blocks

In this section, based on the primitives introduced in Section 2, we describe several new publicly verifiable MPC building blocks that we need for constructing our $\mathsf{P}_{\mathsf{MPC}}$, with full details provided in [25]. We note that these building blocks are of independent interest.

Election methods for parliamentary elections often make use of divisions that produce fractions, which is an issue for encryption schemes and MPC protocols which operate on natural numbers, such as those from Section 2. One common approach [9, 14] to deal with this is to multiply all values by the least common multiple of all divisors used in a computation such that divisions are guaranteed to always produce natural numbers. This can drastically increase the size of numbers, which in turn severely reduces the efficiency gain of the sublinear comparisons protocols $f_{\mathsf{gt}}(), f_{\mathsf{eq}}()$ from Section 2. Therefore, we instead take an alternative approach to deal with fractions by representing our values, where needed, as rational numbers consisting of a numerator n and denominator d. Encrypted rational numbers are denoted as $E^{\text{frac}}_{\mathsf{pk}}(a) := (E_{\mathsf{pk}}(a.n), E_{\mathsf{pk}}(a.d))$ where $a.n$ is the numerator and $a.d$ the denominator of a. We denote by $\textsc{Fraction}(n, d)$ the operation that creates an encrypted rational number with numerator n and denominator d (if the inputs n and/or d are not already encrypted, then they are first encrypted with public constant randomness). Based on this representation, we design and implement MPC components for basic arithmetic computations on encrypted rational numbers, including addition, multiplication, and comparisons.

Based on $f_{\mathsf{max}}()$ (see Section 2), we propose the method `getMaxFraction` that takes a list of k encrypted fractions and returns another list of the same length with $E_{\mathsf{pk}}(1)$ at the index of the maximal fraction and $E_{\mathsf{pk}}(0)$ everywhere else, where if there are multiple maxima, only the last one in the list is marked $E_{\mathsf{pk}}(1)$.

Election methods often need to deal with breaking ties. For this purpose, Cortier et al. [9] proposed an algorithm that finds the maximum in a list and additionally takes care of tie-breaking by scaling values and adding small tie-breaking values. While this scaling idea is conceptually simple, care must be taken to obtain a correct implementation. As we discuss in the full version [25], we found cases where directly applying the tie-breaking mechanism described in [9] in our setting, where fractions are represened by their numerator and denominator, which leads to an incorrect output. We address this problem in our implementation `getMaxFractionByRank` shown in Figure 1. This algorithm ad-

procedure GETMAXFRACTIONBYRANK($l = (E^{\text{frac}}_{\mathsf{pk}}(v_i))_{i=1}^{k}, k, r = (E_{\mathsf{pk}}(r_i))_{i=1}^{k}$)
  $E^{\text{frac}}_{\mathsf{pk}}(\text{c\_max\_val}) = E^{\text{frac}}_{\mathsf{pk}}(v_1)$
  $E_{\mathsf{pk}}(\text{c\_max\_idx}) = E_{\mathsf{pk}}(1)$
  $E_{\mathsf{pk}}(\text{c\_max\_r}) = E_{\mathsf{pk}}(r_1)$
  for $i = 2, \ldots, k$ **do**
    $E_{\mathsf{pk}}(m_{\max}) = E^{\text{frac}}_{\mathsf{pk}}(\text{c\_max\_val}) \odot E_{\mathsf{pk}}(\text{c\_max\_val}.d) \odot E_{\mathsf{pk}}(v_i.d) \odot k \oplus E_{\mathsf{pk}}(\text{c\_max\_r})$
    $E_{\mathsf{pk}}(m_i) = E^{\text{frac}}_{\mathsf{pk}}(v_i) \odot E_{\mathsf{pk}}(v_i.d) \odot E_{\mathsf{pk}}(\text{c\_max\_val}.d) \odot k \oplus E_{\mathsf{pk}}(r_i)$
    $E_{\mathsf{pk}}(\text{set}) = f_{\mathsf{gt}}(E_{\mathsf{pk}}(m_i), E_{\mathsf{pk}}(m_{\max}))$
    $E_{\mathsf{pk}}(\text{c\_max\_val}) = E_{\mathsf{pk}}(\text{set}) \odot E_{\mathsf{pk}}(v_i) \oplus (1 - E_{\mathsf{pk}}(\text{set})) \odot E_{\mathsf{pk}}(\text{c\_max\_val})$
    $E_{\mathsf{pk}}(\text{c\_max\_idx}) = E_{\mathsf{pk}}(\text{set}) \odot E_{\mathsf{pk}}(i) \oplus (1 - E_{\mathsf{pk}}(\text{set})) \odot E_{\mathsf{pk}}(\text{c\_max\_idx})$
    $E_{\mathsf{pk}}(\text{c\_max\_r}) = E_{\mathsf{pk}}(\text{set}) \odot E_{\mathsf{pk}}(r_i) \oplus (1 - E_{\mathsf{pk}}(\text{set})) \odot E_{\mathsf{pk}}(\text{c\_max\_r})$
  result $= (f_{\mathsf{eq}}(E_{\mathsf{pk}}(i), E_{\mathsf{pk}}(\text{c\_max\_idx})))_{i=1}^{k}$
  return result

Fig. 1: Algorithm to find a maximum in a list of fractions, including tie breaking by rank.

procedure FLOORDIVISION($E_{\mathsf{pk}}(a)$, $E_{\mathsf{pk}}(b)$, u)
  $length = \text{BITLENGTH}(u)$
  $E_{\mathsf{pk}}(lower) = E_{\mathsf{pk}}(0)$
  for $i = 1, \ldots, length$ **do**
    $E_{\mathsf{pk}}(j) = E_{\mathsf{pk}}(lower) \oplus E_{\mathsf{pk}}(2^{length-i})$
    $E_{\mathsf{pk}}(gt) = f_{\mathsf{gt}}(E_{\mathsf{pk}}(a), E_{\mathsf{pk}}(j) \odot E_{\mathsf{pk}}(b))$
    $E_{\mathsf{pk}}(lower) = E_{\mathsf{pk}}(lower) \oplus (2^{length-i} \odot E_{\mathsf{pk}}(gt))$
  return $E_{\mathsf{pk}}(lower)$

