



More Efficient Lattice-Based Electronic Voting from NTRU

Patrick Hough^{a,1}  , Caroline Sandsbråten²   and Tjerand Silde²  

¹ University of Oxford, Mathematical Institute, Oxford, United Kingdom

² Norwegian University of Science and Technology, Department of Information Security and Communication Technology, Trondheim, Norway

Abstract. In recent years, there has been much focus on developing core cryptographic primitives based on lattice assumptions, driven by the NIST call for post-quantum key encapsulation and digital signature algorithms. However, more work must be conducted on efficient privacy-preserving protocols based on quantum-safe assumptions.

Electronic voting is one such privacy-preserving protocol whose adoption is increasing across the democratic world. E-voting offers both a fast and convenient alternative to postal voting whilst further ensuring cryptographic privacy of votes and offering full verifiability of the process. Owing to the sensitivity of voting and its infrastructure challenges, it is crucial to ensure security against quantum computers is baked into e-voting solutions.

We present an e-voting scheme from quantum-safe assumptions based on the hardness of the RLWE and NTRU lattice problems, providing concrete parameters and an efficient implementation. Our design achieves a factor $5.3\times$ reduction in ciphertext size, $2.5\times$ reduction in total communication cost, and $2\times$ reduction in total computation time compared to the state-of-the-art lattice-based voting scheme by Aranha et al. (ACM CCS 2023). We argue that the efficiency of this scheme makes it suitable for real-world elections.

Our scheme makes use of non-ternary NTRU secrets to achieve optimal parameters. In order to compute the security of our design, we extend the ternary-NTRU work of Ducas and van Woerden (ASIACRYPT 2021) by determining the concrete fatigue point (for general secrets) of NTRU to be $q = 0.0058 \cdot \sigma^2 \cdot d^{2.484}$ (above which parameters become *overstretched*) for modulus q , ring dimension d , and secrets drawn from a Gaussian of parameter σ . We consider this relation to be of independent interest and demonstrate its significance by improving the efficiency of the (partially) blind signature scheme by del Pino and Katsumata (CRYPTO 2022).

Keywords: Lattice Cryptography · Electronic Voting · NTRU

1 Introduction

With the advent of quantum computers, all public key primitives based on the hardness of factoring or computing discrete logarithms will be deemed insecure.

To mitigate this, there has been an international effort to replace these primitives with ones based on assumptions conjectured to be secure against quantum adversaries. This process, led by the National Institute of Standards and Technology (NIST) in the US, has recently concluded with the selection of standards for post-quantum key encapsulation

E-mail: patrickhough@pm.me (Patrick Hough), caroline.sandsbraten@ntnu.no (Caroline Sandsbråten), tjerand.silde@ntnu.no (Tjerand Silde)

^aWork done in part while visiting the Norwegian University of Science and Technology.



and digital signature algorithms. Three of four standards [SAB⁺22, LDK⁺22, PFH⁺22] are built from *structured lattice assumptions*; Ring Short Integer Solution (RSIS) [Ajt96, PR06, LM06], Ring Learning With Errors (RLWE) [Reg05, LPR10], and NTRU [HPS98], or the *module* versions of the two former assumptions (MSIS/MLWE) [LS15]. Despite this standardization effort, there is still much work to be done in designing *privacy-preserving* primitives from quantum-safe assumptions. In this work, we focus on one of these; electronic voting (e-voting). More precisely, we consider internet voting which allows for fully remote ballot casting via a voter’s device as opposed to voting machines at a polling station, though our framework could in theory be implemented in this setting also. Herein, ‘e-voting’ should be read as synonymous with internet voting.

Electronic Voting. E-voting has become increasingly prevalent with the first experiments for democratic elections beginning around the turn of the millennium. The first binding election to be carried out online was for the Arizona primary in 2000 [CBS00]. In 2005, Estonia offered internet voting nationally [Vin15] and in 2023, over 63% of the votes cast in Estonian parliamentary elections were cast online¹. Switzerland used its Swiss Post voting system in the 2023 national elections for the first time [Swi23] and continues to be one of the leaders in e-voting uptake. Ontario, Canada increasingly offers online voting with 177 municipalities exclusively using online voting in the 2018 municipal elections [CAE19]. In Australia, over 650,000 online voters participated in the 2021 state election in New South Wales [New21]. E-voting is also used in the cryptographic community, where the International Association for Cryptologic Research (IACR) is using Helios [Adi08] for their elections².

E-voting has a number of attractive advantages. Analysis of Estonian local elections in 2017 showed the per-vote cost of online ballots was a factor $2\times$ to $10\times$ cheaper than election-day paper ballots [KCK⁺18]. Moreover, the 2023 Estonian parliamentary elections revealed that the environmental impact (CO2 emissions) of paper ballots was 180 times higher than that for online ballots and its adoption has resulted in a high voter satisfaction and turnout rate [Sol01, SMPS16].

E-voting offers to enhance both the *integrity* and *privacy* of voting. The first attractive property is verifiability; both individual and universal. Individual verifiability allows a voter to check that their ballot was recorded correctly in the final count, whilst universal verifiability allows anyone to check that parties involved in ballot processing carried out their tasks correctly. While this represents a great bolstering to the integrity of the voting process, it can be executed whilst preserving the privacy of voters and their ballots. The second significant property enabled by e-voting is the distribution of the ballot processing. A distributed decryption ensures privacy of ballots since a single honest decryption server prevents a connection between voters and ballots.

Moreover, we emphasise the importance of *long-term* privacy of electronically cast ballots, where encrypted ballots submitted today potentially can be tied to users in the future if we can break the encryption scheme. The increasing deployment of e-voting protocols currently outpaces solutions providing security against quantum computers. Evidenced by the rapid adoption of the PQC NIST standards, the urgency to defend against a potential quantum adversary is plain to see. One possible solution for e-voting is to deploy schemes providing *everlasting privacy* [HMMP23], however, these schemes do not offer integrity against quantum computers, creating problems down the road whenever quantum computers are available. It is not certain that the public will be the first to know when classical assumptions are breakable, and it is thus essential that quantum security be baked into the designs for e-voting from the outset, and particularly to construct e-voting schemes based on assumptions conjectured to be secure against quantum adversaries.

¹valimised.ee/en/archive/statistics-about-internet-voting-estonia

²iacr.org/elections/eVoting

From a desire to achieve these privacy enhancements has emerged the ‘mix-and-decrypt’ paradigm. Here, multiple servers verifiably shuffle encrypted ballots before they are decrypted in a distributed manner. This is commonly achieved by combining a verifiable mix-net [Cha81] with a distributed public-key encryption scheme. Many previous voting designs have used this structure and in 2019, Switzerland (via its national postal service Swiss Post), deployed a national e-voting infrastructure using this paradigm and the Bayer-Groth mix-net [BG12].

Voting from Quantum-Safe Assumptions. Despite much success in developing e-voting protocols, only a few are based on assumptions believed to be quantum-safe. The most notable works are the schemes by del Pino et al. [dPLNS17], Aranha et al. [ABG⁺21], Farzaliyev et al. [FWK21], and Aranha et al. [ABGS23], the latter being the most efficient scheme based on (Ring) SIS and (Ring) LWE. Of these schemes, only the last one satisfies the golden mix-and-decrypt standard for general ballots. Even then, the communication cost of this scheme is around two orders of magnitude greater than the one employed by Swiss Post based on classical assumptions.

The inefficiency of the state-of-the-art scheme in [ABGS23] stems from the need to decrypt ballots correctly being hindered by a few key features of their design.

1. In order to optimise the computational cost of lattice-based protocols and to rely on reductions to worst-case problems, one would like to use polynomial rings whose dimension is a power of two. This imposes the first constraint on parameters ([ABGS23] uses ring dimension 4096).
2. The mixing and distributed decryption stages both use homomorphic operations on encrypted ballots. This has the effect of increasing the noise within each ciphertext. To accommodate this, one must use a ring with a larger modulus to ‘soak up’ this extra noise. This larger modulus itself increases the size of objects in the scheme and decreases the security of the underlying assumption, in turn requiring a larger ring dimension to ensure that the concrete instantiation is still secure.
3. The distributed decryption process requires decryption servers to use so-called ‘noise-drowning’ [BD10] to ensure that decryption shares do not leak anything about the server’s decryption key. Further, each decryption server must prove, in zero-knowledge, that they have applied this noise drowning operation. So far, proving knowledge of such a large element over lattices, can only be done using ‘approximate’ proofs where one can give only approximate guarantees about the size of the noise-drowning term [BBC⁺18]. The noise drowning hugely increases the noise in ciphertexts, and the loose zero-knowledge proof further pushes up parameters if correct ballot decryption is to be ensured.

Unfortunately, all three of these features appear to be crucial in realising the coveted mix-and-decrypt framework in a quantum-secure fashion; the power-of-two ring dimension allows for a highly optimised implementation and the noise-growing and homomorphic operations are fundamental to constructing their mix-net and distributed decryption building blocks. Despite two decades of lattice-based distributed decryption design, no efficient alternative to noise-drowning has been found.

Given the large efficiency gap between classical schemes and the work in [ABGS23], it is hard to see how significant efficiency improvements can be found without a new approach.

Voting from NTRU. In privacy-preserving protocols, zero-knowledge proofs (ZKPs) are deployed to verify the honest actions of parties, and usually dominate the communication cost. Observation of the recent NIST PQC standards reveals that both the RLWE and RSIS problems are used, however, there is a third long-standing problem which appears in

the Falcon digital signature [PFH⁺20]; NTRU. Crucially, NTRU ciphertexts contain only a single ring element with three secret elements, giving rise to simpler ZKP relations when compared to their two-component RLWE-based counterparts such as the BGV encryption scheme [BGV12], as used in [ABGS23], which contains five secret elements. Thus, NTRU might appear to be an attractive candidate problem from which to design a voting protocol.

However, whilst the hardness of the RLWE and RSIS problems are well understood, the picture for NTRU is less clear. Recall, the NTRU problem [HPS98]. Let R_q be a polynomial ring of dimension d and modulus q and sample polynomials f and g with coefficients from some discrete Gaussian D_σ^d . Informally, the NTRU problem is to recover f and g given h , where $h = g/f \in R_q$. In recent years, it has been shown that NTRU is vulnerable to a unique attack when defined for so-called ‘overstretched’ parameters [ABD16, CJL16] i.e. when the modulus q is very large compared to d . Whilst a line of recent works [KF17, DvW21] has made progress in understanding the parameters for which this attack applies, it is not clear how the size of the secrets f and g (parametrized by σ) influence the feasibility of the attack. Thus, designs using large parameters as found in privacy-preserving lattice constructions tend to use the RLWE and RSIS problems for which such behaviour is well understood.

1.1 Our Contribution

We propose an electronic voting protocol in the mix-and-decrypt paradigm based on the RLWE and NTRU lattice assumptions. We construct each building block from the ground up by presenting an NTRU-RLWE-based verifiable distributed decryption scheme and verifiable mix-net. Moreover, via an in-depth analysis of the NTRU problem, we provide a concrete description of the hardness of the NTRU problem for general secret sizes. We demonstrate the significance of this relation in allowing optimal parameter selection for our scheme and other NTRU-RLWE-based works in which large parameters are necessary. Finally, we give an efficient implementation of our voting scheme demonstrating significant efficiency gains over the state-of-the-art in both communication and computational cost.

Verifiable Mix-Net from NTRU. We present a verifiable mix-net for NTRU ciphertexts comprising a series of shuffle servers that each apply a secret permutation to the set of input ciphertexts. As long as at least one shuffle server is honest, the set of input ciphertexts cannot be pair-wise matched to the set of output ciphertexts. Our mix-net is simpler than the one in [ABGS23] owing to the single-element NTRU ciphertexts which yield a cleaner protocol than two-element BGV ciphertexts do. Furthermore, one proves knowledge of fewer secret objects when applying ZKPs for verifiability.

Verifiable Distributed Decryption from NTRU. We present a distributed decryption protocol based on a variant of the NTRUEncrypt scheme of Steinfeld and Stehlé [SS11], proving its security using both the RLWE and NTRU assumptions. This allows for more favourable parameters owing to a computationally secure public NTRU key (vs. a statistically secure one in [SS11]). We then apply an exact zero-knowledge proof (ZKP) in order to prove the well-formedness of decryption shares. In particular, this proof proves knowledge of the large noise drowning term needed to prevent leakage of the decryption key and does so in an *exact* fashion. That is, our ZKP (which is a modification of the one by Bootle et al. [BLNS21]) proves a tight bound on the size of the noise drowning term. To our knowledge, this is the first exact amortized ZKP of a ‘large’ secret vector for lattice relations and may be of independent interest. We note that while this makes the proof of distributed decryption larger, it allows for a less restrictive correctness condition, leading to better global parameters throughout the scheme, more than making up for any additional communication cost incurred by this proof.

Table 1: Per vote comparison to [ABGS23] of ciphertexts, shuffle proofs, decryption proofs, and overall with four servers. Shuffles are sequential, while decryption is done in parallel.

Scheme	Ciphertexts	Shuffle	Dist. Dec.	Total
[ABGS23] [KB]	80	370	157	2188
Our [KB]	15	130	85	875
[ABGS23] [ms]	0.74	261	138	1182
Our [ms]	0.20	62	328	576

NTRU Security Analysis. We build upon the work of Ducas and van Woerden [DvW21], on NTRU hardness, to analyse the so-called ‘overstretched attack’ against NTRU when the norm of the secrets grows with respect to the dimension and modulus. We stress that [DvW21] does give an asymptotic fatigue point for general NTRU but only a concrete relation for ternary secrets. Employing the scripts provided in [DvW21], our analysis shows that when we increase the standard deviation σ , then q can be increased with the *square* of this increase before reaching the fatigue point. Concretely, for σ and ring dimension d our experiments suggest a fatigue point for modulus q given by the following expression

$$q = 0.0058 \cdot \sigma^2 \cdot d^{2.484}.$$

Note, by following a similar asymptotic analysis to that in [DvW21], we confirm that the influence of σ on the fatigue point must indeed manifest only in the leading constant and not in the exponent of d .

To demonstrate the importance of this quadratic relationship, we recompute parameters for the recent (partially) blind signature by del Pino and Katsumata [dPK22], improving its efficiency compared to the original scheme, which uses ternary secrets. Most significantly, for this work, the fatigue relation’s quadratic nature allows parameters to reach the required security level without needing to increase the ring dimension used in our voting protocol (which would significantly impact performance).

A New Lattice-Based E-Voting Design from NTRU. Our main contribution is presenting a new lattice-based e-voting protocol following the standard mix-and-decrypt framework which also supports general ballots. Our design combines our NTRU-based verifiable mix net and distributed decryption schemes. Moreover, we call on our analysis of the NTRU problem to choose fine-tuned concrete parameters. Crucially, when choosing the NTRU secret keys, we can drop the ring dimension down to 2048 from 4096 and modulus down to 59 bits from 78 bits as used in [ABGS23] whilst maintaining a 128-bit security level. This would not have been possible without the quadratic nature of the fatigue relation. We provide an efficient C++ implementation of our design.

Overall, we reduce the voting protocol’s communication by $2.5\times$ and computation by $2\times$ over [ABGS23], see Table 1 for a comparison and Section 5 for more details. It is interesting to note that, when compared to their classically secure counterparts, quantum-safe replacements typically come with a $30\times$ communication cost (e.g. ECDH vs CRYSTALS Kyber/ML-KEM). Comparing our voting scheme to ElGamal-based schemes often used in practice, we incur a cost of at most $20\times$ in ciphertext size, suggesting that our design may be approaching what can be optimally achieved.

1.2 Related Works

Lattice-Based Electronic Voting. In [ABGS23], the authors provide a verifiable mix-net and verifiable distributed decryption protocol based on BGV, showing for the first time that lattice-based electronic voting can be practical for real-world systems. We build directly upon their framework and conduct a more detailed comparison in Section 5. This work utilises the verifiable shuffle of known commitment openings by [ABG⁺21]; a building block we adopt³. del Pino et al. [dPLNS17] gives a practical scheme based on homomorphic counting, but it does not scale well for systems with more complex ballots.

A shuffle by [CMM19] was implemented in [FWK21]; however, it is less efficient than [ABGS23]. More theoretical works include [HMS21], [Str19], and [CGGI16], but none of these are efficient enough to be considered for practical deployment. Moreover, [CMM19, FWK21, HMS21] do not consider the decryption of ballots, which would heavily impact the parameters of the protocols in practice. Finally, [BHM20] gives a fast decryption mix-net, but it cannot achieve universal verifiability and is thus unsuitable for real-world elections.

NTRU Cryptanalysis. The most relevant work analysing NTRU fatigue is that of Ducas and van Woerden [DvW21]. It is important to acknowledge that this sits atop a line of work in recent years. The concurrent works [ABD16, C JL16] showed, for the first time, that NTRU security is more subtle than simply finding a notably short vector in a lattice. These works exploit the specific algebraic structure of the NTRU lattice to gain an advantage on standard lattice reduction for so-called ‘overstretched’ parameter regimes.

This work was closely followed by Kirchner and Fouque [KF17], who showed that improved attacks were, in fact, only due to the geometric existence of an unusually dense sublattice of large dimension within the NTRU lattice. Moreover, their analysis concludes that q larger than $d^{2.783+o(1)}$ already lies in the overstretched range (for ternary secrets). This bound was improved upon by the work of [DvW21] as discussed in Section 4.

1.3 Paper Organization

We begin in Section 2 by introducing some background material including notation, the NTRU, RSIS, and RLWE lattice assumptions, and the necessary building blocks used in our e-voting design; NTRU encryption, and the BDLOP commitment scheme.