Fig. 2: Floor Division to calculate $E_{\mathsf{pk}}(\lfloor \frac{a}{b} \rfloor)$ where u is a known upper bound.

ditionally takes encrypted ranks $r = (E_{\mathsf{pk}}(r_i))_{i=1}^{k}$ as input, where the $(r_1, \ldots, r_k)$ form a permutation of $0, \ldots, k-1$, and first scales all ciphertexts q_i by a certain value, adds the encrypted ranking r_i to the scaled q_i, and then continues just as `getMaxFraction`. By the scaling the addition of r_i does not change the output if the q_i are not tied. But, if any of the inputs q_i are equal, then the party with the highest rank r_i will have the greater (encrypted) value after the addition.

Finally, in Figure 2 we introduce a new MPC algorithm for computing the floor division $E_{\mathsf{pk}}(\lfloor \frac{a}{b} \rfloor)$ from two encrypted natural numbers $E_{\mathsf{pk}}(a), E_{\mathsf{pk}}(b)$ and a publicly known upper bound $u \geq \lfloor \frac{a}{b} \rfloor$ of the result. Compared to the floor division MPC algorithm presented in [14], we require u but can be much more efficient by performing a binary search instead of iterating over a full set of values.

4 MPC Protocol for the Sainte-Laguë Method

The Sainte-Laguë method (also called Webster method) is a seat allocation method, i.e., a procedure that describes how a given number of seats is allocated to a set of parties depending on the number of votes each party has received. The Sainte-Laguë method is used by parliamentary elections in many countries, for example, Indonesia, New Zealand, Nepal, Sweden, Norway, Germany, and Kosovo. As part of computing the election result, these elections run the Sainte-Laguë method multiple times on different inputs. For example, the

official evaluation of the final seat distribution of the German Bundestag of the election in 2021 required running the Sainte-Laguë method 23 times (in addition to several other steps, as explained in Section 1). Hence, in order to obtain an efficient tally-hiding voting system for these elections, it is crucial to design a heavily optimized MPC component for computing the Sainte-Laguë method. In this section, we first give both a general overview of the Sainte-Laguë method and then present our efficient tally-hiding MPC algorithm, including several optimizations and variations.

4.1 Computing a Sainte-Laguë Distribution

There are essentially two distinct (but provably equivalent [21]) algorithms for computing the seat allocation following the Sainte-Laguë method, one based on highest quotients and one on finding suitable denominators. Both algorithms take as input the number of seats n_{seats} to be distributed and, for each party $j \in [n_{\text{parties}}]$, the total number of votes n^j_{votes} that party j has received. They return the number of seats assigned to each party.

- **Highest-Quotients.** For $i \in [n_{\text{seats}}], j \in [n_{\text{parties}}]$ compute the quotients $q^j_i := \frac{n^j_{\text{votes}}}{2i+1}$. Let M be the list of the n_{seats} highest quotients. Then party j is assigned k seats, where k is the number of quotients in M that belong to j, i.e., quotients of the form $q^j_i, i \in [n_{\text{seats}}]$.
- **Suitable-Denominator.** Given a *suitable denominator* d, the number n^j_{seats} of seats assigned to party j is computed as $n^j_{\text{seats}} = \lfloor \frac{n^j_{\text{votes}}}{d} \rceil$, where $\lfloor \cdot \rceil$ denotes rounding to the closest integer (rounding of .5 can be chosen to be either up or down and can be chosen differently for each j). A denominator d is *suitable* if the result of this computation leads to the number of desired total seats, i.e., if $\sum_j n^j_{\text{seats}} = n_{\text{seats}}$. To find a suitable denominator, one generally starts with an arbitrary denominator d, e.g., $d = \left\lfloor \frac{\sum_j n^j_{\text{votes}}}{n_{\text{seats}}} \right\rceil$, checks the corresponding number of seats that would be assigned, and then tweaks d until a suitable value has been found.

For both algorithms, there might be ties that would need to be resolved. E.g., in the highest-quotients algorithm, there might be two equal quotients while there is only enough space left in M for one of them. In the suitable-denominator algorithm, it can happen that all suitable denominators are such that the quotients of multiple parties end on .5 and some of which need to be rounded up while others need to be rounded down to achieve an overall sum of n_{seats}. The Sainte-Laguë method does not define any specific tie-breaking mechanism. Instead, elections using this method additionally need to specify how they handle ties.

1: **procedure** AddSeatBasic($(q = E_{\mathsf{pk}}^{\mathrm{frac}}(q_{\mathrm{current}}^{j}))_{j=0}^{n_{\mathrm{parties}}-1}$, $s = (E_{\mathsf{pk}}(n_{\mathrm{seats}}^{j}))_{j=0}^{n_{\mathrm{parties}}-1}$)
2: $t = (E_{\mathsf{pk}}(m_j))_{j=0}^{n_{\mathrm{parties}}-1} = \textsc{getMaxFraction}(q)$
3: **for** $j \in [n_{\mathrm{parties}}]$ **do**
4: $d = E_{\mathsf{pk}}(q_j.d) \oplus 2 \odot t_j$
5: $q_j = \textsc{Fraction}(E_{\mathsf{pk}}(q_j.n), d)$ ▷ Update q
6: $s_j = s_j \oplus t_j$ ▷ Update seats (s)
7: **return** q, s

Fig. 3: One iteration step of SLQBasic.

4.2 Tally-Hiding MPC Realization of Sainte-Laguë

We want to construct a tally-hiding MPC component that takes as inputs $E_{\mathsf{pk}}(n_{\mathrm{votes}}^{j})$ for each party as well as publicly known values n_{parties} and n_{seats}[4], and computes the encrypted Sainte-Laguë seat distribution $E_{\mathsf{pk}}(n_{\mathrm{seats}}^{j})$. As an initial insight, we observe that basing the MPC protocol on the suitable-denominator algorithm is generally very inefficient: This algorithm has to iterate over several potential denominators d until a suitable one is found. Since the number of iterations required to find d would reveal non-trivial information about the secret inputs, the MPC protocol would rather have to be constructed such that it always uses an apriori fixed number m of iterations (some of which will discard their results if a suitable divisor has already been found by a previous iteration) where m must be chosen sufficiently large such that, for all possible input sequences, a suitable divisor d is guaranteed to always be found. This worst case approximation introduces a lot of additional overhead.

Therefore, we have constructed a basic tally-hiding MPC realization SLQBasic of the Sainte-Laguë method following the highest-quotients approach: each party j is assigned its current quotient q_{current} (see the description of the highest-quotients algorithm) and seats n_{seats}^{j} thus far. Figure 3 shows this excerpt of a single iteration step. Note that this SLQBasic uses the fast `getMaxFraction` algorithm in all iterations, and hence, breaks ties via a fixed mechanism that always assigns the seat to the party with the highest index j.

Support for Breaking Ties by Lot. Many elections use more involved tie-breaking algorithms than the default one implemented by SLQBasic. For example, for many parliamentary elections, e.g., elections in Indonesia, Sweden, and Germany, whenever several parties are tied for a seat, then a new lot is drawn to resolve the tie. A more general tally-hiding MPC implementation SLQCustomTiebreaking for this election does not only have to support this tie-breaking mechanism but also has to keep secret whether any lots were drawn and what the result was. In particular, to build such a SLQCustomTiebreaking we need to first extend/modify the iteration step AddSeatBasic shown in Figure 3, obtaining a new subroutine AddSeatTieBreaking which takes as additional input

[4] As we explain in our full version [25], all MPC algorithms presented in this and the next section can be extended to run with a secret number of seats n_{seats}, as long as an upper limit of seats is known.

an encrypted ranking of parties $r = (E_{\mathsf{pk}}(r_0), \ldots, E_{\mathsf{pk}}(r_{n_{\text{parties}}-1}))$ where r is a uniformly chosen permutation of $0, \ldots, n_{\text{parties}} - 1$, and then resolves ties based on that ranking.

We construct AddSeatTieBreaking by making use of `getMaxFractionByRank` as presented in Section 3. That is, we replace the call to `getMaxFraction` in Line 2 of AddSeatBasic by our algorithm `getMaxFractionByRank` which takes as additional input the ranking r. Based on this AddSeatTieBreaking, we have constructed two versions of a SLQCustomTiebreaking MPC component which implement Sainte-Laguë. In essence, these MPC components first compute, in each iteration, an encrypted ranking r that encodes the result of tie-breaking and then use AddSeatTieBreaking with that r. There are two main optimizations that we introduce in both cases: First, for tie-breaking by lot, we run a distributed randomness generation protocol [22] for each iteration to then compute r based on the results. Since this step is input/vote independent, it can be pre-computed even before the election. Secondly, observe that if a tie occurs between m parties in one iteration of the quotient approach while there are at least m seats to be distributed, then all parties in the tie will obtain a seat in the next m iterations, i.e., it does not actually matter how this tie is broken. Hence, only ties during the last $n_{\text{parties}} - 1$ iterations need to be handled by AddSeatTieBreaking, while otherwise we use the faster AddSeatBasic algorithm.

4.3 Sainte-Laguë based on Floor Division

While our MPC algorithms SLQBasic and SLQCustomTiebreaking based on the highest-quotients approach are already practical in terms of efficiency, they always use n_{seats} iterations to assign all seats and thus do not scale overly well for elections where a high number of seats n_{seats} needs to be allocated. To improve performance in such cases, we propose a new algorithm for computing the Sainte-Laguë method which we call *floor division method*. Our floor division method is different from the highest-quotient and the suitable-denominator methods and allows us to construct an MPC component, called SLQFloorDiv, that requires n_{parties} instead of n_{seats} many iterations, and thus, is more efficient in the common case that the number of seats exceeds the number of parties. In what follows, we first present the floor division method and then describe our MPC component SLQFloorDiv.

Description of our Method. Intuitively, the main idea of our floor division method is to replace the initial iterations and hence seat assignments of the quotient method by computing an under- and an overestimation of the final seat allocation, and then run only the final (at most n_{parties} many) iterations of the quotient method to add/remove a seat from both of these initial estimations until exactly n_{seats} many seats are assigned. As we prove one of the resulting final seat distributions will be the correct Sainte-Laguë distribution, and it can be determined efficiently which one is correct.

Concretely, for each party j compute $m_j := \lfloor \frac{n^j_{\text{votes}} \cdot n_{\text{seats}}}{n_{\text{votes}}} \rfloor$. For the underestimation case, we start by assigning m_j seats to party j. Note that $s^{\mathsf{min}}_{\mathsf{initial}} :=$

$\sum_{j \in [n_{\text{parties}}]} m_j$ might be smaller than n_{seats}, but not smaller than $n_{\text{seats}} - n_{\text{parties}}$. Hence, in order to distribute exactly n_{seats} seats in total, we distribute the remaining $n_{\text{seats}} - s^{\mathsf{min}}_{\mathsf{initial}}$ ($\leq n_{\text{parties}}$) seats to the parties by executing the final iterations of the highest-quotients method (and the desired tie-breaking mechanism). That is, starting from the intermediate quotients $q^j_{m_j} := \frac{n^j_{\text{votes}}}{2m_j+1}$ instead of starting from the initial q^j_0 for each party j. For the overestimation case, we start by assigning $m_j + 1$ seats to party j, which might result in at most n_{parties} additional seats being assigned beyond the desired total of n_{seats}. To remove those seats, we use a reverse variant of the highest-quotients algorithm. For this purpose, we again initialize the quotients as $q^j_{m_j}$ and then, in each iteration step, determine the minimal current quotient and remove a seat from the corresponding party (using the desired tie-breaking mechanism). Then, we update the quotient of that party by reducing the denominator by 2.[5] This continues until only a total of n_{seats} seats is distributed.