In Section 3 we present our electronic voting scheme. Section 3.1 provides an overview of the well-established mix-and-decrypt framework. Section 3.2 introduces the core, passively-secure e-voting construction Π_{PVote} . This is included to aid the reader’s intuition before presenting the full scheme. Section 3.3 contains the full, actively-secure construction Π_{AVote} whose constituent algorithms are detailed in Figs. 5 and 6. This subsection also makes explicit, the implicit verifiable mixing Π_{AMix} and verifiable distributed decryption Π_{ADDec} building blocks at the core of our design. The section is rounded off with a security analysis of the Π_{AMix} and Π_{ADDec} building blocks. We dedicate Section 3.4 to providing the details of zero-knowledge proofs used in our actively secure voting scheme.

Before we can set concrete parameters for our voting scheme, Section 4 takes a necessary interlude to closely examine the NTRU assumption, providing a fatigue point relation for *general* NTRU secrets (Eq. (3)) supported by experimental data displayed in Figs. 7 and 8. This section closes with a discussion of the implications of Eq. (3) for existing NTRU-based constructions in the literature.

³We remark that in an unpublished work by Jonathan Bootle, Vadim Lyubashevsky, and Antonio Merino-Gallardo (shared over private communication), they found a flaw in the shuffle proof by Aranha et al. [ABG⁺21, ABGS23] that we build upon. However, they also propose a solution similar to the product proof by Costa, Martínez, and Morillo [CMM19], where the verifier sends two more public challenges at the start of the protocol, which are incorporated in the proofs of linearity that follows. This change is needed for soundness, but it does not impact the performance of the shuffle.

In Section 5 we show how, informed by the analysis of Section 4, one can set concrete parameters for our voting scheme. We give a sample parameter set yielding the communication costs shown in Table 3. Additionally, we provide an efficient implementation and display the resulting timings in Table 4. We close by discussing some related future research paths.

2 Preliminaries

Here we detail the essential tools employed in our constructions. We recall standard lattice results and necessary cryptographic building blocks. We begin with some notation.

Notation. For a set S and distribution D , “ $\leftarrow S$ ” and “ $\leftarrow D$ ” denote the processes of uniformly sampling from S and sampling from (or executing) D , respectively. We denote by $\text{Perm}[i]$ the set of permutations of the integers $\{1, \dots, i\}$.

Adversarial Model. We assume a static, active adversary who has full control of corrupted parties including access to their internal tapes and the ability to determine outputs. When analysing the security of our voting scheme, properties relating to ballot privacy are proven assuming all but one shuffle server and one decryption server are corrupt⁴. Properties relating to integrity (correctness) hold even if all parties are corrupt.

We assume a trusted setup whereby a single trusted entity creates the public key under which votes are encrypted and secret key shares that are passed to decryption servers. This is common in many voting systems but we note that one could use the techniques of [RST⁺22] to create a distributed key generation for NTRU.

While our work constructs and implements a voting scheme from quantum-safe assumptions, we do not claim that it is post-quantum secure. Particularly, we do not prove security in the Quantum Random Oracle Model (QROM). See [ABGS23, Appendix B] for a more detailed discussion on this distinction as it relates to the building blocks we use.

2.1 Lattices

The Ring $\mathbb{Z}[x]/(x^d + 1)$. Consider the rings $R = \mathbb{Z}[x]/\phi$ and $R_q = \mathbb{Z}_q[x]/\phi$, where $\phi = (x^d + 1)$ for d an integer power of 2 and q a prime. Elements in both rings are polynomials of degree at most $d - 1$, with those in the latter ring having coefficients between $-(q - 1)/2$ and $(q - 1)/2$. We denote elements of \mathbb{Z} and R by lower-case letters, vectors in R^k by bold lower-case letters, and matrices in $R^{(k \times \ell)}$ by bold upper-case letters. For a positive real σ , let $D_{\mathbb{Z}^d, \sigma}$ denote the discrete Gaussian distribution over \mathbb{Z}^d . To make the notation simple, we denote $a \leftarrow D_\sigma$ to mean that the coefficient vector of $a \in R_q$ is sampled from $D_{\mathbb{Z}^d, \sigma}$. For $a, b \in R$, we have that $\|ab\|_\infty \leq \|a\|_1 \cdot \|b\|_\infty$ and $\|ab\|_\infty \leq \|a\|_2 \cdot \|b\|_2$. Let S_ν denote the set of all elements $a \in R$ such that the absolute norm is $\|a\|_\infty \leq \nu$.

We use the following standard results for Gaussian vectors:

Lemma 1 (Tail Bounds [MR04, Lyu12]). *For any real $t > 0$ and $t' > 1$, we have*

$$\begin{aligned} \Pr[x \leftarrow D_{\mathbb{Z}^n, \sigma} : \|x\|_\infty > t\sigma] &< 2n \cdot 2^{-\frac{\log e}{2} \cdot t^2}, \\ \Pr[x \leftarrow D_{\mathbb{Z}^n, \sigma} : \|x\|_2 > t'\sigma\sqrt{n}] &< 2^{n \cdot (\frac{\log e}{2}(1-t'^2) + \log t')}. \end{aligned}$$

⁴In many real-world designs, each server may perform both a shuffling and decryption operation and thus privacy holds if at least one of these servers is honest

Rejection Sampling. In lattice-based cryptography in general, and in our zero-knowledge protocols in particular, we would like to output vectors $\mathbf{z} = \mathbf{y} + \mathbf{v}$ such that \mathbf{z} is independent of \mathbf{v} , and hence, \mathbf{v} is masked by the vector \mathbf{y} . Here, \mathbf{y} is sampled according to a Gaussian distribution \mathcal{N}_σ^k with standard deviation σ and we want the output vector \mathbf{z} to be from the same distribution. The procedure is shown in Figure 1.

Here, $1/M$ is the probability of success, and M is computed as

$$\max \frac{\mathcal{N}_\sigma^k(\mathbf{z})}{\mathcal{N}_{\mathbf{v},\sigma}^k(\mathbf{z})} \leq \exp \left[\frac{24\sigma\|\mathbf{v}\|_2 + \|\mathbf{v}\|_2^2}{2\sigma^2} \right] = M \quad (1)$$

where we use the tail bound from Lemma 1, saying that $|\langle \mathbf{z}, \mathbf{v} \rangle| < 12\sigma\|\mathbf{v}\|_2$ with probability at least $1 - 2^{-100}$. Hence, for $\sigma = 11\|\mathbf{v}\|_2$, we get $M \approx 3$. This is the standard way to choose parameters, see e.g. [BLS19]. However, if the procedure is only done once for the vector \mathbf{v} , we can decrease the parameters slightly, to the cost of leaking only one bit of information about \mathbf{v} from given \mathbf{z} .

In [LNS21], Lyubashevsky et al. suggest to require that $\langle \mathbf{z}, \mathbf{v} \rangle \geq 0$, and hence, we can set $M = \exp(\|\mathbf{v}\|_2/2\sigma^2)$. Then, for $\sigma = 0.675\|\mathbf{v}\|_2$, we get $M \approx 3$. In Fig. 1, we use the pre-determined bit b to denote if we only use \mathbf{v} once or not, with the effect of rejecting about half of the vectors before the sampling of uniform value μ in the case $b = 1$ but allowing a smaller standard deviation.

Rej($\mathbf{z}, \mathbf{v}, b, M, \sigma$)

1. if $b = 1$ and $\langle \mathbf{z}, \mathbf{v} \rangle < 0$, **return 0**
2. $\mu \xleftarrow{\$} [0, 1)$
3. if $\mu > \frac{1}{M} \cdot \exp \left[\frac{-2\langle \mathbf{z}, \mathbf{v} \rangle + \|\mathbf{v}\|_2^2}{2\sigma^2} \right]$, **return 0**
4. **return 1**

Figure 1: Rejection Sampling.

The NTRU Problem. We give the historical presentation of the NTRU problem as it is more convenient for our analysis in Section 4. We note that some works refer to this problem as the ‘search/decisional short polynomial ratio’ problem [LTV12, SXY18]. Furthermore, one can consider the so-called ‘module’ NTRU problem [CKKS19, CPS⁺20], which considers the ratio of matrices of polynomials \mathbf{F} and \mathbf{G} . Our analysis and applications can naturally be extended to the module setting, so for ease of presentation, we use the basic (polynomial) NTRU formulation [HPS98].

Definition 1 (Search/Decision NTRU). Let $q > 2$ be a prime, d be the ring dimension, and $D_{\sigma_{\text{NTRU}}}$ be a distribution over R_q . Sampling $(f, g) \leftarrow D_{\sigma_{\text{NTRU}}}^2$ with rejection if f is not invertible in R_q , define $h = g/f \in R_q$. The search-NTRU $_{q,d,\sigma_{\text{NTRU}},t}$ problem is, given h , to recover any pair (f', g') such that $h = g'/f' \in R_q$ and $\|f', g'\|_2 \leq t \cdot \sigma_{\text{NTRU}}$. The decision-NTRU $_{q,d,\sigma_{\text{NTRU}}}$ problem is, given h , to decide if h is computed as $h = g/f$ for $(f, g) \leftarrow D_{\sigma_{\text{NTRU}}}^2$ or if h is sampled uniformly from R_q .

The RLWE and RSIS Problems. We define the standard lattice-hardness problems over rings [Ajt96, Reg05, LPR10].

Definition 2 (Ring Learning with Errors). Let $q > 2$ be a prime, d be the ring dimension, $D_{\sigma_{\text{RLWE}}}$ be a distribution over R_q , and \mathcal{A} a PPT algorithm that makes at most Q oracle

queries. Then the advantage of \mathcal{A} in solving the ring learning with errors $\text{RLWE}_{d,q,Q,\sigma_{\text{RLWE}}}$ problem is defined as

$$\text{Adv}_{d,q,Q,\sigma_{\text{RLWE}}}^{\text{RLWE}}(\mathcal{A}) = \left| \Pr[\mathcal{A}^{\mathcal{O}_{\text{RLWE}}}(d, q, \sigma_{\text{RLWE}}) \rightarrow 1] - \Pr[\mathcal{A}^{\mathcal{O}_{\mathfrak{s}}}(d, q, \sigma_{\text{RLWE}}) \rightarrow 1] \right|,$$

where oracles $\mathcal{O}_{\text{RLWE}}$ and $\mathcal{O}_{\mathfrak{s}}$ are defined as

- $\mathcal{O}_{\text{RLWE}}$: Samples $a \leftarrow R_q$, $(s_1, s_2) \leftarrow D_{\sigma_{\text{RLWE}}}^2$, and then output $(a, as_1 + s_2)$;
- $\mathcal{O}_{\mathfrak{s}}$: Samples $(a, b) \leftarrow R_q \times R_q$, and then output (a, b) .

Definition 3 (Ring Short Integer Solutions). Let $q > 2$ be a prime, d be the ring dimension, $\|\cdot\|$ a norm, and $\beta \in \mathbb{R}^+$ a positive integer. The $\text{RSIS}_{d,q,\beta}$ problem is, given a uniformly random $a \in R_q$, find $s_1, s_2 \in R_q$ such that $as_1 + s_2 = 0 \in R_q$ and $0 \leq \|s_1, s_2\| \leq \beta$.

2.2 Building Blocks

NTRU Encryption. In this work, we will use the provably secure variant of the NTRU cryptosystem first presented by Steinfeld and Stehlé in [SS11]. This scheme relies on the hardness of both the RLWE and NTRU assumptions. Note we make two minor modifications to ensure perfectly correct decryption: (1) encryption randomness is sampled from a bounded distribution, and (2) the secret keys f and g are rejected unless their ℓ_2 norm is below a given bound. When sampled accordingly, this limitation has only a negligible effect on the completion probability of the key generation algorithm and the entropy of resulting keys.

Setup. Let $p \ll q$ be primes and d a power of two which define the rings R_p and R_q . Messages lie in R_p . Let $\sigma_{\text{NTRU}} \in \mathbb{R}^+$ and $D_{\sigma_{\text{NTRU}}}$ a discrete Gaussian distribution over R with standard deviation σ_{NTRU} , $t \in (1, 2]$ and $\nu \in \mathbb{N}$. Let the setup parameters be $\text{sp} = (d, p, q, \sigma_{\text{NTRU}}, t, \nu)$. The encryption scheme is described in Fig. 2.

Key Generation $\text{KeyGen}_{\text{NTRU}}(\text{sp})$. Given input $\text{sp} = (d, p, q, \sigma_{\text{NTRU}}, t, \nu)$:

1. $f, g \leftarrow D_{\sigma_{\text{NTRU}}}$; if $f \notin R_q^\times$ or $f \not\equiv 1 \in R_p$, resample.
2. If $\|f\|_2, \|g\|_2 > t \cdot \sqrt{d} \cdot \sigma_{\text{NTRU}}$, restart.
3. Return $\text{sk} = f$, $\text{pk} = h := g/f \in R_q$.

Encryption $\text{Enc}_{\text{NTRU}}(m, \text{pk})$. Given message $m \in R_p$ and public key $\text{pk} = h$:

1. Sample encryption randomness $s, e \leftarrow S_\nu$.
2. Return ciphertext $c = p \cdot (hs + e) + m \in R_q$.

Decryption $\text{Dec}_{\text{NTRU}}(c, \text{sk})$. Given ciphertext c and key $\text{sk} = f$:

1. Return message $m = (f \cdot c \bmod q) \bmod p$.
-

Figure 2: Adapted NTRUEncrypt [SS11].

Lemma 2 (NTRUEncrypt Security). Let $p \cdot d \cdot t \cdot \sigma_{\text{NTRU}}(2\nu + 1/2) < \lfloor q/2 \rfloor$. Then the encryption scheme in Fig. 2 is (perfectly) correct. Moreover, assuming the hardness of the $\text{NTRU}_{q,d,\sigma_{\text{NTRU}}}$ and $\text{RLWE}_{d,q,Q,\chi}$ problems, the scheme is IND-CPA secure.

The proof of Lemma 2 closely follows the one given in [SS11] and a similar scheme (also assuming the NTRU assumption) is presented in [LTV12]. We provide a sketch proof of our scheme for completeness.

Proof. We first consider correctness. Examine the decryption operation $(f \cdot c \bmod q) \bmod p$. For messages to be recovered correctly, we need that no modular reduction occur modulo q inside the bracket. That is, for perfectly correct decryption we require $\|f \cdot c\|_\infty \leq \lfloor q/2 \rfloor$. We have that

$$\begin{aligned} \|f \cdot c\|_\infty &\leq p\|gs\|_\infty + p\|fe\|_\infty + \|fm\|_\infty \\ &\leq p \cdot \|g\|_1 \cdot \nu + p \cdot \|f\|_1 \cdot \nu + \|f\|_2 \cdot \sqrt{d} \cdot \|m\|_\infty \\ &\leq p \cdot \sqrt{d} \cdot \|g\|_2 \cdot \nu + p \cdot \sqrt{d} \cdot \|f\|_2 \cdot \nu + \|f\|_2 \cdot \sqrt{d} \cdot \|m\|_\infty \\ &\leq p \cdot dt \cdot \sigma_{\text{NTRU}} \cdot (2\nu + 1/2), \end{aligned}$$

as assumed in the statement of the lemma. IND-CPA security is proven using a hybrid argument, in two steps:

1. The hardness of the $\text{NTRU}_{q,d,\sigma_{\text{NTRU}}}$ allows the simulator to change the public key $h = g/f \in R_q$ to a uniformly sampled h .
2. The simulator next changes the challenge ciphertext $c^* = p \cdot (hs + e) + m \in R_q$ to $c^* = u + m \in R_q$, where u is sampled uniformly at random from R_q . This change is indistinguishable by the $\text{RLWE}_{d,q,1,\chi}$ assumption.

In this final hybrid, the advantage of the adversary is exactly $1/2$ since c^* is uniform over R_q independent of the message m . Thus, we have

$$\text{Adv}^{\text{IND-CPA}}(\mathcal{A}) \leq \text{Adv}_{q,d,\sigma_{\text{NTRU}},t}^{\text{decision-NTRU}}(\mathcal{A}) + \text{Adv}_{d,q,1,\chi}^{\text{RLWE}}(\mathcal{A})$$

□

The BDLOP Commitment Scheme. Here we recall the BDLOP commitment scheme from [BDL⁺18]. For simplicity, we present the scheme instantiated over rings instead of modules, committing to only one ring element at a time. The scheme is parametrised by $B_{\text{Com}}, \sigma_{\text{Com}} \in \mathbb{R}^+$, and challenge space \mathcal{C} whose difference set is $\vec{\mathcal{C}}$.

Setup : Samples uniformly random a_1, a_2, a_3 from R_q and outputs the public commitment key $\text{pk}_{\mathcal{C}}$ defined as:

$$\text{pk}_{\mathcal{C}} = \begin{bmatrix} \vec{a}_1 & 0 \\ \vec{a}_2 & 1 \end{bmatrix} = \begin{bmatrix} 1 & a_1 & a_2 & 0 \\ 0 & 1 & a_3 & 1 \end{bmatrix}.$$

Com($\text{pk}_{\mathcal{C}}, x$) : On input a public commitment key $\text{pk}_{\mathcal{C}}$ and an element x in R_q , samples a vector $\vec{r} \in R_q^3$ such that $\|\vec{r}\|_\infty \leq B_{\text{Com}}$, and computes the commitment as:

$$\text{com} = \begin{bmatrix} \vec{c}_1 \\ \vec{c}_2 \end{bmatrix} = \begin{bmatrix} 1 & a_1 & a_2 & 0 \\ 0 & 1 & a_3 & 1 \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \\ r_3 \\ x \end{bmatrix} = \llbracket x \rrbracket.$$

It outputs the commitment com and the opening $d = (x, \vec{r}, 1)$.

$\text{Open}(\text{pk}_C, \text{com}, d)$: On input a public commitment key pk_C , the commitment com and the opening $d = (x, \vec{r}, f)$ where $f \in \bar{\mathcal{C}}$. It verifies:

$$f \cdot \text{com} \stackrel{?}{=} \begin{bmatrix} 1 & a_1 & a_2 & 0 \\ 0 & 1 & a_3 & 1 \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \\ r_3 \\ f \cdot x \end{bmatrix},$$

and $\forall i \in [3]: \|r_i\|_2 \stackrel{?}{\leq} 4 \cdot \sigma_{\text{Com}} \sqrt{d}$. It outputs 1 if the relations hold and 0 otherwise.