Finally, to figure out which result is the correct Sainte-Laguë distribution, we evaluate the underestimation case an additional time to compute the next seat that would be assigned. If the corresponding quotient is less than all the initial quotients $q^j_{m_j}$ of the underestimation, then the result computed based on the underestimation is the correct seat distribution. Otherwise, the result computed based on the overestimation is the correct seat distribution. In the full version [25], we show the following result for our algorithm:

Lemma 1 (Correctness of SLQFloorDiv). *The algorithm SLQFloorDiv as presented above is correct, i.e., always outputs the seat allocation according to the Sainte-Laguë method with the desired tie-breaking.*

Tally-Hiding MPC Component. Using our building blocks described in Section 3 and the other building blocks from Section 2, most of our tally-hiding MPC protocol for computing the above algorithm for Sainte-Laguë is straightforward. The main issue left to be solved is that the number of iterations that our algorithm needs to add/subtract seats from the MPC protocol reveals initial seat assignment (this would reveal non-trivial information about the inputs/votes). We solve this by always using n_{parties} iterations in our MPC protocol, which is an upper bound on the number of iterations that are needed.

Our benchmarks of single runs of the Sainte-Laguë algorithms show that this variant of the Sainte-Laguë method is indeed faster than SLQCustomTiebreaking for larger numbers of seats and smaller numbers of parties. For example, we have the following runtime for ten parties: To distribute 60 resp. 100 seats using SLQCustomTiebreaking, the runtime is $6.7h$ resp. $12.6h$ while SLQFloorDiv only needs $4.7h$ resp. $5h$. However, for smaller numbers of, say, 20 seats, SLQCustomTiebreaking is faster with $1.6h$ instead of SLQFloorDiv, which needs $4.6h$. We compare benchmarks for further values of n_{seats} and n_{parties} in Appendix A.2.

[5] It might happen that all $m_j + 1$ seats are removed from a party j. In that case, this party is ignored in the following iterations. Note that this special case is non-trivial to implement in our SLQFloorDiv MPC component since we cannot reveal the values m_j or the quotients.

5 Election of the German Parliament (Bundestag)

The election of the German federal parliament, the *Bundestag*, which consists of at least 598 seats is a combination of proportional representation and first-past-the-post voting. Each voter has two votes: a constituency vote (called *first vote*) given towards an individual candidate, who is typically but not necessarily also a member of a party, and a vote for state-specific party lists (called *second vote*) which determines the proportions of parties in the parliament. The first votes are evaluated for each of the 299 constituencies individually: The candidate with the most votes wins the constituency and is guaranteed a seat in the parliament, called a *direct mandate*.[6] Each constituency belongs to exactly one of the 16 German states, say state $l \in L$ where L is the set of all states. We denote by $s^{\mathrm{d}}_{j,l}$ the total number of direct mandates that candidates of party j win in state l.

Let $v_{j,l}$ be the number of second votes for party j in state l and $v_j := \sum_{l \in L} v_{j,l}$ the total number of second votes for party j. In the next step, the baseline of 598 seats of the parliament are assigned to the states in proportion to their number of inhabitants; we call this the *first top distribution* and refer to these seats as *state seats*. A party j can obtain state seats if v_j is at least 5% of all second votes, j has obtained at least $\sum_{l \in L} s^{\mathrm{d}}_{j,l} \geq 3$ direct mandates, or j represents a special minority. Let $\mathfrak{S}$ be the set of parties that are allowed to obtain state seats. Then, for each state l, the state seats are assigned to parties $j \in \mathfrak{S}$ following the Sainte-Laguë method based on $v_{j,l}$. The resulting seats are called *quota seats*, denoted by $s^{\mathrm{q}}_{j,l}$ for party j and state l. We call this distribution the *first low distribution*. It usually happens in several states l that a party $j \in \mathfrak{S}$ wins more direct mandates and hence guaranteed seats for their candidates than the party actually receives in terms of quota seats, i.e., $s^{\mathrm{d}}_{j,l} > s^{\mathrm{q}}_{j,l}$. In such cases, the overall size of the parliament is increased, and the seat assignment to parties is updated such that *(i)* parties have enough seats for all their candidates with direct mandates and *(ii)* the number of seats given to party j in the final parliament is "close" to the Sainte-Laguë seat distribution based on v_j (up to 3 additional seats, called *overhang seats*, are tolerated). This is computed via the following procedure.

Let $s^{\min}_j := \sum_{l \in L} \left(\max(\lceil \frac{s^{\mathrm{d}}_{j,l} + s^{\mathrm{q}}_{j,l}}{2} \rceil, s^{\mathrm{d}}_{j,l}) \right)$ be a lower bound for the seats that party $j \in \mathfrak{S}$ will receive. Compute $d_{\mathrm{no}} := \min_{j \in \mathfrak{S}} \left(\frac{v_j}{s^{\min}_j - 0.5} \right)$. Then, for each party $j \in \mathfrak{S}$, it is determined whether there is a state $l \in L$ such that $t(j,l) := s^{\mathrm{d}}_{j,l} - s^{\mathrm{q}}_{j,l} > 0$. This value is also called the *threatening overhang* of party j in state l. Based on these values, one computes a set of divisors: $D_{\mathrm{overh}} = \{ \frac{v_j}{s^{\min}_j - i} \mid i \in \{0.5, 1.5, 2.5, 3.5\}, j \in \mathfrak{S} \text{ and } \exists l : \ t(j,l) > 0, t(j,l) + 1 > i \}$. Let d_{overh} be the fourth smallest element of D_{overh} and set $d := \min(d_{\mathrm{no}}, d_{\mathrm{overh}})$.

Then, as in the suitable-denominator algorithm (c.f. Section 4), party $j \in \mathfrak{S}$ receives $n^j_{\mathrm{seats}} := \lfloor \frac{n^j_{\mathrm{votes}}}{d} \rceil$ seats, where .5 is always rounded up, i.e., in this step,

[6] Ties for the first place and hence the direct mandate are resolved by lot. Ties in any of the following iterations of the Sainte-Laguë method are also resolved by lot with one exception discussed below.

ties are resolved by giving every tied party a seat. The resulting distribution is called the *second top distribution*. Next, for each party $j \in \mathfrak{S}$, these n^j_{seats} are assigned to individual states following the Sainte-Laguë method weighted by $v_{j,l}$, resulting in the *second low distribution*.

In addition to those n^j_{seats} seats, (some) parties further receive $\alpha_j := s^{\min}_j - n^j_{\text{seats}}(\leq 3)$ overhang seats to cover a possibly remaining surplus of direct mandates. These overhang seats are then also distributed to states according to the smallest α_j many values from the following set: $O^j_{\text{overh}} = \{\frac{v_{j,l}}{s^{\min}_{j,l} - i} \mid i \in \{0.5, 1.5, 2.5\}, l \in L, i < \alpha_j\}$.

All seats assigned to party $j \in \mathfrak{S}$ in state l are then assigned to candidates as follows. First, candidates of party j that won a direct mandate in a constituency of state l obtain a seat. The remaining seats, if any, are assigned to the party candidate list for that state, starting with the first one and skipping any candidates that have already obtained a direct mandate. Finally, if there are any direct mandates for candidates that do not belong to a party from $\mathfrak{S}$, then each of these candidates receives a seat that is added to the parliament. The resulting set of seats defines the updated size of the parliament.

Our Tally-Hiding Realization. We construct our e-voting system by following the general Ordinos approach, except for one difference. The original Ordinos framework proposed in [19] was designed for elections without electoral constituencies or with just a single constituency where all votes are treated equally. We capture the existence of several constituencies in the German parliamentary elections, where the result also depends on the constituency that a vote was submitted in, via the following small changes: The list of eligible voters that is published during the setup phase now additionally assigns each voter to a constituency. Ballots are extended to additionally contain (in plain) the identifier of the constituency they were cast in such that everyone can check whether ballots were cast for the correct constituency. Encrypted ballots are aggregated per constituency and then evaluated via (the MPC component for) f_{res}. We note that this difference in settings also slightly changes the meaning of full tally-hiding: For elections without electoral constituencies, only the number of submitted votes (since this is public on B) and the final result become known. In the setting with electoral constituencies, only the number of submitted votes *per constituency* and the final result becomes known. As part of our security proof (cf. Theorem 2), we define full tally-hiding for our setting and verify that the original proofs of Theorem 1 carry over in a natural way to our setting using the same preconditions.

We hence instantiate the (modified) Ordinos approach for the German election by using the threshold Paillier encryption scheme $\mathcal{E}$, choice space $C_{\mathsf{single}} \times C_{\mathsf{single}}$, standard NIZKPs $\pi^{\mathsf{KeyShareGen}}$ and π^{Enc} [10] and any standard EUF-CMA-secure signature scheme from the literature, result function $f_{\text{res-Ger}}$ for the German parliamentary election as described above, and importantly our new MPC protocol $\mathsf{P}_{\mathsf{MPC\text{-}Ger}}$ for $f_{\text{res-Ger}}$ described next.

Constructing $\mathsf{P}_{\mathsf{MPC\text{-}Ger}}$. We have constructed $\mathsf{P}_{\mathsf{MPC\text{-}Ger}}$ using the components from Sections 2 and 3 to compute the full evaluation procedure for German par-

liament, as described above. This includes all small details and special cases, e.g., computing and changing the final parliament size, determining and distributing up to 3 overhang seats per party, and exempting parties from obtaining state seats iff they did not win 5% of the total second votes, won less than 3 direct mandates, and are not representing a special minority.

Of course, capturing the full complexity of the election evaluation of the German parliament in an MPC protocol $\mathsf{P}_{\text{MPC-Ger}}$ comes at a hefty cost in terms of performance and hence runs the risk of becoming impractical. We have therefore spent considerable effort into carefully optimizing $\mathsf{P}_{\text{MPC-Ger}}$ by, among others, the following: *(i)* Computing the election result requires multiple iterations of Sainte-Laguë. We use both SLQCustomTiebreaking and SLQFloorDiv, depending on the number of seats and candidates that has to be processed in the current iteration. *(ii)* We have constructed $\mathsf{P}_{\text{MPC-Ger}}$ in such a way that, as far as possible, substeps such as repeated state-wise operations can be computed in parallel. We have performed benchmarks for various numbers of threads, which demonstrate that this is a major factor in improving performance, see Table 1. *(iii)* We first compute and reveal the set of parties that will obtain at least one seat in the parliament. This allows us to tailor the following computations to this specific set of parties and thus save time by not having to perform the same operations for (dozens of) parties that will not obtain a seat. As part of Theorem 2, we argue that this construction is still a secure MPC protocol as, intuitively, the intermediate output can be computed from/is part of the final result. *(iv)* By proposing a different algorithm for computing the final number of seats for each party in the German parliament based on an encrypted divisor $d = \min(d_{\text{no}}, d_{\text{overh}})$, we can use an efficient binary search on encrypted data to obtain this value.