Lemma 3 (Commitment Security [BDL⁺18]). *The BDLOP commitment scheme is hiding if the RLWE problem is hard for vectors of ℓ_∞ norm B_{Com} over a lattice of dimension $2 \cdot d$ and binding if the RSIS problem is hard for vectors of ℓ_2 norm $16\sigma_{\text{Com}}\sqrt{\kappa d}$ over a lattice of dimension $2 \cdot d$.*

3 The Voting Scheme

A *cryptographic voting scheme* is usually defined in terms of the algorithms for election setup, casting ballots, and counting cast ballots. We need algorithms for shuffling and distributed decryption to model the counting process accurately. To make such a scheme verifiable (actively secure), we also need a mechanism to enforce the fact that the encryption, shuffling, and decryption algorithms are computed honestly. This section presents an NTRU-based voting protocol in the well-established ‘mix-and-decrypt’ paradigm, comprising new verifiable distributed decryption and mix-net protocols. We refer to Chapter 14 in the book by Gjøsteen [Gjø22] and Appendix H in the full version of Aranha et al. [ABGS22] for a more thorough description of e-voting schemes.

3.1 Voting Scheme Overview

Setup Phase. A trusted party runs the key generation algorithm for the PKE scheme with distributed decryption. In this work, we will assume a trusted key generation and leave the design of a distributed key generation algorithm for NTRU to future work, as a trusted setup is typical for many voting schemes⁵. The generated public parameters sp are given to every participant, while the decryption key shares dk_j are distributed amongst the decryption servers.

Casting Phase. Each voter instructs their voting device to cast their chosen ballot. The device encrypts the ballot under the public key pk to create a ciphertext c , and it computes a *ballot proof*. The standard way to do this is to use a verifiable encryption scheme such as the one presented in [LNP22], proving that the submitted ciphertext contains a genuine ballot in zero-knowledge. We remark that it is indeed required for the voters to prove plaintext knowledge to ensure security of the voting scheme, to prevent attacks such as copying and re-randomizing ballots from other voters.

Counting Phase. This is divided into three sequential processes. First, encrypted ballots are passed through a series of shuffle servers.

The ξ_1 shuffle servers $\mathcal{S}_1, \dots, \mathcal{S}_{\xi_1}$ consecutively run the shuffle algorithm of the set of encrypted ballots $\{c_i^{(k-1)}\}$, passing the shuffled and re-encrypted ballots $\{c_i^{(k)}\}$ to the next shuffle server. They also generate a shuffle proof which anyone can verify. We may refer to this whole shuffle process as the *mix-net*.

⁵One could adapt the techniques of [RST⁺22] to create a distributed key generation for NTRU.

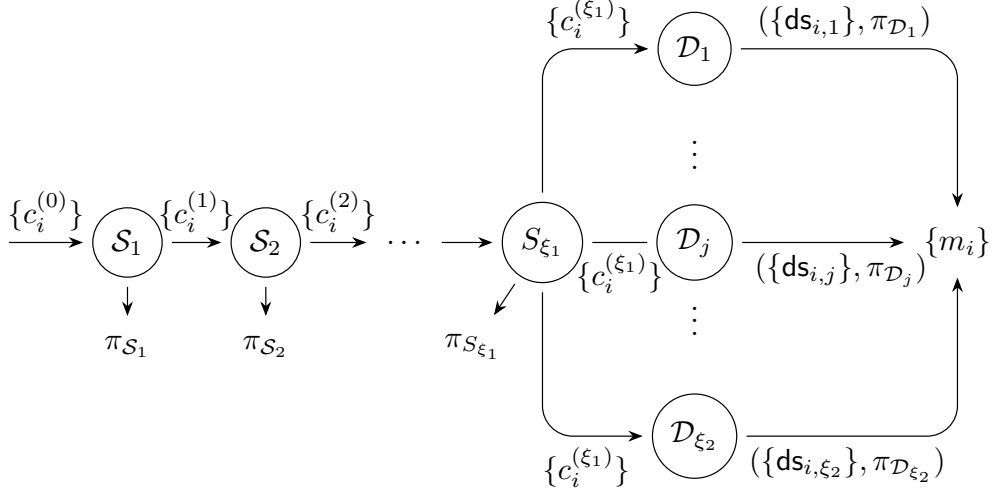


Figure 3: The voting protocol with verifiable mix-net and distributed decryption, adapted from [ABGS23, Figure 1] with shuffle servers \mathcal{S}_i and decryption servers \mathcal{D}_j .

Each of the ξ_2 decryption servers \mathcal{D}_j receives the output of each shuffle server and verifies the corresponding shuffle proofs. Only after verifying each proof does a decryption server begin decryption. \mathcal{D}_j then computes a set of partial decryption shares $\{\text{ds}_{ij}\}$, one for each of the ciphertexts. Finally, it creates a proof of decryption to guarantee that it computed its shares correctly. Each decryption server passes its shares to the combining algorithm **Comb**.

The **Comb** algorithm performs the task of recovering the ballots. The **Comb** algorithm verifies the decryption proofs after receiving all decryption shares from decryption servers. If all decryption proofs are verified, the ballots are recovered by combining the decryption shares.

A schematic of these processes, in the malicious setting, is shown in Fig. 3. This figure is adapted from [ABGS23] and shows the voting protocol beginning with input of a set of encrypted ballots and finishing with a set of ballots in plaintext. We note that some works consider an auditor who verifies the processes at each step by checking the proofs provided. This is a stylistic design choice. For the purposes of this paper, it is helpful to think of the proofs as providing verifiability of each phase by any third party and by the component servers before carrying out their roles.

3.2 Passively Secure Scheme

Here we present our passively secure voting scheme. Whilst our ultimate goal is to give a verifiable (actively secure) voting scheme, we first isolate the core, passively secure skeleton for clarity of presentation. We begin by defining the algorithms and syntax of this construction.

Definition 4 (Passively Secure Voting Scheme). Let τ be the number of voters, ξ_1 the number of shuffle servers, and ξ_2 the number of decryption servers. A passively secure cryptographic voting scheme Π_{Vote} consists of five algorithms (**KeyGen**, **Cast**, **Mix**, **DDec**, **Comb**).

KeyGen(sp) \rightarrow ($\text{pk}, \text{sk}, \{\text{dk}_j\}_{j \in [\xi_2]}$): On input setup parameters sp , it returns a public encryption key pk , a secret key sk and a set of ξ_2 secret decryption key shares $\{\text{dk}_j\}_{j \in [\xi_2]}$.

$\text{Cast}(\text{pk}, v) \rightarrow c$: On input a public key pk and vote v it returns an encrypted ballot c .

$\text{Mix}(\{c_i\}_{i \in [\tau]}) \rightarrow \{\hat{c}_i\}_{i \in [\tau]}$: On input a set of encrypted ballots $\{c_i\}_{i \in [\tau]}$ it returns a set of encrypted ballots $\{\hat{c}_i\}_{i \in [\tau]}$.

$\text{DDec}_j(\{c_i\}_{i \in [\tau]}, \text{dk}_j) \rightarrow \{\text{ds}_{i,j}\}$: On input a set of encrypted ballots $\{c_i\}_{i \in [\tau]}$ and a decryption key dk_j , it returns a set of decryption shares $\text{ds}_j = \{\text{ds}_{i,j}\}_{i \in [\tau]}$.

$\text{Comb}(\{c_i\}_{i \in [\tau]}, \{\text{ds}_j\}_{j \in [\xi_2]}) \rightarrow \{v\}_{i \in [\tau]}$: On input a set of encrypted ballots $\{c_i\}_{i \in [\tau]}$ and a set of decryption shares $\{\text{ds}_j\}_{j \in [\xi_2]}$, it outputs a set of votes $\{v\}_{i \in [\tau]}$.

We instantiate the algorithms, present our passively secure voting scheme and give an overview in Fig. 4.

Setup. Let $p \ll q$ be primes and d a power of two which define the rings R_p and R_q . Votes lie in R_p . Let $\sigma_{\text{NTRU}}, B_{\text{Dec}}, B_{\text{Drown}} \in \mathbb{R}^+$, $t \in (1, 2]$, and $\nu, \tau, \xi_1, \xi_2 \in \mathbb{N}$. Let $\text{sp} = (d, p, q, \sigma_{\text{NTRU}}, t, \nu, \tau, \xi_1, \xi_2)$.

3.3 Actively secure scheme

We present our actively secure (verifiable) voting scheme. This can be seen as a natural extension of the passive protocol Π_{Vote} by adding verifiability to the mixing and distributed decryption processes. This is done by applying the zero-knowledge proofs of Section 3.4 so that the outputs of Mix_A and DDec_A , now include a proof of shuffle π_S and a proof of decryption π_D respectively. Note the use of ‘A’ in algorithm/protocol suffixes to indicate the actively-secure variant (as opposed to those using ‘P’ i.e. passively-secure variants).

Now, any third party can verify that the mixing and distributed decryption processes were carried out correctly without compromising the privacy or integrity of the voting system. As usual, we assume a trusted dealer for key generation and leave the construction of an NTRU-based distributed key generation for other applications to future work.

This construction implicitly defines a verifiable mixing with verifiable distributed decryption from NTRU. We consider these to be of independent interest and give an overview as stand-alone protocols.

Verifiable Mixing. Our aim here is, given a set of input ciphertexts, to generate a new set of ciphertexts that decrypts to the same set of plaintexts. Crucially, input-output ciphertext correspondence must be obscured. Additionally, we would like any third party to verify that this process has been performed correctly without compromising the privacy of the mix.

For this, we apply the proof of [ABG⁺21], which allows one to prove a shuffle of openings of the lattice commitments in Section 2.2. We denote this proof system Π_{Shuf} . Since NTRU ciphertexts only contain a single element, we can import their scheme without modification where the committed messages are ciphertexts. We also employ the Π_{SMALL} proof systems described in Section 3.4 to prove that the new ciphertext noise is sufficiently bounded. At a high level, our verifiable mixing of NTRU ciphertexts c_1, \dots, c_τ re-randomises the input ciphertexts and then permutes their order where the permutation is only known to the shuffle server:

1. The mixing server creates encryptions c'_1, \dots, c'_τ of 0 and commits to these as $\llbracket c'_i \rrbracket$ for each $i \in [\tau]$. Run the Π_{SMALL} protocol to prove that each committed ciphertext is honestly computed.
2. Adding the original ciphertexts c_i to these commitments homomorphically yields commitments $\llbracket \hat{c}_i \rrbracket$ to ciphertexts with the same plaintext as in c_1, \dots, c_τ , now with fresh randomness.

KeyGen(sp). On input system parameters sp :

1. $(\text{sk} = f, \text{pk} = h) \leftarrow \text{KeyGen}_{\text{NTRU}}(d, p, q, \sigma_{\text{NTRU}}, t)$.
2. For $j \in [\xi_2 - 1]$, $\text{dk}_j \leftarrow \mathcal{U}(R_q)$ and set $\text{dk}_{\xi_2} = \text{sk} - \sum_{j=1}^{\xi_2-1} \text{dk}_j \pmod q$.
3. Return (pk, sk) and key shares $\{\text{dk}_j\}_{j \in [\xi_2]}$.

Cast(pk, v). On input the public key pk and a vote $v \in R_p$:

1. Compute $c = \text{Enc}_{\text{NTRU}}(\text{pk}, v)$.
2. Return encrypted ballot c .

Mix($\{c_i\}_{i \in [\tau]}$). On input encrypted ballots $\{c_i\}_{i \in [\tau]}$:

1. For each $i \in [\tau]$, compute $c'_i = \text{Enc}_{\text{NTRU}}(\text{pk}, 0)$.
2. For each $i \in [\tau]$, compute $\hat{c}_i = c_i + c'_i \pmod q$.
3. Sample a random permutation $\pi \leftarrow \text{Perm}[\tau]$.
4. Return re-encrypted ballots $\{\hat{c}_{\pi(i)}\}_{i \in [\tau]}$.

DDec $_j(\{c_i\}_{i \in \tau}, \text{dk}_j)$. On input a set of encrypted ballots $\{c_i\}_{i \in \tau}$ and a decryption key share dk_j :

1. For each $i \in [\tau]$, sample $E_{ij} \leftarrow S_{B_{\text{Drown}}}$ and compute share $\text{ds}_{ij} = \text{dk}_j \cdot c_i + p \cdot E_{ij} \pmod q$.
2. Return the set of decryption shares $\text{ds}_j = \{\text{ds}_{ij}\}_{i \in [\tau]}$.

Comb($\{c_i\}_{i \in [\tau]}, \{\text{ds}_j\}_{j \in [\xi_2]}$). On input encrypted ballots $\{c_i\}_{i \in [\tau]}$ and decryption shares $\{\text{ds}_j = \{\text{ds}_{ij}\}_{i \in [\tau]}\}_{j \in [\xi_2]}$:

1. For each $i \in [\tau]$, $v_i = \left(\sum_{j \in [\xi_2]} \text{ds}_{ij} \pmod q \right) \pmod p$.
 2. Return the set of votes $\{v_i\}_{i \in [\tau]}$.
-

Figure 4: The passively-secure voting scheme Π_{PVote} .

3. The server now reveals the openings \hat{c}_i in a randomly permuted order and runs the Π_{Shuf} protocol to prove that these are indeed a permutation of the correct openings of the commitments.

We note that verification of the shuffle proof should be done before any ballot decryption begins. This can be seen as a first step of the DDec algorithm or as part of a global verification process carried out by an auditor. For simplicity of presentation and since this is covered in [ABGS23], we omit this from the full protocol.

Verifiable Distributed Decryption. Our aim here is, given a set of input ciphertexts, to generate a set of decryption shares to extract the encrypted plaintexts when all the shares are combined. Furthermore, each decryptor must prove they decrypted their decryption share correctly using their secret key share. Therefore, in the active setting, the public key contains a commitment $\llbracket \text{dk}_j \rrbracket$ to each secret key share dk_j , and each decryptor holds an opening to precisely one of the commitments. The verifiable distributed decryption

protocol works as follows:

1. For each $i \in [\tau]$, the decryptor samples a noise value $E_{ij} \leftarrow S_{B_{\text{Drown}}}$, computes a decryption share $\text{ds}_{ij} = \text{dk}_j \cdot c_i + p \cdot E_{ij}$ and commits to the noise as $\llbracket E_{ij} \rrbracket$.
2. For each $i \in [\tau]$, it uses the Π_{Lin} protocol to prove that the linear decryption equation above is computed honestly with respect to $\llbracket \text{dk}_j \rrbracket$ and $\llbracket E_{ij} \rrbracket$.
3. For each $i \in [\tau]$, it uses the Π_{Bnd} protocol to prove that $\llbracket E_{ij} \rrbracket$ is an honestly created commitment and the committed value is bounded by B_{Drown} .

We wish to emphasise the importance of our Π_{Bnd} proof system. Π_{Bnd} is a modification of the ZKP by Bootle et al. in [BLNS21]). Crucially, it proves a tight bound on the size of the noise drowning term E_{ij} . To our knowledge, this is the first exact ZKP of a ‘large’ secret vector for lattice relations and may be of independent interest. Though a more costly proof, proving an exact bound on E_{ij} will lead to more efficient global parameters in Section 5. We remark that one might consider using the LNP proof system of [LNP22] for Π_{Bnd} . However, proofs arising from the LNP framework scale linearly with the witness while our design saves a square root factor by comparison. Furthermore, LNP is designed for proving 2-norms directly and one would need to perform some witness encoding to prove infinity norms, further reducing efficiency.

Setup. Let $p \ll q$ be primes and d a power of two which define the rings R_p and R_q . Votes lie in R_p . Let $\sigma_{\text{NTRU}}, B_{\text{Dec}}, B_{\text{Drown}}, B_{\text{Com}}, B_{\text{Small}} \in \mathbb{R}^+$, $t \in (1, 2]$, and $\nu, \tau, l, \xi_1, \xi_2 \in \mathbb{N}$. Let $\text{sp} = (d, p, q, \sigma_{\text{NTRU}}, t, \nu, \tau, l, \xi_1, \xi_2, B_{\text{Dec}}, B_{\text{Drown}}, B_{\text{Com}})$.

3.4 Zero-Knowledge Proofs

We present the proof systems needed in the actively secure voting protocol. We wish to highlight, in particular, how we adapt the amortised proof of boundedness Π_{BND} as compared to previous works [ABGS23]. The crucial observation here is that whilst we get slightly larger proofs there, the *exact* guarantees provided by the proof allow for better parameters to be chosen for the overall voting scheme. This leads to a net reduction in communication costs.

Amortized Proof of Shortness. In Step 2 of the shuffle, we use Π_{Small} to prove that we have committed to well-formed encryptions of zero. The protocol Π_{Small} produces a proof that a batch of equations $\vec{A}\vec{s}_i = \vec{t}_i$ for $i \in [\ell]$ is satisfied for a set of secret vectors \vec{s}_i with ℓ_∞ norm bounded by ν . The exact relation for the proof system is:

$$\mathcal{R}_{\text{SMALL}} := \left\{ (x, w) \mid \begin{array}{l} x := (\text{pk}_C, \{\text{com}_i\}_{i \in [\ell]}) \wedge w := (\{d_i = (u_i, \vec{r}_i, f_i)\}_{i \in [\ell]}) : \\ \forall i \in [\ell] : \|u_i\|_\infty \leq \nu \wedge \text{Open}(\text{pk}_C, \text{com}_i, d_i) \end{array} \right\}.$$

The proof of shortness π_{SMALL} is quite involved, combining error-correcting codes, Merkle trees, Lagrange interpolation and proximity testing, and we refer to [ABGS23] for details. For batch size ℓ of ternary secret vectors, the proof size is given in [ABGS23, Eq. (1)] as

$$(3vd + (3\ell + 2)\eta) \log_2 q + 2\lambda\eta(1 + \log_2 \gamma) \text{ bits},$$

using an $[\gamma, \mu, \iota]$ Reed-Solomon Code with code-length γ , message length μ and minimal distance ι where $\mu = d(k + 2) + \eta \leq \gamma < q$ for encoding randomness of length η . λ is the security parameter. The soundness error of the proof is given as

$$2 \cdot \max \left\{ 2 \left(\frac{\mu'}{\gamma - \eta} \right)^\eta, \frac{1}{q - \ell} + \left(1 - \frac{\mu' - \mu}{6\gamma} \right)^\eta, 2 \cdot \left(1 - \frac{2(\mu' - \mu)}{3\gamma} \right)^\eta, \frac{18\ell}{q - \ell} \right\},$$

for a choice of message length μ' such that $\mu \leq \mu' \leq \gamma < q$.

KeyGen_A(sp). On input system parameters sp:

1. Run $((\mathbf{pk}, \mathbf{sk}), \{\mathbf{dk}_j\}_{j \in [\xi_2]}) \leftarrow \text{KeyGen}(\text{sp})$.
2. For $j \in [\xi_2]$, compute the commitments and openings $(\llbracket \mathbf{dk}_j \rrbracket, \vec{r}_{\mathbf{dk}_j}) \leftarrow \text{Com}(\mathbf{dk}_j)$.
3. Return $\mathbf{pk}_A = (\mathbf{pk}, \llbracket \mathbf{dk}_1 \rrbracket, \dots, \llbracket \mathbf{dk}_{\xi_2} \rrbracket)$, $\mathbf{sk}_A = \mathbf{sk}$, and key shares $\{\mathbf{dk}_{A,j} = (\mathbf{dk}_j, \vec{r}_{\mathbf{dk}_j})\}_{j \in [\xi_2]}$.

Cast_A(pk_A, v). On input a public key \mathbf{pk}_A and a vote v in R_p , retrieving \mathbf{pk} from \mathbf{pk}_A :

1. Return $c \leftarrow \text{Cast}(\mathbf{pk}, v)$.

Mix_A({c_i}_{i ∈ [τ]}). On input a set of encrypted ballots $\{c_i\}_{i \in [\tau]}$:

1. For $i \in [\tau]$, compute $c'_i \leftarrow \text{Enc}_{\text{NTRU}}(\mathbf{pk}, 0)$ using encryption randomness (s'_i, e'_i) .
2. For $i \in [\tau]$, commit to c'_i as $\text{com}'_i := \llbracket c'_i \rrbracket \leftarrow \text{Com}(\mathbf{pk}_C, c'_i)$ where $\vec{r}_{c'_i}$ is the commitment randomness used. Then denoting

$$\mathbf{A}_M = \begin{bmatrix} 1 & a_{1,1} & a_{1,2} & 0 & 0 \\ 0 & 1 & a_{2,2} & ph & p \end{bmatrix},$$

and $\mathbf{s}_{c'_i} = [\vec{r}_{c'_i}, s'_i, e'_i]^T$ compute $\pi_{\text{Small}_i} \leftarrow \Pi_{\text{Small}}$ for matrix \mathbf{A}_M , input vector $\mathbf{s}_{c'_i}$, targets com'_i , and bound B_{Small} . Set $\pi_{\text{Small}} := \{\pi_{\text{Small}_i}\}_{i \in [\tau]}$.

3. For $i \in [\tau]$, compute $\hat{c}_i = c_i + c'_i$. Sample $\pi \leftarrow \text{Perm}([\tau])$, and compute $\pi_{\text{Shuf}} \leftarrow \Pi_{\text{Shuf}}$ with input commitments $\{\llbracket \hat{c}_i \rrbracket\}_{i \in [\tau]}$, randomness $\{\vec{r}_{c'_i}\}_{i \in [\tau]}$, ciphertexts $\{\hat{c}_i\}_{i \in [\tau]}$, and permuted ciphertexts $\{\hat{c}_{\pi(i)}\}_{i \in [\tau]}$.
 4. Return $(\{\hat{c}_{\pi(i)}\}_{i \in [\tau]}, \pi_S)$, where $\pi_S = (\{\text{com}'_i\}_{i \in [\tau]}, \pi_{\text{Small}}, \pi_{\text{Shuf}})$.
-

Figure 5: Π_{AVote} key generation, casting, and shuffle.

Proof of Shuffle. In Step 3 of the shuffle, we use Π_{Shuf} to prove that a set of committed values is a permutation of public values.

The committed values will be the \hat{c}_i values in our context. The verifier can construct these from the c_i and the $\llbracket c'_i \rrbracket$. The public values are the \hat{c}_i . The proof then convinces the verifier that output ciphertexts are a genuine permutation of the re-randomized input ciphertexts. The exact relation for the proof system is:

$$\mathcal{R}_{\text{SHUF}} := \left\{ (x, w) \mid \begin{array}{l} x := (\{\{\text{com}_i, \bar{u}_i\}_{i \in [\tau]}\} \wedge w := (\{d_i = (u_i, \vec{r}_i, f_i)\}_{i \in [\tau]}, \rho) : \\ \rho \in \text{Perm}[\tau] \wedge \forall i \in [\tau] : u_i = \bar{u}_{\rho(i)} \wedge \text{Open}(\mathbf{pk}_C, \text{com}_i, d_i) \end{array} \right\}.$$

The proof of shuffle π_{SHUF} is computed as follows [ABG⁺21, Section 4]:

1. Hash the statement to get a uniform value and then convert all commitments and messages to u'_i and \bar{u}'_i (the commitments are additionally homomorphic).
2. For all $i \in [\tau-1]$, sample random values θ_i and commit to random linear combinations of the form $\llbracket D_i \rrbracket = \llbracket \theta_{i-1} \cdot u'_i + \theta_i \cdot \bar{u}'_i \rrbracket$ (where $\theta_0 = \theta_\tau = 0$).
3. Hash the commitments to get a uniform challenge β . Then for all $i \in [\tau]$ compute s_i to solve the linear system for β .
4. For all $i \in [\tau]$, compute proofs of linearity for the commitment equations of the form $\llbracket D_i \rrbracket = s_{i-1} \llbracket u'_i \rrbracket + s_i \cdot \bar{u}'_i$ (where $s_0 = \beta$ and $s_\tau = (-1)^\tau \beta$).

$\text{DDec}_{A,j}(\{c_i\}_{i \in [\tau]}, \text{dk}_{A,j})$. On input a set of ciphertexts $\{c_i\}_{i \in [\tau]}$ and decryption key share $\text{dk}_{A,j} = (\text{dk}_j, \vec{r}_{\text{dk}_j})$:

1. For $i \in [\tau]$, sample $E_{ij} \leftarrow S_{B_{\text{Drown}}}$, and compute $\text{ds}_{ij} = \text{dk}_j \cdot c_i + p \cdot E_{ij}$.
2. For $i \in [\tau]$, compute $(\llbracket E_{ij} \rrbracket, \vec{r}_{E_{ij}}) \leftarrow \text{Com}(E_{ij}, \text{pk}_C)$ and use the Π_{Lin} protocol to compute a proof $\pi_{\text{Lin}_{ij}}$ for the linear relation $\text{ds}_{ij} = \text{dk}_j \cdot c_i + p \cdot E_{ij}$.
3. Apply the amortized proof of boundedness Π_{Bnd} , to create a proof π_{Bnd} that, for all $i \in [\tau]$, $\|E_{ij}\|_\infty \leq B_{\text{Drown}}$ and $\|\vec{r}_{E_{ij}}\|_\infty \leq B_{\text{Com}}$.
4. Return $\text{ds}_j := (\{\text{ds}_{ij}\}_{i \in [\tau]}, \pi_{\mathcal{D}})$, where $\pi_{\mathcal{D}} = (\{\llbracket E_{ij} \rrbracket\}_{i \in [\tau]}, \{\pi_{\text{Lin}_{ij}}\}_{i \in [\tau]}, \pi_{\text{Bnd}})$.

$\text{Comb}_A(\{c_i\}_{i \in [\tau]}, \{\text{ds}_j\}_{j \in [\xi_2]})$. On input encrypted ballots $\{c_i\}_{i \in [\tau]}$ and decryption shares $\{\text{ds}_j\}_{j \in [\xi_2]}$:

1. Parse ds_j as $(\{\text{ds}_{ij}\}_{i \in [\tau]}, \pi_{\mathcal{D}_j})$, and verify the proofs $\pi_{\text{Lin}_{ij}}$ and $\pi_{\text{Bnd},ij}$, returning \perp if either fails to verify.
2. Compute

$$v_i = \left(\sum_{j \in [\xi_2]} \text{ds}_{ij} \pmod{q} \right) \pmod{p}.$$

3. Return the set of votes $\{v_i\}_{i \in [\tau]}$.
-

Figure 6: Π_{AVote} distributed decryption and combining.

The verifier accepts if all proofs of linearity are valid. This proof π_{SHUF} consists of one ring element, one commitment and one proof of linearity per shuffled element. Using the proof of linearity π_{LIN} from above, the size of π_{SHUF} is $\tau d(2k \log_2(4\sigma_{\text{LIN}}) + 3 \log_2 q)$ bits.

We use the following challenge set in our proof of linearity Π_{Lin} .

Challenge Set. Let κ be such that $\binom{d}{\kappa} \cdot 2^\kappa > 2^\lambda$ and define $\mathcal{C}_\kappa = \{c \in R_q \mid \|c\|_\infty = 1 \wedge \|c\|_1 = \kappa\}$ and $\bar{\mathcal{C}}_\kappa = \{c - c' \mid c, c' \in \mathcal{C}_\kappa \wedge c \neq c'\}$.

Proof of Linearity. In Step 2 of DDec , we prove well-formedness of the linear decryption share. The protocol Π_{Lin} produces a proof that a committed value v is a multiple of another committed value u with respect to a public scalar g . In our setting, we will prove $\llbracket E_{ij} - p^{-1} \text{ds}_{ij} \rrbracket = -p^{-1} c_i \llbracket \text{dk}_j \rrbracket$.

The exact relation for the proof system is:

$$\mathcal{R}_{\text{LIN}} := \left\{ (x, w) \left| \begin{array}{l} x := (\text{pk}_C, \text{com}_u, \text{com}_v, g) \wedge \\ w := (d_u = (u, \vec{r}_u, f_u), d_v = (v, \vec{r}_v, f_v)) : \\ u = g \cdot v \wedge \text{Open}(\text{pk}_C, \text{com}_u, d_u) \wedge \text{Open}(\text{pk}_C, \text{com}_v, d_v) \end{array} \right. \right\}.$$

The proof of linearity π_{LIN} is computed as follows [BDL⁺18]:

1. Sample vectors \vec{y}_u and \vec{y}_v of length k over R_q according to $D_{\sigma_{\text{LIN}}}$ and compute $\vec{w}_u = \vec{a}_1 \cdot \vec{y}_u$ and $\vec{w}_v = \vec{a}_1 \cdot \vec{y}_v$ and $t = g \cdot \vec{a}_2 \cdot \vec{y}_u - \vec{a}_2 \cdot \vec{y}_v$.
2. Hash $(\vec{w}_u, \vec{w}_v, t)$ to c in \mathcal{C}_κ , and compute $\vec{z}_u = \vec{y}_u + c \cdot \vec{r}_u$, $\vec{z}_v = \vec{y}_v + c \cdot \vec{r}_v$.
3. Rejection sample with respect to (\vec{y}_u, \vec{z}_u) , and (\vec{y}_v, \vec{z}_v) . If it outputs 1 then output $\pi_{\text{LIN}} = (c, \vec{z}_u, \vec{z}_v)$ and otherwise restart by sampling new (\vec{y}_u, \vec{y}_v) .

The verifier checks if $\|\vec{z}_u, \vec{z}_v\|_2 \leq 2\sigma_{\text{LIN}}\sqrt{k \cdot d}$ and if the hash of $(\vec{a}_1 \cdot \vec{z}_u - c \cdot c_{u,1}, \vec{a}_1 \cdot \vec{z}_v - c \cdot c_{v,1}, g \cdot \vec{a}_2 \cdot \vec{z}_u - \vec{a}_2 \cdot \vec{z}_v + c_{v,2} + g \cdot c_{u,2})$ equals c . It outputs 1 if all checks verify, and otherwise it outputs 0.

Using the improved rejection sampling techniques from [LNS21], we set $\sigma_{\text{LIN}} = B_{\text{Com}} \cdot \kappa\sqrt{d}$. The size of π_{LIN} is $2kd \log_2(4\sigma_{\text{LIN}})$ bits.

Amortized Proofs of Boundedness. In Step 3 of DDec, we use Π_{Bnd} to prove boundedness on noise drowning terms E_{ij} . The main idea is that we use a bit-decomposition E_i to produce a long vector with small entries, allowing the application Π_{Small} to prove an exact bound on E . We define Π_{Bnd} as an adapted version of the Π_{Small} protocol.

The previous work by Aranha et al. [ABGS23] used the amortised relaxed proofs by Baum et al. [BBC⁺18] to get smaller proof sizes at the cost of slightly increasing the overall parameters of the voting scheme because of the slack inherent in the proof system. In practice, this leads to a slightly larger modulus q but does not impact the ring dimension d . However, in our setting, we get better parameters in practice for the whole scheme when giving exact proofs of boundedness, even though the proofs themselves are larger. The precise relation of the proof system, with batch size ℓ' and secret vectors bounded in the ℓ_∞ norm by B_{Drown} , is:

$$\mathcal{R}_{\text{BND}} := \left\{ (x, w) \mid \begin{array}{l} x := (\text{pk}_C, \{\text{com}_i\}_{i \in [\ell']}) \wedge w := (\{d_i = (u_i, \vec{r}_i, f_i)\}_{i \in [\ell']}) : \\ \forall i \in [\ell'] : \|u_i\|_\infty \leq B_{\text{Drown}} \wedge \text{Open}(\text{pk}_C, \text{com}_i, d_i) \end{array} \right\}.$$

Since Π_{Small} is a proof system that scales with the number of possible values of the secret vectors, we use bit decomposition techniques to limit a blow-up in terms of running time, memory usage and proof size, to the cost of proving knowledge of longer secret vectors.

Any integer E between 0 and q can be represented in base b as $E = [b_0 \ b_1 \ \dots \ b_\zeta] \circ [1 \ b \ \dots \ b^\zeta]$ for unique coefficients b_i between 0 and $b - 1$ and $\zeta = \lceil \log_b q \rceil - 1$ where \circ is the dot product. This can be naturally extended to vectors, matrices and modules, particularly for our commitment matrix A . Since the commitment randomness is already short, we only need to decompose the last element in the secret vector, and we can do so in the following way (note that we abuse notation, where after $(*)$ the elements before $|$ are in R_q and the elements after are in \mathbb{Z}_q , but any element in R_q can be represented in \mathbb{Z}_q^d):

$$\begin{aligned} \mathbf{A}_{ij} \mathbf{s}_{ij} &= \begin{bmatrix} 1 & a_{1,1} & \mathbf{a}_{1,2} & | & 0 \\ 0 & 1 & \mathbf{a}_{2,2} & | & 1 \end{bmatrix} \begin{bmatrix} \vec{r}_{E_{ij}} \\ E_{ij} \end{bmatrix} \\ &\stackrel{(*)}{=} \begin{bmatrix} 1 & a_{1,1} & \mathbf{a}_{1,2} & | & 0 & \dots & 0 \\ 0 & 1 & \mathbf{a}_{2,2} & | & 1 & \dots & b^\zeta \end{bmatrix} \begin{bmatrix} \vec{r}_{E_{ij}} \\ E_{0ij} \\ \vdots \\ E_{\zeta ij} \end{bmatrix} = \bar{\mathbf{A}}_{ij} \bar{\mathbf{s}}_{ij}. \end{aligned}$$

Here, the ring element E_{ij} is decomposed, and elements $E_{0ij}, \dots, E_{\zeta ij}$ have integer values between 0 and $b - 1$. We note that these statements are equivalent and that the length of $\bar{\mathbf{A}}_{ij}$ is $d(k + \zeta + 1)$ over \mathbb{Z}_q instead of $d(k + 2)$.

Finally, we use the Π_{Small} protocol to prove ternary secret values as above but with a tweak: the public matrix input to the protocol is $\bar{\mathbf{A}}_{ij}$ instead of \mathbf{A}_{ij} , and we change the coefficient values that we are checking for in the proof. For the first $d \cdot k$ values we are checking for $(0, 1, -1)$ coefficients but for the next $d(\zeta + 1)$ values we are checking for $(0, 1, 2)$ coefficients instead (this is a minor tweak of line 3 in [ABGS23, Figure 5] that does not impact the performance of the protocol in any way, these values are initially arbitrary to the proof system). Since the other secret parts are ternary, we have that $\zeta = \lceil \log_3 B_{\text{Drown}} \rceil - 1$.

3.5 Security Analysis

We analyse the security of our verifiable voting scheme presented in Figs. 5 and 6. We examine the shuffle and distributed decryption protocols individually and prove their security. We emphasise that this is the standard approach in modern e-voting schemes based on mix-nets and is similar to the protocol used by Swiss Post for voting in Switzerland (they have a mix-net with key-switching so that there is only one decryption server needed in the end, while we only do mixing and then provide a distributed decryption protocol; these approaches are equivalent). Then, these standard primitives can be composed into a secure voting scheme as described in Chapter 14 in the book by Gjøsteen [Gjø22] and Appendix H in the full version of Aranha et al. [ABGS22]. Hence, we inherit their arguments for integrity, privacy and verifiability of the overall e-voting scheme. Our main contribution is more efficient building blocks that fit into this gold standard framework for e-voting, leading to a lattice-based e-voting scheme with improved performance over the state-of-the-art by Aranha et al. [ABGS23]. We *do not* propose a new model for e-voting.