We provide full details of $\mathsf{P}_{\text{MPC-Ger}}$, including all of our optimizations in the full version [25]. We have the following:

Theorem 2 (Security). *Let $\mathsf{P}_{\text{MPC-}Ger}$ be our MPC protocol from above. Then, the Ordinos instance using the components mentioned above is a verifiable*[7] *and fully tally-hiding e-voting system for the election of the German parliament.*

We prove this theorem in the full version [25]. As part of this, we define full tally-hiding for elections with constituencies, re-verify the original proofs of the Ordinos framework for our setting, and show that our $\mathsf{P}_{\text{MPC-Ger}}$ is a secure, publicly verifiable MPC protocol for $f_{\text{res-Ger}}$.

Benchmarks: We have benchmarked our system using the election data for the German parliament in 2021 available at [11]. This election had 61,181,072 eligible voters, 46,854,508 valid submitted ballots, and 47 parties with 6211 candidates running in 299 constituencies. With each trustee running on an ES-PRIMO Q957 (64bit, i5-7500T CPU @ 2.70GHz, 16 GB RAM) using 8 cores, we can evaluate (and verify, as explained in Appendix A.4) the German parlia-

[7] We actually show that our voting system achieves the stronger notion of accountability as well. That is, if the result turns out to be incorrect, then a misbehaving party causing this mistake can be identified and be held accountable.

mentary elections based on this real-world data in about a day. For more details on our setup and further benchmarks,see Appendix A and our full version [25].

6 Related Work and Conclusion

Various tally-hiding e-voting systems for a wide variety of election types have been proposed so far, e.g., [4,5,9,12,14,15,19,23,27]. For simple types of elections such as single vote (every voter submits a single vote for the candidate of their choice with the winner(s) being the candidate(s) with the most votes), it has been demonstrated that they can be performed in a verifiable tally-hiding manner, even at a large scale (see, e.g., [9,14,15,19]). However, many real-world elections, notably political ones, are much more complex and have proven to be a challenge for tally-hiding systems.

Preferential Elections: An important class of complex real-world elections are preferential ones. Tally hiding has already been studied for several voting methods from this class, with such systems typically being viable at a small to medium scale but often being impractical for large-scale applications. For example: *(i)* Recent advances in tally-hiding e-voting have managed to support instant-runoff voting (IRV) for small numbers of candidates [15, 23]. However, none of these systems remain practical for more than 6 candidates. *(ii)* Cortier et. al [9] have proposed the first MPC component that can be used to construct a fully tally-hiding voting system for single transferable vote (STV), a preferential voting method somewhat similar to IRV. However, they state that the computational cost of the resulting system would be too high for large-scale elections. *(iii)* For the Condorcet Schulze election scheme, Hertel et.al. [14] proposed an Ordinos instantiation that can handle small numbers of candidates, however, already needs about 9 days to compute an election result for 20 candidates (and essentially arbitrary numbers of voters). Cortier et. al [9] proposed an alternative tally-hiding MPC component for computing Condorcet Schulze, which is faster for small numbers of voters but can be extrapolated to also require 9 days for 20 candidates as soon as there are $\sim$32.000 voters.

Parliamentary Elections With Party-Based Seat Allocations: As already explained in the introduction, prior to our work, it had not been investigated for any election from this class, whether and in how far, it can be performed in a tally-hiding manner. In this work, we have proposed several new tally-hiding building blocks, as well as the first verifiable tally-hiding voting system for an election from this class, namely, the German parliament. Our results serve as an important foundational feasibility study, which, perhaps surprisingly and for the first time, demonstrates that even such a complex and large-scale real-world election can, in principle be performed in a verifiable fully tally-hiding manner. It is interesting future work to use our building blocks and ideas to construct tally-hiding voting systems for further elections from this class.

Acknowledgements: This research was supported by the DFG through grant KU 1434/11-1, by the Carl Zeiss Foundation, and by the Centre for Integrated Quantum Science and Technology (IQ^{ST}).

A Appendix

A.1 Details of the Setup for our Benchmarks

We use a Paillier key of size 2048 bits. The setup for our benchmarks consists of three trustees communicating over a local network. Each trustee ran on an ESPRIMO Q957 (64bit, i5-7500T CPU @ 2.70GHz, 16 GB RAM). As in [19], the benchmarks of our MPC protocols start with an already aggregated tally. Küsters et al. [19] showed for the MPC protocols in their Ordinos instances that the number of trustees does not influence the benchmarks in a noticeable way and that, due to the sublinear communication complexity of the comparison protocols, there is no significant difference between a local network and the Internet. Since our MPC protocols are based on the same primitives and basic building blocks as used by [19], the same is also true for our MPC protocols. Our benchmarks therefore focus mostly on the number of candidates/parties which is the main factor for the performance of our protocols.

A.2 Comparison of SLQCustomTiebreaking and SLQFloorDiv

We present benchmarks for both SLQFloorDiv and SLQCustomTiebreaking (cf. Section 4) in Figure 4. As the figure shows, SLQCustomTiebreaking is linear in the number of seats. While SLQFloorDiv has a larger overhead depending on the number of parties, it is nearly constant in the number of seats.

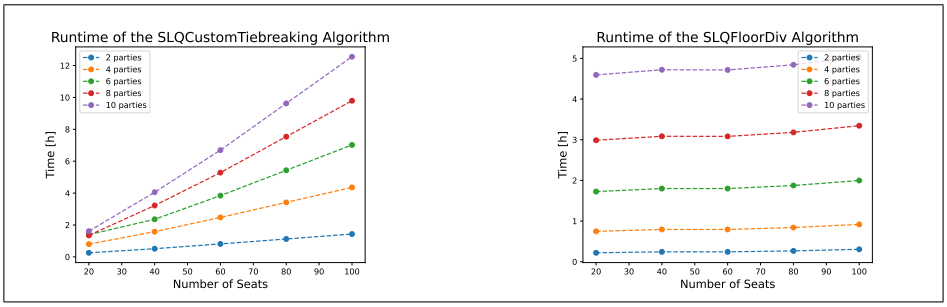

Fig. 4: Benchmarks for one execution of SLQFloorDiv and SLQCustomTiebreaking from Section 4.