Verifiable Mixing Protocol. We analyse the security of the verifiable shuffle protocol implicitly defined by the tuple of algorithms $\Pi_{\text{AMix}} := (\text{KeyGen}_A, \text{Cast}_A, \text{Mix}_A, \text{Dec}_{\text{NTRU}})$. We conduct this analysis against the set of standard definitions in Section A for verifiable mixing. Thus, our security results refer to ‘mixing’ properties, but we emphasise that this is another term for shuffling in this context. Note that security is analysed concerning a single shuffle server. In the context of our voting protocol, ballot privacy requires that at least one shuffle server is honest. We say that Π_{AMix} is *secure* if it satisfies the properties of mixing completeness, mixing soundness, and mixing simulatability.

Lemma 4 (Π_{AMix} Completeness). *If the protocols Π_{Small} and Π_{Shuf} are complete then Π_{AMix} always terminates. Moreover, if the input ciphertexts c_i have noise bounded by B_{Dec} , and the total noise added in Mix_A is bounded by B_{Mix} such that $(B_{\text{Mix}} + B_{\text{Dec}}) \leq \lfloor q/2 \rfloor$, then the output ciphertexts \hat{c}_i decrypt to the same messages as c_i .*

Proof. Since Π_{Small} and Π_{Shuf} are complete, the protocol will finish and the verifier will accept the output of the mix. Since $(B_{\text{Mix}} + B_{\text{Dec}}) \leq \lfloor q/2 \rfloor$, it follows that decryption is correct. Thus, Π_{AMix} is complete. \square

Lemma 5 (Π_{AMix} Soundness). *Suppose the input ciphertexts, output ciphertexts, and commitments in the shuffle are*

$$(c_1, \dots, c_\tau, \hat{c}_1, \dots, \hat{c}_\tau, \llbracket c'_1 \rrbracket, \dots, \llbracket c'_\tau \rrbracket),$$

respectively. Let Ext_1 be a knowledge extractor for the protocol Π_{Small} with success probability ϵ_1 and let Ext_2 be a knowledge extractor for the protocol Π_{Shuf} with success probability ϵ_2 . Then we can construct a knowledge extractor Ext_0 that succeeds with probability $\epsilon_0 \leq \epsilon_1 \cdot \epsilon_2$ in extracting (1) a permutation π , (2) encryption randomness s'_i, e'_i for $i \in [\tau]$, and (3) commitment randomness $\mathbf{r}_{c'_i}$ for $i \in [\tau]$ such that

- *The $\llbracket c'_i \rrbracket$ are commitments, using commitment randomness $\mathbf{r}_{c'_i}$, to ciphertexts of 0, say c'_i , with encryption randomness s'_i, e'_i .*
- $\|(s'_i, e'_i)\|_\infty \leq B_{\text{Mix}}$
- $\hat{c}_{\pi(i)} = c_i + c'_i$

Proof. First, note that the conditions above exactly constitute a correct mix so that if an adversary provides us

$$(c_1, \dots, c_\tau, \hat{c}_1, \dots, \hat{c}_\tau, \llbracket c'_1 \rrbracket, \dots, \llbracket c'_\tau \rrbracket),$$

and we are able to extract as in the statement of the Lemma, then the adversary must have performed a correct mix. The main observation is that both extractors Ext_1 and Ext_2 must succeed. Before verifying the shuffle proof π_{Shuf} , the verifier sets

$$\llbracket \hat{c}_i \rrbracket := \llbracket c_i \rrbracket_0 + \llbracket c'_i \rrbracket_{r_{c'_i}}$$

Here 0 and $r_{c'_i}$ are the commitment randomness values resp. The verifier then runs the π_{Shuf} verification checks against this set of commitments. This ensures that, if the proof verifies, the commitments com_i in the shuffle proof statement use exactly the same commitment randomness as used in the commitments to the c'_i .

If the extractor Ext_1 succeeds, we are able to extract τ randomness vectors $\mathbf{s}_{c'_i}$ bounded by B_{Small} , which gives us the randomness for both the commitments and ciphertexts used in the protocol. However, if the adversary is able to cheat in Π_{Shuf} then the output ciphertexts will be different to the ciphertexts we extract from Π_{Small} , and hence, we have not yet extracted all objects as in the Lemma.

If the extractor Ext_2 succeeds, we are able to extract both the permutation π and τ randomness vectors $\mathbf{r}_{c'_i}$ used in the commitments. However, if the adversary is able to cheat in Π_{Small} then the output ciphertexts might have more noise than $(B_{\text{Dec}} + B_{\text{Mix}})$ and lead to decryption failures, and hence, we have not yet extracted all objects in the Lemma.

We conclude that it is both necessary and sufficient that both extractors succeed at the same time to extract witnesses with respect to the same set of output ciphertexts and proofs to extract both the randomness used to encrypt, the randomness used to commit, and the permutation used to shuffle, and hence, to extract all objects in the Lemma. We thus conclude that $\epsilon_0 \leq \epsilon_1 \cdot \epsilon_2$. □

Lemma 6 (Π_{AMix} Simulatability). *Suppose the protocols Π_{Small} and Π_{Shuf} are honest-verifier zero-knowledge and that Com is hiding, Then there exists a simulator for Π_{AMix} such that for any distinguisher Adv_0 with advantage ϵ_0 , there exist adversaries $\text{Adv}_1, \text{Adv}_2$ against the HVZK of the Π_{Small} and Π_{Shuf} protocols, and an adversary Adv_3 against hiding of the commitment scheme with advantage ϵ_3 , with advantages ϵ_1, ϵ_2 , and ϵ_3 respectively such that $\epsilon_0 \leq \epsilon_1 + \epsilon_2 + \epsilon_3$. The runtime of $\text{Adv}_1, \text{Adv}_2, \text{Adv}_3$ are the same as of Adv_0 .*

Proof. The simulator is given a set of input ciphertexts and a set of output ciphertexts from an honest mix. The simulator simulates the zero-knowledge proofs Π_{Small} and Π_{Shuf} using the appropriate simulators and replaces the commitments to the ciphertexts with commitments to zero.

The claim about the simulator follows from a hybrid argument. In the following games, we denote by E_i the event that the adversary wins in Game i .

Game 0. This is the real-world protocol. We have that

$$\Pr[E_0] = \epsilon_0.$$

Game 1. Here, we replace the Π_{Shuf} arguments with simulated arguments. We have that, by the honest verifier zero-knowledge property of Π_{Shuf} ,

$$|\Pr[E_1] - \Pr[E_0]| \leq \epsilon_1.$$

Game 2. Here we replace the Π_{Small} by simulated arguments. We have that by the honest verifier zero-knowledge of Π_{Small}

$$|\Pr[E_2] - \Pr[E_1]| \leq \epsilon_2.$$

Game 3. Here we replace the commitments to ciphertexts with commitments to zero. We have that by the hiding property the commitment scheme that

$$|\Pr[E_3] - \Pr[E_2]| \leq \epsilon_3.$$

After the changes, we are left with a simulator for the actively secure protocol that outputs a proof that is independent of any secrets, and so the advantage ϵ_0 of the adversary Adv_0 is thus

$$\epsilon_0 \leq \epsilon_1 + \epsilon_2 + \epsilon_3,$$

as claimed. \square

Verifiable Distributed Decryption Protocol. We now analyse the security of the PKE with distributed decryption implicitly defined by the tuple of algorithms $\Pi_{\text{ADDec}} := (\text{KeyGen}_A, \text{Cast}_A, \text{Dec}_{\text{NTRU}}, \text{DDec}_A, \text{Comb}_A)$. We say that Π_{ADDec} is *secure* if it satisfies the properties of threshold correctness, threshold verifiability, and distributed decryption simulatability. For completeness, we give the full definitions of these notions in Section B. Since many of these properties rely on building blocks used in previous works, we provide a proof sketch here and refer the reader to [ABGS23] for the full arguments. We will however make parameter constraints explicit to aid in the performance analysis of Section 5.

Lemma 7 (Π_{ADDec} Threshold Correctness). *Suppose Π_{Lin} and Π_{Bnd} are complete, and the total noise in each ciphertext $(1 + 2^{\text{sec}}/p\xi_2)B_{\text{Dec}}$ is less than $\lfloor q/2 \rfloor$, then Π_{ADDec} is threshold correct.*

Proof. Examining the threshold correctness, define the predicate $P_{\text{sk}_A}(\cdot)$ so that $P_{\text{sk}_A}(c) = 1$ if and only if $\|\text{sk}_A \cdot c\|_\infty < B_{\text{Dec}}$. Then given a set of adversarially generated ciphertexts $\{c\}_{i \in [\tau]}$ satisfying $P_{\text{sk}_A}(c_i) = 1$ for all $i \in [\tau]$, we have that $\left\| \sum_{j \in [\xi_2]} \text{ds}_{ij} \right\|_\infty < q/2$ and so the Comb algorithm will return the correct decryption of c_i . The completeness of Π_{Lin} and Π_{Bnd} ensure that the arguments will be accepted, thus, Π_{ADDec} is threshold correct. \square

Lemma 8 (Π_{ADDec} Threshold Verifiability). *Let Adv_0 be an adversary against threshold verifiability of Π_{ADDec} with advantage ϵ_0 . Then there exists adversaries Adv_1 and Adv_2 against soundness for Π_{Lin} and Π_{Bnd} , respectively, with advantages ϵ_1 and ϵ_2 , such that $\epsilon_0 \leq \epsilon_1 + \epsilon_2$.*

Proof. As in the definition of threshold verifiability, we only consider ciphertexts such that $P_{\text{sk}_A}(c) = 1$. That is, ciphertexts whose noise is bounded by B_{Dec} . Note that if Comb accepts a ciphertext for which decryption is incorrect then, for some j , no relation $\text{ds}_{ij} = \text{dk}_j \cdot c_i + pE_{ij}$ holds for an E_{ij} of norm at most B_{Drown} .

This can happen in one of two ways. Either the proof of the linear relation $\text{ds}_{ij} = \text{dk}_j \cdot c_i + pE_{ij}$ is incorrect or the proof of the bound on E_{ij} is incorrect. In the first case one has an adversary Adv_1 against the soundness of Π_{Lin} and in the second case an adversary Adv_2 against the soundness of Π_{Bnd} . \square

Lemma 9 (Π_{ADDec} Simulatability). *Suppose NTRUencrypt cryptosystem is IND-CPA secure, Π_{Lin} and Π_{Bnd} are honest-verifier zero-knowledge, and Com is hiding. Then there exists a simulator such that for any distinguisher Adv_0 for this simulator with advantage ϵ_0 , there exists an adversary Adv_1 against the HVZK of Π_{Lin} , an adversary Adv_2 against the HVZK of Π_{Bnd} , and an adversary Adv_3 against the hiding of Com with advantages $\epsilon_1, \epsilon_2, \epsilon_3$ respectively such that $\epsilon_0 \leq \epsilon_1 + \epsilon_2 + \epsilon_3$. The runtime of $\text{Adv}_1, \text{Adv}_2$, and Adv_3 are the same as of Adv_0 .*

Proof. The claim about the simulator follows from a hybrid argument. In the following games, we denote by E_i the event that that the adversary wins in Game i .

Game 0. This is the real world protocol. We have that

$$\Pr[E_0] = \epsilon_0.$$

Game 1. Here, we replace the Π_{Lin} arguments with simulated arguments. We have that, by the honest verifier zero-knowledge property of Π_{Lin} ,

$$|\Pr[E_1] - \Pr[E_0]| \leq \epsilon_1.$$

Game 2. Here we replace the Π_{Bnd} by simulated arguments. We have that by the honest verifier zero-knowledge of Π_{Bnd} that

$$|\Pr[E_2] - \Pr[E_1]| \leq \epsilon_2.$$

Game 3. Here we replace the commitments noise E_{ij} with random commitments. We have that by the hiding property the commitment scheme that

$$|\Pr[E_3] - \Pr[E_2]| \leq \epsilon_3.$$

Game 4. Here we replace decryption shares ds_{ij} with random shares. By the choice of the statistical security parameter sec , which is chosen so that ds_{ij} is statistically indistinguishable from uniform, we have that

$$|\Pr[E_4] - \Pr[E_3]| = 0.$$

After the changes, we are left with a simulator for the actively secure protocol outputting a proof that is independent of any secrets, and so the advantage ϵ_0 of adversary Adv_0 is thus bounded by

$$\epsilon_0 \leq \epsilon_1 + \epsilon_2 + \epsilon_3.$$

□

4 NTRU Hardness

Before choosing concrete parameters for implementing our new voting scheme, it is clear that we needed to better understand the NTRU problem's hardness. This section contains that in-depth analysis which informs our parameter choices in Section 5.

Research on the security of the NTRU problem has revealed a significant improvement in the performance of lattice reduction algorithms when applied to NTRU lattices for so-called *overstretched* parameters. More precisely, analysis carried out over a series of works [ABD16, KF17, LW20] shows this weakening of NTRU occurs when the modulus q is very large compared to the ring dimension d and when secrets are small. Naturally, these works seek to determine the turning point at which q becomes large enough for such attacks to apply. We refer to this as the *fatigue point*.

4.1 Extending the NTRU Analysis

Until recently, only an asymptotic result was known about the position of the fatigue point, determined by Kirchner and Fouque as $q = d^{2.783+o(1)}$ [KF17].

Ducas-van Woerden Analysis. In their recent paper, Ducas and van Woerden [DvW21] improve on the asymptotic result of Kirchner and Fouque who showed in their analysis that the NTRU lattice, denoted $\Lambda^{h,q}$, contains an exceptionally dense sublattice, denoted $\Lambda^{g,f}$, of low volume which gives a constraint on the basis profile via Lemma 10 below⁶. Thus, the fatigue point for ternary secrets is narrowed down to $q = d^{2.484+o(1)}$. Building on this

⁶Ducas and van Woerden [DvW21] considers both NTRU (see Definition 1), and the module version of NTRU and use matrix notation in their paper. We continue to use polynomial ring notation to keep this consistent with our notation.

result, the authors perform an average-case analysis (rather than a worst-case bound) based on the volume of the relevant lattices and sublattices to arrive at a concrete prediction of the fatigue point. To facilitate their analysis, they identify two lattice reduction events that distinguish standard regimes from their overstretched counterparts.

Secret Key Recovery (SKR): The event in which a vector as short as the secret key is inserted into the lattice basis.

Dense Sublattice Discover (DSD): The event in which a vector of the dense sublattice is inserted into the basis.

A DSD event has been shown to shortly precede SKR by a cascading of further DSD events or enabling decryption of fresh ciphertexts. A further distinction between DSD events that are triggered at large positions $\kappa \leq 2d - \beta$ in the basis; DSD-LL (lucky lift), and those triggered at positions $\kappa = d + k - \beta$ for $0 < k \ll d$ and BKZ block size β , DSD-PT, in a run of progressive BKZ on NTRU lattices for fixed parameters $d = 127, \sigma^2 = 2/3$ and several moduli q . We redefine them below.

DSD-LL: For a few instances the DSD_κ event is triggered at large positions κ , up to $2d - \beta$. The inserted dense vector \mathbf{v} is again significantly longer than the secret key, but it has an unexpectedly short projection $\pi_\kappa(\mathbf{v})$ on the BKZ block $[\kappa : k\beta]$.

DSD-PT: The DSD_κ event is triggered at positions $\kappa = d + k - \beta$ for $0 < k \ll d$. The inserted dense vector \mathbf{v} is often significantly longer than the secret key, but still shorter than the q -vectors, and the projection of \mathbf{v} , $\pi_\kappa(\mathbf{v}) < \|\mathbf{b}_\kappa^*\|$.

They introduce the new DSD-PT event named after the Pataki and Tural in Lemma 10 and give an asymptotic analysis in Section 3 of their paper before giving an average-case analysis to construct a concrete estimator and compare this estimate with experiments in Section 4 and Section 5, respectively. The DSD-LL events are found to be a rare occurrence, and they speculate that the DSD-LL events are artefacts of using modest parameters in their experiments and are, for the most part, excluded from their analysis [DvW21].

Lemma 10 (Pataki and Tural [PT08]). *Let Λ be a d -dimensional lattice with basis $\mathbf{b}_0, \dots, \mathbf{b}_{d-1}$. For any k -dimensional sublattice $\Lambda' \subset \Lambda$ we have*

$$\text{vol}(\Lambda') \geq \min_J \prod_{j \in J} \|\mathbf{b}_j^*\|^2,$$

where J ranges over the k -size subsets of $\{0, \dots, d-1\}$ and \mathbf{b}_j^* is the Gram-Schmidt orthogonalization of \mathbf{b}_j .

Ducas and van Woerden give a concrete average-case estimate for the intersection of lattice volumes and the log-expectation of the volume of the dense NTRU sublattice $\Lambda^{g,f}$ in Eq. (2). They assume that all coefficients of the dense sublattice basis are sampled according to an independent continuous Gaussian distribution with standard deviation σ . They prove this estimate in [DvW21, Section 4.2] and verify these estimates experimentally with variance $\sigma^2 = 2/3$.

$$\mathbb{E}[\ln(\text{vol}(\Lambda^{g,f}))] = \frac{1}{2}d(\ln(2\sigma^2) + \psi(d)) + \sum_{i=0}^{d-1} \left[\psi\left(\frac{2d-i}{2}\right) - \psi(d) \right], \quad (2)$$

where ψ is the digamma function.