A.3 Benchmarks of the Evaluation of the Elections for the German Bundestag in 2021

In Table 1 we present our benchmarks of the evaluation of the elections for the German Bundestag in 2021 using real-world data available at [11]. Each row in the table represents one main step of the algorithm, where each of these main steps is executed in sequence. Within each individual step, it is possible

# Threads per Trustee	1	2	4	8	16	32
Single-member constituency seats	40.03 h	20.04 h	10.06 h	5.06 h	2.56 h	1.32 h
Determine which parties enter the Bundestag	71 min	36 min	18 min	9 min	5 min	3 min
First low distribution	23.38 h	11.77 h	6.07 h	3.25 h	1.82 h	1.1 h
Minimal number of seats per party	11.68 h	5.84 h	2.93 h	1.46 h	0.74 h	0.36 h
Second top distribution	2.81 h	2.0 h	1.42 h	1.19 h	1.19 h	1.19 h
Second low distribution	77.06 h	38.53 h	19.4 h	12.34 h	6.3 h	5.93 h
Assigning overhang seats	6.67 h	3.33 h	2.2 h	1.14 h	1.14 h	1.14 h
Computing the final result	4 min	2 min	1 min	1 min	0 min	0 min
Total Runtime	163 h	82 h	42 h	24.3 h	13.8 h	11.1 h

Table 1: Benchmarks of the election for the German Bundestag in 2021 using real-world data available at [11] with different numbers of available parallel threads for each trustee.

to leverage parallelism. We show the resulting runtime for various numbers of threads/cpu cores. Further benchmarks are presented in the full version [25].

A.4 Verification of the Election

Verification of an election following the Ordinos approach essentially consists of two main tasks: Firstly, checking the correctness of the ballots submitted to B including verification of the ballot NIZKPs π^{Enc} for the choice space. Secondly, verifying that the MPC protocol $\mathsf{P}_{\mathsf{MPC}}$ was executed correctly.

The first task can be performed on the fly for each new ballot submitted to B while the election is still running. Notably, we use a NIZKP π^{Enc} from [10] that is standard and employed by many e-voting systems since it is very efficient and fast to verify. The second step requires checking certain data, including further NIZKPs, that is published on B while $\mathsf{P}_{\mathsf{MPC\text{-}Ger}}$ is running. Notably, all trustees also perform all of the same verification checks as part of running $\mathsf{P}_{\mathsf{MPC\text{-}Ger}}$. Hence, not only is it possible for an external observer to perform verification of $\mathsf{P}_{\mathsf{MPC\text{-}Ger}}$ in parallel to $\mathsf{P}_{\mathsf{MPC\text{-}Ger}}$ being executed. An external observer will also be done with this verification as soon as the end result is returned by $\mathsf{P}_{\mathsf{MPC\text{-}Ger}}$ because he has to perform strictly less work than the trustees running $\mathsf{P}_{\mathsf{MPC\text{-}Ger}}$. We therefore only had to benchmark the runtime of $\mathsf{P}_{\mathsf{MPC\text{-}Ger}}$ to obtain the overall time for *both computing and verifying* the election result of our system proposed in Section 5.

References

1. Adida, B.: Helios: Web-based Open-Audit Voting. In: Proceedings of the 17th USENIX Security Symposium. pp. 335–348. USENIX Association (2008)
2. Benaloh, J.: Verifiable Secret Ballot Elections. Ph.D. thesis, Yale University (1987)
3. Benaloh, J., Byrne, M.D., Eakin, B., Kortum, P.T., McBurnett, N., Pereira, O., Stark, P.B., Wallach, D.S., Fisher, G., Montoya, J., Parker, M., Winn, M.: Star-vote: A secure, transparent, auditable, and reliable voting system. In: 2013 Electronic Voting Technology Workshop / Workshop on Trustworthy Elections, EVT/WOTE '13. USENIX Association (2013)
4. Benaloh, J., Moran, T., Naish, L., Ramchen, K., Teague, V.: Shuffle-Sum: Coercion-Resistant Verifiable Tallying for STV Voting. IEEE Transactions on Information Forensics and Security **4**(4), 685–698 (2009). https://doi.org/10.1109/TIFS.2009.2033757
5. Canard, S., Pointcheval, D., Santos, Q., Traoré, J.: Practical Strategy-Resistant Privacy-Preserving Elections. In: Computer Security - ESORICS 2018. Lecture Notes in Computer Science, vol. 11099, pp. 331–349. Springer (2018). https://doi.org/10.1007/978-3-319-98989-1_17
6. Chaum, D., Carback, R., Clark, J., Essex, A., Popoveniuc, S., Rivest, R.L., Ryan, P.Y.A., Shen, E., Sherman, A.T.: Scantegrity II: end-to-end verifiability for optical scan election systems using invisible ink confirmation codes. In: 2008 USENIX/ACCURATE Electronic Voting Workshop, EVT 2008, Proceedings. USENIX Association (2008)
7. Clarkson, M.R., Chong, S., Myers, A.C.: Civitas: Toward a Secure Voting System. In: 2008 IEEE Symposium on Security and Privacy (S&P 2008). pp. 354–368. IEEE Computer Society (2008). https://doi.org/10.1109/SP.2008.32
8. Cortier, V., Galindo, D., Glondu, S., Izabachène, M.: Election Verifiability for Helios under Weaker Trust Assumptions. In: Computer Security - ESORICS 2014. Proceedings, Part II. Lecture Notes in Computer Science, vol. 8713, pp. 327–344. Springer (2014). https://doi.org/10.1007/978-3-319-11212-1_19
9. Cortier, V., Gaudry, P., Yang, Q.: A Toolbox for Verifiable Tally-Hiding E-Voting Systems. In: Computer Security - ESORICS 2022. Lecture Notes in Computer Science, vol. 13555, pp. 631–652. Springer (2022). https://doi.org/10.1007/978-3-031-17146-8_31
10. Damgård, I., Jurik, M., Nielsen, J.B.: A generalization of Paillier's public-key system with applications to electronic voting. International Journal of Information Security **9**(6), 371–385 (2010). https://doi.org/s10207-010-0119-9
11. Der Bundeswahlleiter: Wahl zum 20. Deutschen Bundestag am 26. September 2021: Heft 3 Endgültige Ergebnisse nach Wahlkreisen (2021), https://bundeswahlleiter.de/dam/jcr/cbceef6c-19ec-437b-a894-3611be8ae886/btw21_heft3.pdf, https://www.bundeswahlleiter.de/bundestagswahlen/2021/ergebnisse/opendata/csv/
12. Haines, T., Pattinson, D., Tiwari, M.: Verifiable Homomorphic Tallying for the Schulze Vote Counting Scheme. In: Verified Software. Theories, Tools, and Experiments - VSTTE 2019. Lecture Notes in Computer Science, vol. 12031, pp. 36–53. Springer (2019). https://doi.org/10.1007/978-3-030-41600-3_4
13. Heather, J.: Implementing STV securely in Prêt à Voter. In: IEEE 20th Computer Security Foundations Symposium (CSF 2007). pp. 157–169. IEEE Computer Society (2007). https://doi.org/10.1109/CSF.2007.22