In [DvW21, Section 5] they compare their concrete predictions with experiments. They ran progressive BKZ 2.0 with 8 tours on NTRU with parameters $d = 127, \sigma^2 = 2/3$ for

several moduli q , accounting for both SKR and DSD-PT events. For BKZ block sizes $\beta > 30$ their experiments correspond nicely to their predictions, while for $\beta \leq 30$ their DSD-PT estimate is slightly pessimistic compared to the experiments. To further verify their fatigue point estimate, they did more experiments in larger, but still feasible dimensions. They found that the concrete estimator for the dense sublattice volume in Eq. (2) reasonably follows the observed experiments.

Through careful observation of the occurrence of these events, Ducas and van Woerden use their predictive model to determine the concrete fatigue point of NTRU with ternary secrets to be $q = 0.004 \cdot d^{2.484}$ for $d > 100$. One can use the scripts provided⁷ in their work to estimate the concrete hardness of NTRU. We also affirm their predictive model by running real experiments on low-dimensional instances to confirm this relation. It is worth noting that this is not the same NTRU estimator as the one provided in the lattice estimator⁸ maintained by Martin Albrecht which implements the 2016 estimate [ADPS16].

Beyond Ternary Secrets. The reader may have noticed that the discussion of the fatigue point, thus far, only focuses on the modulus and dimension of the ring. Recalling Definition 1 reminds us that f and g need not always be ternary. Indeed, many NTRU-based constructions use secrets with non-ternary coefficients. Let us consider f and g generated according to a Gaussian distribution D_σ of standard deviation σ . For convenience, the analysis of [DvW21] models the ternary secret case by sampling $f, g \leftarrow D_\sigma$ with $\sigma^2 = 2/3$. Varying σ can model any secret key size, and thus we will herein consider that f and g are always sampled according to some Gaussian D_σ . The natural question arises:

Does the choice of (secret size) σ influence the fatigue point, and if so, what is its impact?

To get some intuition on this, we recall the work of Steinfeld and Stehlé [SS11] in which the authors show how selecting σ sufficiently large gives rise to a public key $h = g/f$ that is statistically indistinguishable from uniform when f and g are sampled from D_σ . Moreover, they show that using such parameters allows one to remove the NTRU assumption from a proof of the NTRU cryptosystem. The σ needed for statistical security depends on the size of q and d . This suggests that fixing q and d and increasing σ makes the NTRU problem harder.

This observation goes some way to answering the first part of our question since it is clear that, for a sufficiently large σ , both SKR and DSD become ineffective.

Whilst using statistically uniform public keys provides peace of mind, this practice comes with significant efficiency losses. In addition to much larger key sizes, conditions for a cryptosystem’s correctness can become much more constraining. Note that for correct decryption of the `NTRUEncrypt` cryptosystem defined in Fig. 2, one needs the relation $\|p(gs + fe) + fm\|_\infty < q/2$ to hold. Clearly, using larger secrets f and g thus leads to less favourable parameters by pushing up the modulus q .

We, therefore, have a balancing act that needs to be performed when setting NTRU parameters; to keep parameters small whilst avoiding the attacks affecting overstretched regimes. Fortunately, the script provided in [DvW21] also allows for NTRU hardness estimations using any choice of σ though no analysis is performed in their work outside the ternary case. Nevertheless, their estimator provides a tool, much like the LWE estimator of Albrecht et al. [APS15], to analyse the concrete hardness of any given NTRU parameter set.

A More General Fatigue Relation. In our analysis, we are interested in answering the second part of the above question. In particular, we would like to know by *how much*

⁷See github.com/WvanWoerden/NTRUFatigue for their code.

⁸Lattice estimator: github.com/malb/lattice-estimator.

an increase in NTRU secret size affects the position of the fatigue point for a given ring dimension.

A simple calculation, following the analysis of [DvW21], Section 3.2, confirms that the asymptotic relation $q = d^{2.484+o(1)}$ holds regardless of the value of σ . This suggests that if the value of σ plays a role in the concrete, average case relation, it manifests in the leading constant. We can thus infer that, for some function ψ and constant c , the fatigue point is given by

$$q = c \cdot \psi(\sigma) \cdot d^{2.484}.$$

In order to determine the nature of ψ , we consider a range of $\sigma \in [2, 2^2, \dots, 2^{20}]$. For each σ , we perform a loglog-linear regression on the estimated fatigue points overall prime ring dimensions $199, \dots, 499$. This mimics the calculations of [DvW21] used in the ternary case. For a full explanation of why this is a sensible range to examine, we refer the reader to Section 5.3 of that work.

Next, we consider the predicted fatigue points as a geometric series. This reveals the predicted average-case fatigue point to be

$$q = 0.0058 \cdot \sigma^2 \cdot d^{2.484}. \quad (3)$$

The precision of this relation across all σ considered is very high. Whilst we could extend this part of our analysis to larger σ , it is highly unlikely that, for cryptographic applications, one would need to take σ higher than 2^{20} . We note also that, setting $\sigma^2 = 2/3$, we recover the fatigue point determined for the ternary case [DvW21].

This gives a definitive answer to our question about the impact of σ on the fatigue point. To give more gravity to this prediction, we also run a series of experiments for computable ring dimensions to validate this estimate. The results are displayed in Fig. 7.

We also give a second figure (Fig. 8) in which we plot q/σ^2 along the vertical axis. This reveals the accuracy of the preceding constant by showing how closely bunched the estimations and experiments are when normalised across varying σ . As Ducas and van Woerden observed in the ternary case, the estimator is slightly pessimistic, predicting a fatigue point roughly 15% lower than the one suggested by actual experiments. They give potential explanations for this discrepancy, pointing to the slope parameter used in the estimator, which is not well calibrated for such small block sizes. In practice, though, this small error only translates to a difference of 2-3 in the block size needed to run BKZ and thus hardly affects the predicted security. Importantly, our experiments show that this error remains constant even at larger moduli.

The Significance of σ^2 . Having determined the impact of σ on the concrete fatigue point for NTRU, we reflect on the structure of Eq. (3). As an illustrating example, let us return to the decryption correctness constraint for the NTRU cryptosystem. This can be written as $\sigma \cdot \mathcal{F}(p, d, \nu) < q$ for some function \mathcal{F} . Suppose for a given parameter set (d, q, σ) , this constraint is satisfied, but the corresponding NTRU instance does not provide adequate security. Let us then increase q by a constant factor δ , say. According to the constraint, this gives room for an increase in σ to $\delta \cdot \sigma$. Then Eq. (3) tells us that the new fatigue point for the set $(d, \delta \cdot q, \delta \cdot \sigma)$ increases by a factor of δ^2 . The important observation here is that, while increasing q in the first move might weaken the NTRU instance, the same increase permitted for σ actually gives rise to a *net increase in the hardness of the instance*. In summary, Eq. (3) tells us that it is possible to ‘win’ this cat-and-mouse game for NTRU that so often arises when setting lattice parameters.

We now consider how our analysis might be applied to existing works to refine parameter choices.




Figure 7: Experimental fatigue point values for a range of σ , calculated using BKZ with 8 tours on matrix NTRU instances. The straight-coloured lines show the estimated values using the (modified) estimator from [DvW21].




Figure 8: Experimental values for q/σ^2 illustrate that the fatigue point, when adjusted for σ , is modelled by $q/\sigma^2 = 0.0058 \cdot d^{2.484}$.

4.2 Implications for Existing Work

While the authors of [DvW21] note that parameters used in the NTRU-based NIST finalists are still secure to the degrees claimed, many works in the literature use different sets, some of which may fall foul of the dense sublattice attack, and thus, one needs to use the techniques described in the previous section to set parameters.

We examine an existing primitive in which the parameters fall short of providing claimed security levels. However, as suggested by Eq. (3), we can carefully re-select parameters so that a small increase in the secrets yields the security bump-up needed.

Blind Signatures [dPK22]. del Pino and Katsumata present a lattice-based (partially) blind signature using trapdoor sampling. In the (round optimal) construction given, a user passes the message to be signed in a *blind* way so that the signer does not learn the message they sign. This is done by committing to the message and then proving the well-formedness of this commitment. We will call this the *first flow*. The signer then creates an output passed back to the user (*second flow*). Finally, the user computes a signature for its original message using this response message from the signer.

In the first flow, Pino and Katsumata employ the NTRU-based linear homomorphic commitment scheme (LinHC) of [Kat21] to ensure the soundness and overall Quantum Random Oracle Model (QROM) security of the well-formedness proof. One must, therefore, choose parameters so that the relevant NTRU instance is hard. The choice of $d = 2048$, $q = 2^{66}$, and ternary NTRU secrets is informed by the constraint requiring straight-line extractability of the proof system. However, as we have observed, such large moduli run the risk of taking a parameter set into overstretched territory. Moreover, these values give rise to only 63 bits of security when run through the estimator of [DvW21] rather than 128.

To rectify this situation, there are two common strategies; either one can increase the ring dimension used throughout the scheme or use sufficiently large NTRU secrets that the corresponding public key is *statistically* indistinguishable from uniform.

Doubling the ring dimension in [dPK22] from 2048 to 4096 (to retain the implementation benefits of a power-of-two dimension) and computing the other parameters accordingly, 128 bits of security is reached at the cost of doubling the sign-request flow (69.2MB), doubling the returned ‘pre-signature’, and doubling the user’s final signature size to 200 KB.

The alternative method, using a statistically uniform public commitment key turns out to be impossible whilst satisfying all parameter constraints simultaneously.

We now exhibit the benefits of the relation Eq. (3), as revealed by our analysis, when applied to the problem of setting NTRU parameters with the same ring dimension $d = 2048$, we increase σ_{NTRU} (secret size). This has the effect of pushing up the modulus needed to facilitate the straight-line extraction condition. The reader might observe that increasing q reduces the hardness of the problem again. However, Eq. (3) reveals that it is possible to ‘win’ this cat-and-mouse game since the fatigue point increases *quadratically* with the size of the secrets. We thus propose the following parameters to ensure 128 bit security is reached:

$$q \approx 2^{74}, \quad p \approx 2^{41}, \quad \sigma_{\text{NTRU}} = 13,$$

where p is the prime used to commit to the witness in the LinHC protocol. Fortunately, this change only has a small effect on the total communication cost. In the first flow, the user signing query increases from 34 MB to 35.4 MB, and the sizes of the user’s pre-signature and final signature output are unchanged. This significantly improves the sizes that arise from changing the ring dimension and avoids doubling the final signature.

Summary. Simply increasing the size of the NTRU secrets may be all that is needed to ensure the correct security threshold is reached. In other settings, this also pushes up the

modulus over which a scheme is defined, as in the examples above. However, the scheme may also rely on other hardness assumptions, such as RLWE, as in our voting scheme, which is defined over the same ring. Now, the RLWE problem may no longer be hard for the adjusted parameters, and one may need to increase the ring *dimension* to find parameters for which both problems are hard. This can make what was an efficient scheme into one that cannot be deployed in practice.

Clearly, such balancing acts must be approached with a good understanding of the hardness of NTRU instances. We aim to further demonstrate the advantages of this approach when setting concrete parameters for our voting scheme in Section 5 where our fine-grained analysis allows us to dramatically bring down the overall communication cost.

5 Performance

We analyse the practical performance of our voting scheme. We begin by identifying all system parameters and any constraints that apply to them. These are displayed in Table 6. Next, we compute a sample set of parameters that satisfy the necessary constraints and give rise to a minimum of 128 bits of security. Table 2 displays these values. Finally, using these parameters, we compute the concrete communication cost of our voting system. The resulting sizes are compared to the previous work of [ABGS23], revealing a significant improvement in the state-of-the-art for cryptographic voting from quantum-safe assumptions. These results are displayed in Table 1.

5.1 Setting Parameters

We begin by collecting all parameters of the scheme and noting any constraints applying to them in Table 6.

Next, we closely examine the constraint needed for the correct (perfect) decryption of votes as performed by the **Comb** algorithm. This turns out to be the most influential constraint on the overall efficiency of the scheme. In particular, this constraint informs our choice of the global ring dimension d and modulus q , which most directly affect the communication sizes.

Decryption Correctness. After passing through the mix-net of ξ_1 shuffle servers, a ciphertext is of the form

$$c = p(h \sum_{k \in [\xi_1]} s_k + \sum_{k \in [\xi_1]} e_k) + m,$$

where the encryption randomness terms s_k and e_k are sampled from S_ν . Next, this ciphertext is passed to a decryption server, which computes a decryption share of the form $ds_j = f_j \cdot c + pE_j$. Then the **Comb** algorithm, on collecting $\{ds_j\}_{j \in [\xi_2]}$, outputs

$$v' = \left(\sum_{j \in [\xi_2]} ds_j \pmod{q} \right) \pmod{p}.$$

In order for the result of this process to yield the original ballot cast, we require the infinity norm of the sum here to be bounded by $\lfloor q/2 \rfloor$. It follows that a sufficient constraint for this correct decryption is:

$$p \cdot d \cdot t \cdot \sigma_{\text{NTRU}} \cdot (2\xi_1 \cdot \nu + 1/2)(1 + 2^{\text{sec}}) < \lfloor q/2 \rfloor, \quad (4)$$

where t is the rejection parameter in $\text{KeyGen}_{\text{NTRU}}$.

Table 2: Sample parameter set.

Parameter	Explanation	Value
λ	Computational security parameter	128
d	Ring dimension	2048
q	Ciphertext and commitment modulus	$\approx 2^{59}$
sec	Statistical security parameter	40
p	Plaintext modulus	2
t	KeyGen _{NTRU} rejection parameter	1.058
ν	Infinity norm of encryption randomness	1
ξ_1, ξ_2	Number of shuffle and decryption servers	4
B_{Com}	Infinity norm of commitment randomness	1
B_{Dec}	Noise in ciphertext	262144
B_{Drown}	Infinity norm of noise drowning term E_{ij}	$\approx 2^{55}$
σ_{NTRU}	Standard deviation for encryption secret key	7.12
η	Reed-Solomon encoding randomness length	325
ℓ_{Small}	Proof batch size in Π_{Small}	9830
ℓ_{Bnd}	Proof batch size in Π_{Bnd}	12288
μ_{Small}	Reed-Solomon message length in Π_{Small}	10565
μ_{Bnd}	Reed-Solomon message length in Π_{Bnd}	8517
μ'_{Small}	Reed-Solomon message dimension in Π_{Small}	23988
μ'_{Bnd}	Reed-Solomon message dimension in Π_{Bnd}	181550
γ_{Small}	Reed-Solomon code length in Π_{Small}	26616
γ_{Bnd}	Reed-Solomon code length in Π_{Bnd}	198668

Computational Security. Having chosen parameters satisfying the constraints of Table 6, we must ensure that the underpinning lattice problems are sufficiently hard for these parameters.

For RLWE we follow standard convention by using the estimator [APS15]. This estimates the cost of BKZ conservatively by focusing only on the cost of a single uSVP oracle call, a core operation in BKZ. The number of such calls required has been estimated to be $8d$ for a lattice dimension d , and we follow this estimate.

To determine the NTRU problem’s hardness, we use the analysis of Section 4. Having settled on a ring dimension d and modulus q giving sufficient hardness of the RLWE problem, we use (4) to determine the maximum standard deviation permissible for generating the NTRU secrets (f, g) . Finally, following the procedure described in Section 4, we calculate the estimated hardness of NTRU. Again, we employ the conservative formula $0.292\beta + 16.4 + \log_2(8d)$ used in [DTGW17, SPL⁺17, BIP⁺22] to compute bit-security from blocksize β .

In order to ensure the binding property of the BDLOP commitment schemes we use,

Table 3: Ciphertext, commitment, and proof sizes per voter. Note that the two sizes in [ABGS23] reflect commitments to noise-drowning terms and ciphertexts, respectively.

Scheme	c_i	$\llbracket R_q \rrbracket$	π_{Shuf}	π_{Lin}	π_{Small}	π_{Bnd}
[ABGS23] [KB]	80	80/120	150	35	20	2
Our [KB]	15	30	63	18	22	22

Table 4: Ciphertext and commitment timings. Numbers were obtained averaging over 10^4 executions measured using the cycle counter available on the platform.

Scheme	Com	Open	Enc	Dec	DDec
[ABGS23] [ms]	0.45	2.76	0.74	0.64	1.56
Our [ms]	0.17	0.80	0.20	0.21	0.45

the RSIS problem must be hard. We use the relation due to Micciancio and Regev [MR09], which states that LLL will recover a short vector a vector of 2-norm $2^{(2\sqrt{d}\log_2 q \log_2 \delta)}$. δ is the root Hermite factor and $\delta < 1.0045$ gives rise to at least 128 bits of security. Owing to the horizontally long nature of the commitment matrix used, the hardness of the corresponding RSIS instance easily meets this threshold.

5.2 Sample Parameters and Total Size

Table 2 gives a sample set of parameters generated by following the process described in the previous section. In Table 3, we present the total sizes of objects in our voting scheme and compare them with those of [ABGS23]. We denote the output of each shuffle node by $\pi_{\mathcal{S}_i}$, including ciphertexts, commitments, proofs of shortness, and shuffle proofs. Similarly, we denote the total output of each decryption node as $\pi_{\mathcal{D}_j}$, consisting of decryption shares, commitments, proofs of linearity and boundedness.

Our scheme reduces ciphertext size by over a factor of five. Moreover, the reduction in commitment sizes and constituent proofs leads to shuffle server outputs of 130 KB per vote, which are three times smaller, and decryption server outputs of 85 KB per vote, which are half of those in [ABGS23]. This represents a $2.5\times$ overall efficiency gain over [ABGS23] as summarised in Table 1.

5.3 Benchmarks

We adapt the proof-of-concept implementation by [ABGS23] to fit our scheme since the framework is the same. Our benchmarks were collected on an Intel Kaby Lake Core i7-7700 CPU machine with 64 GB of RAM running single-threaded at 3.6 GHz, with Turbo Boost disabled to reduce measurement variability. This is a similar machine as in [ABGS23]. Our code is available at https://github.com/carrosa/ntru_voting_impl.

We compare the timings in Table 4 and Table 5. Analysing our experiments, each shuffle server takes $(0.20 + 0.17 + 17.5 + 44.2) = 62$ ms and each decryption server takes $(0.17 + 0.45 + 16.9 + 310.5) = 328$ ms. Given four servers, where shuffles are performed sequentially and decryption is performed in parallel, the total time is 576 ms, making our scheme twice as fast as [ABGS23] as summarized in Table 1.

We finally note that the proofs of linearity in the shuffle and the batched proofs of shortness and boundedness during shuffles and decryption can be computed in parallel,

Table 5: Proving and verification times, obtained by computing the average of 100 executions with $\tau = 1000$.

Scheme	π_{Lin}	π_{Shuf}	π_{Small}	π_{Bnd}
[ABGS23] [ms]	(43.4 + 6.4)	(44.9 + 7.9)	(214.4 + 10.0)	(92.7 + 23.9)
Our [ms]	(16.9 + 2.0)	(17.5 + 2.1)	(44.2 + 4.0)	(310.5 + 4.3)

and that powerful servers dedicated to an election with Turbo Boost enabled would most likely outperform our numbers by at least an order of magnitude.

5.4 Future Improvements

We provide some directions for interesting future work:

1. *Return codes.* To extend our scheme and ensure voter verifiability, we need to add return codes to our scheme. This can be done by extending the work of [HS22] from BGV to NTRU. This also includes verifiable encryption [LNP22].
2. *Improved noise analysis.* Our results can possibly be improved using techniques in [AKSY22, BS23b, CSS+22, KLSS23]. We use 40 bits of statistical noise drowning to protect the secret key in the distributed decryption protocol. This can possibly be improved if we choose parameters based on how many ciphertexts we will decrypt or change noise drowning techniques to be Gaussian distributed, compute the Rényi divergence, or analyse hints to estimate the leakage.
3. *Module assumptions.* The new NIST post-quantum key-encapsulation mechanism ML-KEM [SAB+22] and digital signature ML-DSA [LDK+22] are based on the module LWE and SIS problems [LS15]. Our scheme could potentially be improved by instantiating our framework based on these assumptions combined with the more recent module NTRU assumption [CKKS19, CPS+20].
4. *Succinct lattice ZKPs.* Recent development of succinct proof systems from lattice assumptions such as LaBRADOR [BS23a] (which can be made zero-knowledge) could be applied to the shuffle and decryption processes in our scheme to produce sub-linear proof sizes (the overall communication would stay linear since we need to send commitments, ciphertexts and decryption shares for each vote).
5. *Improved parameters in other schemes.* Our extended NTRU analysis might lead to more efficient FHE parameters in [BIP+22] and [Klu22] using the same methodology that led to a more efficient instantiation of `NTRUEncrypt`.

References

- [ABD16] Martin R. Albrecht, Shi Bai, and Léo Ducas. A subfield lattice attack on overstretched NTRU assumptions - cryptanalysis of some FHE and graded encoding schemes. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016, Part I*, volume 9814 of *Lecture Notes in Computer Science*, pages 153–178. Springer, Berlin, Heidelberg, August 2016. doi:10.1007/978-3-662-53018-4_6.
- [ABG⁺21] Diego F. Aranha, Carsten Baum, Kristian Gjøsteen, Tjerand Silde, and Thor Tunge. Lattice-based proof of shuffle and applications to electronic voting. In Kenneth G. Paterson, editor, *Topics in Cryptology – CT-RSA 2021*, volume 12704 of *Lecture Notes in Computer Science*, pages 227–251. Springer, Cham, May 2021. doi:10.1007/978-3-030-75539-3_10.
- [ABGS22] Diego F. Aranha, Carsten Baum, Kristian Gjøsteen, and Tjerand Silde. Verifiable mix-nets and distributed decryption for voting from lattice-based assumptions. Cryptology ePrint Archive, Report 2022/422, 2022. URL: <https://eprint.iacr.org/2022/422>.
- [ABGS23] Diego F. Aranha, Carsten Baum, Kristian Gjøsteen, and Tjerand Silde. Verifiable mix-nets and distributed decryption for voting from lattice-based assumptions. In Weizhi Meng, Christian Damsgaard Jensen, Cas Cremers, and Engin Kirda, editors, *ACM CCS 2023: 30th Conference on Computer and Communications Security*, pages 1467–1481. ACM Press, November 2023. doi:10.1145/3576915.3616683.
- [Adi08] Ben Adida. Helios: Web-based open-audit voting. In Paul C. van Oorschot, editor, *USENIX Security 2008: 17th USENIX Security Symposium*, pages 335–348. USENIX Association, July / August 2008.
- [ADPS16] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - A new hope. In Thorsten Holz and Stefan Savage, editors, *USENIX Security 2016: 25th USENIX Security Symposium*, pages 327–343. USENIX Association, August 2016.
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th Annual ACM Symposium on Theory of Computing*, pages 99–108. ACM Press, May 1996. doi:10.1145/237814.237838.
- [AKSY22] Shweta Agrawal, Elena Kirshanova, Damien Stehlé, and Anshu Yadav. Practical, round-optimal lattice-based blind signatures. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022: 29th Conference on Computer and Communications Security*, pages 39–53. ACM Press, November 2022. doi:10.1145/3548606.3560650.
- [APS15] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015. doi:10.1515/jmc-2015-0016.
- [BBC⁺18] Carsten Baum, Jonathan Bootle, Andrea Cerulli, Rafaël del Pino, Jens Groth, and Vadim Lyubashevsky. Sub-linear lattice-based zero-knowledge arguments for arithmetic circuits. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part II*, volume 10992 of *Lecture Notes in Computer Science*, pages 669–699. Springer, Cham, August 2018. doi:10.1007/978-3-319-96881-0_23.

- [BD10] Rikke Bendlin and Ivan Damgård. Threshold decryption and zero-knowledge proofs for lattice-based cryptosystems. In Daniele Micciancio, editor, *TCC 2010: 7th Theory of Cryptography Conference*, volume 5978 of *Lecture Notes in Computer Science*, pages 201–218. Springer, Berlin, Heidelberg, February 2010. doi:10.1007/978-3-642-11799-2_13.
- [BDL⁺18] Carsten Baum, Ivan Damgård, Vadim Lyubashevsky, Sabine Oechsner, and Chris Peikert. More efficient commitments from structured lattice assumptions. In Dario Catalano and Roberto De Prisco, editors, *SCN 18: 11th International Conference on Security in Communication Networks*, volume 11035 of *Lecture Notes in Computer Science*, pages 368–385. Springer, Cham, September 2018. doi:10.1007/978-3-319-98113-0_20.
- [BG12] Stephanie Bayer and Jens Groth. Efficient zero-knowledge argument for correctness of a shuffle. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 263–280. Springer, Berlin, Heidelberg, April 2012. doi:10.1007/978-3-642-29011-4_17.
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In Shafi Goldwasser, editor, *ITCS 2012: 3rd Innovations in Theoretical Computer Science*, pages 309–325. Association for Computing Machinery, January 2012. doi:10.1145/2090236.2090262.
- [BHM20] Xavier Boyen, Thomas Haines, and Johannes Müller. A verifiable and practical lattice-based decryption mix net with external auditing. In Liqun Chen, Ninghui Li, Kaitai Liang, and Steve A. Schneider, editors, *ESORICS 2020: 25th European Symposium on Research in Computer Security, Part II*, volume 12309 of *Lecture Notes in Computer Science*, pages 336–356. Springer, Cham, September 2020. doi:10.1007/978-3-030-59013-0_17.
- [BIP⁺22] Charlotte Bonte, Iliia Iliashenko, Jeongeun Park, Hilder V. L. Pereira, and Nigel P. Smart. FINAL: Faster FHE instantiated with NTRU and LWE. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part II*, volume 13792 of *Lecture Notes in Computer Science*, pages 188–215. Springer, Cham, December 2022. doi:10.1007/978-3-031-22966-4_7.
- [BLNS21] Jonathan Bootle, Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. More efficient amortization of exact zero-knowledge proofs for LWE. In Elisa Bertino, Haya Shulman, and Michael Waidner, editors, *ESORICS 2021: 26th European Symposium on Research in Computer Security, Part II*, volume 12973 of *Lecture Notes in Computer Science*, pages 608–627. Springer, Cham, October 2021. doi:10.1007/978-3-030-88428-4_30.
- [BLS19] Jonathan Bootle, Vadim Lyubashevsky, and Gregor Seiler. Algebraic techniques for short(er) exact lattice-based zero-knowledge proofs. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part I*, volume 11692 of *Lecture Notes in Computer Science*, pages 176–202. Springer, Cham, August 2019. doi:10.1007/978-3-030-26948-7_7.
- [BS23a] Ward Beullens and Gregor Seiler. LaBRADOR: Compact proofs for R1CS from module-SIS. In Helena Handschuh and Anna Lysyanskaya, editors,

- Advances in Cryptology – CRYPTO 2023, Part V*, volume 14085 of *Lecture Notes in Computer Science*, pages 518–548. Springer, Cham, August 2023. doi:10.1007/978-3-031-38554-4_17.
- [BS23b] Katharina Boudgoust and Peter Scholl. Simple threshold (fully homomorphic) encryption from LWE with polynomial modulus. In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology – ASIACRYPT 2023, Part I*, volume 14438 of *Lecture Notes in Computer Science*, pages 371–404. Springer, Singapore, December 2023. doi:10.1007/978-981-99-8721-4_12.
- [CAE19] Anthony Cardillo, Nicholas Akinyokun, and Aleksander Essex. Online voting in ontario municipal elections: A conflict of legal principles and technology? https://link.springer.com/chapter/10.1007/978-3-030-30625-0_5, 2019. Accessed: 2024-02-27.
- [CBS00] CBS News. Online first in arizona. <https://www.cbsnews.com/news/online-first-in-arizona/>, 2000. Accessed: 27-02-2024.
- [CGGI16] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. A homomorphic LWE based E-voting scheme. In Tsuyoshi Takagi, editor, *Post-Quantum Cryptography - 7th International Workshop, PQCrypto 2016*, pages 245–265. Springer, Cham, 2016. doi:10.1007/978-3-319-29360-8_16.
- [Cha81] David L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2):84–90, February 1981. doi:10.1145/358549.358563.
- [CJL16] Jung Hee Cheon, Jinhyuck Jeong, and Changmin Lee. An algorithm for ntru problems and cryptanalysis of the ggh multilinear map without a low-level encoding of zero. *LMS Journal of Computation and Mathematics*, 19(A):255–266, 2016. doi:10.1112/S1461157016000371.
- [CKKS19] Jung Hee Cheon, Duhyeong Kim, Taechan Kim, and Yongha Son. A new trapdoor over module-NTRU lattice and its application to ID-based encryption. Cryptology ePrint Archive, Report 2019/1468, 2019. URL: <https://eprint.iacr.org/2019/1468>.
- [CMM19] Núria Costa, Ramiro Martínez, and Paz Morillo. Lattice-based proof of a shuffle. In Andrea Bracciali, Jeremy Clark, Federico Pintore, Peter B. Rønne, and Massimiliano Sala, editors, *FC 2019 Workshops*, volume 11599 of *Lecture Notes in Computer Science*, pages 330–346. Springer, Cham, February 2019. doi:10.1007/978-3-030-43725-1_23.
- [CPS⁺20] Chitchanok Chuengsatiansup, Thomas Prest, Damien Stehlé, Alexandre Wallet, and Keita Xagawa. ModFalcon: Compact signatures based on module-NTRU lattices. In Hung-Min Sun, Shiuh-Pyng Shieh, Guofei Gu, and Giuseppe Ateniese, editors, *ASIACCS 20: 15th ACM Symposium on Information, Computer and Communications Security*, pages 853–866. ACM Press, October 2020. doi:10.1145/3320269.3384758.
- [CSS⁺22] Siddhartha Chowdhury, Sayani Sinha, Animesh Singh, Shubham Mishra, Chandan Chaudhary, Sikhar Patranabis, Pratyay Mukherjee, Ayantika Chatterjee, and Debdeep Mukhopadhyay. Efficient threshold FHE with application to real-time systems. Cryptology ePrint Archive, Report 2022/1625, 2022. URL: <https://eprint.iacr.org/2022/1625>.

- [dPK22] Rafaël del Pino and Shuichi Katsumata. A new framework for more efficient round-optimal lattice-based (partially) blind signature via trapdoor sampling. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part II*, volume 13508 of *Lecture Notes in Computer Science*, pages 306–336. Springer, Cham, August 2022. doi:10.1007/978-3-031-15979-4_11.
- [dPLNS17] Rafaël del Pino, Vadim Lyubashevsky, Gregory Neven, and Gregor Seiler. Practical quantum-safe voting from lattices. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017: 24th Conference on Computer and Communications Security*, pages 1565–1581. ACM Press, October / November 2017. doi:10.1145/3133956.3134101.
- [DTGW17] Jintai Ding, Tsuyoshi Takagi, Xinwei Gao, and Yuntao Wang. Ding Key Exchange. Technical report, National Institute of Standards and Technology, 2017. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-1-submissions>.
- [DvW21] Léo Ducas and Wessel P. J. van Woerden. NTRU fatigue: How stretched is overstretched? In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021, Part IV*, volume 13093 of *Lecture Notes in Computer Science*, pages 3–32. Springer, Cham, December 2021. doi:10.1007/978-3-030-92068-5_1.
- [FWK21] Valeh Farzaliyev, Jan Willemson, and Jaan Kristjan Kaasik. Improved lattice-based mix-nets for electronic voting. In Jong Hwan Park and Seung-Hyun Seo, editors, *ICISC 21: 24th International Conference on Information Security and Cryptology*, volume 13218 of *Lecture Notes in Computer Science*, pages 119–136. Springer, Cham, December 2021. doi:10.1007/978-3-031-08896-4_6.
- [Gjo22] Kristian Gjosteen. *Practical Mathematical Cryptography*. CRC Press, 2022.
- [HMMP23] Thomas Haines, Rafieh Mosaheb, Johannes Müller, and Ivan Pryvalov. SoK: Secure E-voting with everlasting privacy. *Proceedings on Privacy Enhancing Technologies*, 2023(1):279–293, January 2023. doi:10.56553/popets-2023-0017.
- [HMS21] Javier Herranz, Ramiro Martínez, and Manuel Sánchez. Shorter lattice-based zero-knowledge proofs for the correctness of a shuffle. In Matthew Bernhard, Andrea Bracciali, Lewis Gudgeon, Thomas Haines, Arian Klages-Mundt, Shin’ichiro Matsuo, Daniel Perez, Massimiliano Sala, and Sam Werner, editors, *FC 2021 Workshops*, volume 12676 of *Lecture Notes in Computer Science*, pages 315–329. Springer, Berlin, Heidelberg, March 2021. doi:10.1007/978-3-662-63958-0_27.
- [HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *Third Algorithmic Number Theory Symposium (ANTS)*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer, June 1998.
- [HS22] Audhild Høgåsen and Tjerand Silde. Return codes from lattice assumptions. *E-VOTE-ID*, 2022. doi:10.15157/diss/025.
- [Kat21] Shuichi Katsumata. A new simple technique to bootstrap various lattice zero-knowledge proofs to QROM secure NIZKs. In Tal Malkin and Chris

- Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part II*, volume 12826 of *Lecture Notes in Computer Science*, pages 580–610, Virtual Event, August 2021. Springer, Cham. doi:10.1007/978-3-030-84245-1_20.
- [KCK⁺18] Robert Krimmer, David Duenas Cid, Iuliia Krivonosova, Priit Vinkel, and Arne Koitmaa. How much does an e-vote cost? cost comparison per vote in multichannel elections in estonia. https://link.springer.com/chapter/10.1007/978-3-030-00419-4_8, 2018. Accessed: 2024-02-27.
- [KF17] Paul Kirchner and Pierre-Alain Fouque. Revisiting lattice attacks on overstretched NTRU parameters. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017, Part I*, volume 10210 of *Lecture Notes in Computer Science*, pages 3–26. Springer, Cham, April / May 2017. doi:10.1007/978-3-319-56620-7_1.
- [KLSS23] Duhyeong Kim, Dongwon Lee, Jinyeong Seo, and Yongsoo Song. Toward practical lattice-based proof of knowledge from hint-MLWE. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part V*, volume 14085 of *Lecture Notes in Computer Science*, pages 549–580. Springer, Cham, August 2023. doi:10.1007/978-3-031-38554-4_18.
- [Klu22] Kamil Kluczniak. NTRU-v-um: Secure fully homomorphic encryption from NTRU with small modulus. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022: 29th Conference on Computer and Communications Security*, pages 1783–1797. ACM Press, November 2022. doi:10.1145/3548606.3560700.
- [LDK⁺22] Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact Knapsacks are collision resistant. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *ICALP 2006: 33rd International Colloquium on Automata, Languages and Programming, Part II*, volume 4052 of *Lecture Notes in Computer Science*, pages 144–155. Springer, Berlin, Heidelberg, July 2006. doi:10.1007/11787006_13.
- [LNP22] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part II*, volume 13508 of *Lecture Notes in Computer Science*, pages 71–101. Springer, Cham, August 2022. doi:10.1007/978-3-031-15979-4_3.
- [LNS21] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. Shorter lattice-based zero-knowledge proofs via one-time commitments. In Juan Garay, editor, *PKC 2021: 24th International Conference on Theory and Practice of Public Key Cryptography, Part I*, volume 12710 of *Lecture Notes in Computer Science*, pages 215–241. Springer, Cham, May 2021. doi:10.1007/978-3-030-75245-3_9.
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in*

- Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23. Springer, Berlin, Heidelberg, May / June 2010. doi:[10.1007/978-3-642-13190-5_1](https://doi.org/10.1007/978-3-642-13190-5_1).
- [LS15] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography*, 75(3):565–599, 2015. doi:[10.1007/s10623-014-9938-4](https://doi.org/10.1007/s10623-014-9938-4).
- [LTV12] Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In Howard J. Karloff and Toniann Pitassi, editors, *44th Annual ACM Symposium on Theory of Computing*, pages 1219–1234. ACM Press, May 2012. doi:[10.1145/2213977.2214086](https://doi.org/10.1145/2213977.2214086).
- [LW20] Changmin Lee and Alexandre Wallet. Lattice analysis on MinTRU problem. Cryptology ePrint Archive, Report 2020/230, 2020. URL: <https://eprint.iacr.org/2020/230>.
- [Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 738–755. Springer, Berlin, Heidelberg, April 2012. doi:[10.1007/978-3-642-29011-4_43](https://doi.org/10.1007/978-3-642-29011-4_43).
- [MR04] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. In *45th Annual Symposium on Foundations of Computer Science*, pages 372–381. IEEE Computer Society Press, October 2004. doi:[10.1109/FOCS.2004.72](https://doi.org/10.1109/FOCS.2004.72).
- [MR09] Daniele Micciancio and Oded Regev. *Lattice-based Cryptography*, pages 147–191. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. URL: https://doi.org/10.1007/978-3-540-88702-7_5.
- [New21] New South Wales Electoral Commission. ivote and 2021 nsw local government elections. <https://elections.nsw.gov.au/about-us/media-centre/news-and-media-releases/ivote-and-2021-nsw-local-government-elections>, 2021. Accessed: 2024-02-27.
- [PFH⁺20] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2020. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions>.
- [PFH⁺22] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In Shai Halevi and Tal Rabin, editors, *TCC 2006: 3rd Theory of Cryptography Conference*, volume 3876 of *Lecture Notes in Computer Science*, pages 145–166. Springer, Berlin, Heidelberg, March 2006. doi:[10.1007/11681878_8](https://doi.org/10.1007/11681878_8).