14. Hertel, F., Huber, N., Kittelberger, J., Küsters, R., Liedtke, J., Rausch, D.: Extending the Tally-Hiding Ordinos System: Implementations for Borda, Hare-Niemeyer, Condorcet, and Instant-Runoff Voting. In: Electronic Voting - 6th International Joint Conference, E-Vote-ID 2021, Proceedings. pp. 269—-284. University of Tartu Press (2021)
15. Huber, N., Küsters, R., Krips, T., Liedtke, J., Müller, J., Rausch, D., Reisert, P., Vogt, A.: Kryvos: Publicly Tally-Hiding Verifiable E-Voting. In: CCS 2022. pp. 1443–1457. ACM (2022). https://doi.org/10.1145/3548606.3560701
16. Juels, A., Catalano, D., Jakobsson, M.: Coercion-resistant Electronic Elections. In: Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society, WPES 2005. pp. 61–70. ACM (2005). https://doi.org/10.1145/1102199.1102213
17. Kiayias, A., Zacharias, T., Zhang, B.: DEMOS-2: Scalable E2E Verifiable Elections without Random Oracles. In: CCS 2015. pp. 352–363. ACM (2015). https://doi.org/10.1145/2810103.2813727
18. Kiayias, A., Zacharias, T., Zhang, B.: End-to-End Verifiable Elections in the Standard Model. In: Advances in Cryptology - EUROCRYPT 2015. Lecture Notes in Computer Science, vol. 9057, pp. 468–498. Springer (2015). https://doi.org/10.1007/978-3-662-46803-6_16
19. Küsters, R., Liedtke, J., Müller, J., Rausch, D., Vogt, A.: Ordinos: A Verifiable Tally-Hiding E-Voting System. In: 2020 IEEE European Symposium on Security and Privacy (EuroS&P 2020). pp. 216–235. IEEE Computer Society (2020). https://doi.org/10.1109/EuroSP48549.2020.00022
20. Küsters, R., Müller, J., Scapin, E., Truderung, T.: sElect: A Lightweight Verifiable Remote Voting System. In: IEEE 29th Computer Security Foundations Symposium (CSF 2016). pp. 341–354. IEEE Computer Society (2016). https://doi.org/10.1109/CSF.2016.31
21. Lijphart, A.: Degrees of proportionality of proportional representation formulas. RIVISTA ITALIANA DI SCIENZA POLITICA **13**(2), 295–305 (1983)
22. Lipmaa, H., Toft, T.: Secure Equality and Greater-Than Tests with Sublinear Online Complexity. In: Automata, Languages, and Programming, ICALP 2013. Lecture Notes in Computer Science, vol. 7966, pp. 645–656. Springer (2013). https://doi.org/10.1007/978-3-642-39212-2_56
23. Ramchen, K., Culnane, C., Pereira, O., Teague, V.: Universally Verifiable MPC and IRV Ballot Counting. In: Financial Cryptography and Data Security, FC 2019. Lecture Notes in Computer Science, vol. 11598, pp. 301–319. Springer (2019). https://doi.org/10.1007/978-3-030-32101-7_19
24. Ryan, P.Y.A., Rønne, P.B., Iovino, V.: Selene: Voting with transparent verifiability and coercion-mitigation. In: Financial Cryptography and Data Security, FC 2016. Lecture Notes in Computer Science, vol. 9604, pp. 176–192. Springer (2016). https://doi.org/10.1007/978-3-662-53357-4_12
25. Wabartha, C., Liedtke, J., Huber, N., Rausch, D., Küsters, R.: Fully Tally-Hiding Verifiable E-Voting for Real-World Elections with Seat-Allocations. Cryptology ePrint Archive, Paper 2023/1289 (2023), https://eprint.iacr.org/2023/1289, Full Version of this Paper
26. Wabartha, C., Liedtke, J., Huber, N., Rausch, D., Küsters, R.: Implementation of our System. (2023), https://github.com/JulianLiedtke/ordinos-bundestag
27. Wen, R., Buckland, R.: Minimum Disclosure Counting for the Alternative Vote. In: E-Voting and Identity, Second International Conference, VoteID 2009. Proceedings. Lecture Notes in Computer Science, vol. 5767, pp. 122–140. Springer (2009). https://doi.org/10.1007/978-3-642-04135-8_8