- [PT08] Gabor Pataki and Mustafa Tural. On sublattice determinants in reduced bases, 2008. URL: <https://arxiv.org/abs/0804.4014>, [arXiv:0804.4014](https://arxiv.org/abs/0804.4014).
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 84–93. ACM Press, May 2005. doi:10.1145/1060590.1060603.
- [RST⁺22] Dragos Rotaru, Nigel P. Smart, Titouan Tanguy, Frederik Vercauteren, and Tim Wood. Actively secure setup for SPDZ. *Journal of Cryptology*, 35(1):5, January 2022. doi:10.1007/s00145-021-09416-w.
- [SAB⁺22] Peter Schwabe, Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Gregor Seiler, Damien Stehlé, and Jintai Ding. CRYSTALS-KYBER. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [SMPS16] Paolo Spada, Jonathan Mellon, Tiago Peixoto, and Fredrik M. Sjöberg. Effects of the internet on participation: Study of a public policy referendum in brazil. *Journal of Information Technology & Politics*, 13(3):187–207, 2016. doi:10.1080/19331681.2016.1162250.
- [Sol01] Frederic I. Solop. Digital democracy comes of age: Internet voting and the 2000 arizona democratic primary election. *PS: Political Science & Politics*, 34(2):289–293, 2001. doi:10.1057/9780230523531_14.
- [SPL⁺17] Minhye Seo, Jong Hwan Park, Dong Hoon Lee, Suhri Kim, and Seung-Joon Lee. EMBLEM and R.EMBLEM. Technical report, National Institute of Standards and Technology, 2017. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-1-submissions>.
- [SS11] Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In Kenneth G. Paterson, editor, *Advances in Cryptology – EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 27–47. Springer, Berlin, Heidelberg, May 2011. doi:10.1007/978-3-642-20465-4_4.
- [Str19] Martin Strand. A verifiable shuffle for the GSW cryptosystem. In Aviv Zohar, Ittay Eyal, Vanessa Teague, Jeremy Clark, Andrea Bracciali, Federico Pintore, and Massimiliano Sala, editors, *FC 2018 Workshops*, volume 10958 of *Lecture Notes in Computer Science*, pages 165–180. Springer, Berlin, Heidelberg, March 2019. doi:10.1007/978-3-662-58820-8_12.
- [Swi23] Swiss Post. Swiss post article on e-voting introduction. <https://post-medien.ch/en/swiss-posts-e-voting-system-to-be-used-for-the-first-time-in-elections-this-autumn-following-further-development-and-successful-hacker-test/>, 2023. Accessed: 27-02-2024.
- [SXY18] Tsunekazu Saito, Keita Xagawa, and Takashi Yamakawa. Tightly-secure key-encapsulation mechanism in the quantum random oracle model. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part III*, volume 10822 of *Lecture Notes in Computer Science*, pages 520–551. Springer, Cham, April / May 2018. doi:10.1007/978-3-319-78372-7_17.

- [Vin15] Priit Vinkel. Remote electronic voting in estonia: Legality, impact, and confidence. *ResearchGate*, 2015. URL: https://www.researchgate.net/profile/Priit-Vinkel/publication/281319274_Remote_Electronic_Voting_in_Estonia_Legality_Impact_and_Confidence/links/55e194ef08aecb1a7cc68462/Remote-Electronic-Voting-in-Estonia-Legality-Impact-and-Confidence.pdf.

A Security of Verifiable Mixing

We define *completeness*, *soundness* and *simulatability* for a mixing protocol Π_{Mix} executed by a prover **Prover**, with respect to a generic encryption scheme $\mathcal{E} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ [ABGS23].

Definition 5 (Mixing Completeness). We say that the mixing protocol Π_{Mix} is *complete* if for honest PPT parties **Prover** and **Verifier** that follows the protocol then **Prover** on input a set of honestly generated ciphertexts will output a new set of ciphertexts together with a proof such that **Verifier** accepts the proof and the output ciphertexts decrypt to the same set of messages as the input ciphertexts. Hence, we want that for all $(\text{pp}, \text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\kappa)$, $\{c_i\}_{i \in [\tau]} \leftarrow \text{Enc}(\text{pk}, \{m_i\}_{i \in [\tau]})$, and $(\{\hat{c}_i\}_{i \in [\tau]}, \pi) \leftarrow \text{Prover}(\text{pp}, \text{pk}, \{c_i\}_{i \in [\tau]})$, we have

$$\Pr \left[\begin{array}{l} \{m_i\}_{i \in [\tau]} = \text{Dec}(\text{sk}, \{\hat{c}_i\}_{i \in [\tau]}) \\ 1 \leftarrow \text{Verifier}(\text{pp}, \text{pk}, \{c_i\}_{i \in [\tau]}, \{\hat{c}_i\}_{i \in [\tau]}, \pi) \end{array} \right] \leq 1 - \epsilon(\lambda),$$

where the probability is taken over **KeyGen**, **Enc** and **Prover**.

Definition 6 (Mixing Soundness). We say that the mixing protocol Π_{Mix} is *sound* if a dishonest PPT adversary **Adv** that can behave arbitrarily on input a set of honestly generated ciphertexts will not be able to output a new set of ciphertexts together with a proof such that an honest **Verifier** accepts the proof but the output ciphertexts decrypt to a different set of messages than the input ciphertexts. Hence, we want that for all $(\text{pp}, \text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\kappa)$, $\{c_i\}_{i \in [\tau]} \leftarrow \text{Enc}(\text{pk}, \{m_i\}_{i \in [\tau]})$, and $(\{\hat{c}_i\}_{i \in [\tau]}, \pi) \leftarrow \text{Adv}(\text{pp}, \text{pk}, \{c_i\}_{i \in [\tau]})$, we have

$$\Pr \left[\begin{array}{l} \{m_i\}_{i \in [\tau]} \neq \text{Dec}(\text{sk}, \{\hat{c}_i\}_{i \in [\tau]}) \\ 1 \leftarrow \text{Verifier}(\text{pp}, \text{pk}, \{c_i\}_{i \in [\tau]}, \{\hat{c}_i\}_{i \in [\tau]}, \pi) \end{array} \right] \leq \epsilon(\lambda),$$

where the probability is taken over **KeyGen**, **Enc** and **Adv**.

Definition 7 (Mixing Simulatability). We say that the mixing protocol Π_{Mix} is *simulatable* if a PPT adversary \mathcal{A} that on input a set of honestly generated ciphertexts can not distinguish between a real execution of the mixing protocol with accepting output and a protocol execution from a PPT simulator \mathcal{S} (given a set honestly mixed output ciphertexts) producing a simulated mixing proof. Hence, we want that

$$\left| \Pr \left[\begin{array}{l} (\text{pp}, \text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\kappa); b \xleftarrow{\$} \{0, 1\} \\ \{c_i\}_{i \in [\tau]} \leftarrow \text{Enc}(\text{pk}, \{m_i\}_{i \in [\tau]}) \\ (\{\hat{c}_i\}_{i \in [\tau]}, \pi^{(0)}) \leftarrow \text{Prover}(\text{pp}, \text{pk}, \{c_i\}_{i \in [\tau]}) \\ (\pi^{(1)}) \leftarrow \mathcal{S}(\text{pp}, \text{pk}, \{c_i\}_{i \in [\tau]}, \{\hat{c}_i\}_{i \in [\tau]}) \\ b' \leftarrow \text{Adv}(\text{pp}, \text{pk}, \{c_i\}_{i \in [\tau]}, \{\hat{c}_i\}_{i \in [\tau]}, \pi^{(b)}) \end{array} \right] - \frac{1}{2} \right| \leq \epsilon(\lambda),$$

where the probability is taken over **KeyGen**, **Enc**, **Prover**, \mathcal{S} and **Adv**.

B Security of Distributed Decryption

Here we define the syntax and security properties for a PKE with distributed decryption [ABGS23].

Definition 8 (PKE with Distributed Decryption). A PKE scheme with distributed decryption consists of five algorithms: key generation (**KeyGen**), encryption (**Enc**), decryption (**Dec**), distributed decryption (**DDec**), and combine (**Comb**), where

KeyGen. On input security parameter 1^λ and number of key-shares ξ_2 , outputs public parameters **pp**, a public key **pk**, a secret key **sk**, and key-shares $\{\text{sk}_j\}$,

Enc. On input pk and messages $\{m_i\}$, outputs ciphertexts $\{c_i\}$,

Dec. On input sk and ciphertexts $\{c_i\}$, outputs messages $\{m_i\}$,

DDec. On input a secret key share sk_{j^*} and ciphertexts $\{c_i\}$, outputs decryption shares $\{\text{ds}_{i,j^*}\}$,

Comb. On input ciphertexts $\{c_i\}$ and decryption shares $\{\text{ds}_{i,j}\}$, outputs either messages $\{m_i\}$ or \perp ,

and pp are implicit inputs to Enc, Dec, DDec and Comb.

Definition 9 (Chosen Plaintext Security). We say that the public key encryption scheme is secure against *chosen plaintext attacks* if an adversary Adv , after choosing two messages m_0 and m_1 and receiving an encryption c of either m_0 or m_1 (chosen at random), cannot distinguish which message c is an encryption of. Hence, we want that

$$\left| \Pr \left[b = b' : \begin{array}{l} (\text{pp}, \text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\kappa) \\ (m_0, m_1, \text{st}) \leftarrow \text{Adv}(\text{pp}, \text{pk}) \\ b \stackrel{\$}{\leftarrow} \{0, 1\}, c \leftarrow \text{Enc}(\text{pk}, m_b) \\ b' \leftarrow \text{Adv}(c, \text{st}) \end{array} \right] - \frac{1}{2} \right| \leq \epsilon(\lambda),$$

where the probability is taken over KeyGen and Enc .

Definition 10 (Threshold Correctness). We say that the public key distributed encryption scheme is *threshold correct* with respect to $P_{\text{sk}}(\cdot)$ if the following probability equals 1:

$$\Pr \left[\begin{array}{l} \text{Comb}(\{c_i\}_{i \in [\tau]}, \{\text{ds}_{i,j}\}_{i \in [\tau], j \in [\xi_2]}) \\ = \\ \text{Dec}(\text{sk}, \{c_i\}_{i \in [\tau]}) \end{array} : \begin{array}{l} (\text{pp}, \text{pk}, \text{sk}, \{\text{sk}_j\}_{j \in [\xi_2]}) \leftarrow \text{KeyGen}(1^\lambda, \xi_2) \\ \{c_1, \dots, c_\tau\} \leftarrow \mathcal{A}(\text{pp}, \text{pk}) \\ \forall i \in [\tau] : P_{\text{sk}}(c_i) = 1, \forall j \in [\xi_2] : \\ \{\text{ds}_{i,j}\}_{i \in [\tau]} \leftarrow \text{DDec}(\text{sk}_j, \{c_i\}_{i \in [\tau]}) \end{array} \right],$$

where the probability is taken over KeyGen and DDec .

Definition 11 (Threshold Verifiability). A PKE scheme with distributed decryption is *threshold verifiable* with respect to $P_{\text{sk}}(\cdot)$ if an adversary \mathcal{A} corrupting $J \subseteq [\xi_2]$ secret key shares $\{\text{sk}_j\}_{j \in J}$ cannot convince Comb to accept maliciously created decryption shares $\{\text{ds}_{i,j}\}_{i \in [\tau], j \in J}$. More concretely, the following probability is bounded by a negligible $\epsilon(\lambda)$:

$$\Pr \left[\begin{array}{l} \text{Dec}(\text{sk}, \{c_i\}_{i \in [\tau]}) \\ \neq \\ \text{Comb}(\{c_i\}_{i \in [\tau]}, \{\text{ds}_{i,j}\}_{i \in [\tau], j \in [\xi_2]}) \\ \neq \\ \perp \end{array} : \begin{array}{l} (\text{pp}, \text{pk}, \text{sk}, \{\text{sk}_j\}_{j \in [\xi_2]}) \leftarrow \text{KeyGen}(1^\lambda, \xi_2) \\ (\{c_1, \dots, c_\tau\}) \leftarrow \mathcal{A}(\text{pp}, \text{pk}, \{\text{sk}_j\}_{j \in J}) \\ \forall i \in [\tau] : P_{\text{sk}}(c_i) = 1, \forall j \notin J : \\ \{\text{ds}_{i,j}\}_{i \in [\tau]} \leftarrow \text{DDec}(\text{sk}_j, \{c_i\}_{i \in [\tau]}) \\ \{\text{ds}_{i,j}\}_{i \in [\tau], j \in J} \leftarrow \mathcal{A}(\{\text{ds}_{i,j}\}_{i \in [\tau], j \notin J}) \end{array} \right],$$

where the probability is taken over KeyGen and DDec .

Definition 12 (Distributed Decryption Simulatability). A PKE scheme with distributed decryption is *simulatable* with respect to $P_{\text{sk}}(\cdot)$ if an adversary \mathcal{A} corrupting $J \subsetneq [\xi_2]$ secret key shares $\{\text{sk}_j\}_{j \in J}$ cannot distinguish the transcript of the decryption protocol from a simulation by a simulator Sim which only gets $\{\text{sk}_j\}_{j \in J}$ as well as correct decryptions as input. More concretely, the following probability is bounded by a negligible $\epsilon(\text{sec})$:

$$\left| \Pr \left[b = b' : \begin{array}{l} (\text{pp}, \text{pk}, \text{sk}, \{\text{sk}_j\}_{j \in [\xi_2]}) \leftarrow \text{KeyGen}(1^\lambda, \xi_2) \\ (\{c_1, \dots, c_\tau\}) \leftarrow \mathcal{A}(\text{pp}, \text{pk}, \{\text{sk}_j\}_{j \in J}) \\ \forall i \in [\tau] : P_{\text{sk}}(c_i) = 1 \\ \{\text{ds}_{i,j}^0\} \leftarrow \text{DDec}(\{\text{sk}_j\}_{j \in [\xi_2]}, \{c_i\}_{i \in [\tau]}) \\ \{\text{ds}_{i,j}^1\} \leftarrow \text{Sim}(\text{pp}, \{\text{sk}_j\}_{j \in J}, \{c_i, \text{Dec}(\text{sk}, c_i)\}_{i \in [\tau]}) \\ b \stackrel{\$}{\leftarrow} \{0, 1\}, b' \leftarrow \mathcal{A}(\{\text{ds}_{i,j}^b\}_{i \in [\tau], j \in [\xi_2]}) \end{array} \right] - \frac{1}{2} \right|,$$

where the probability is taken over KeyGen , DDec , Sim .

C Parameter Constraints

Here we describe the parameters used in our electronic voting scheme. Table 6 lists these and makes explicit any constraints that apply to them. These constraints inform the choice of concrete values computed in Section 5.2.

Table 6: System parameters and constraints.

Parameter	Explanation	Constraints
λ	Computational security parameter	≥ 128
sec	Statistical security parameter	≥ 40
d	Ring dimension of R_p and R_q	d a power of two
p	Plaintext modulus	p a small prime
q	Ciphertext and commitment modulus	Prime $q = 1 \pmod{2d}$ s.t. $\ \sum_{j \in \xi_2} ds_j\ _\infty \leq \lfloor q/2 \rfloor$
t	KeyGen _{NTRU} rejection parameter	Set for rej. prob. $< 1/1000$ (Lemma 1)
k	Length of binding vector in BDLOP commitment	$k > 2$
\mathcal{C}	Challenge space for linear ZK proofs of commitments	$\mathcal{C} = \{c \in R_q \mid \ c\ _\infty = 1, \ c\ _1 = \kappa\}$
κ	Maximum ℓ_1 -norm of elements in \mathcal{C}	$2^\kappa \cdot \binom{d}{\kappa} > 2^\lambda$
ξ_1, ξ_2	Number of shuffle and decryption-servers	At least two servers
B_{Com}	Bound on the commitment noise	So that SIS is hard
B_{Dec}	Noise in ciphertext	$B_{\text{Dec}} = p \cdot d \cdot t \cdot \sigma_{\text{NTRU}} \cdot (2\xi_1 \cdot \nu + 1/2)$
B_{Drown}	Infinity norm of noise drowning term E_{ij}	$B_{\text{Drown}} = 2^{\text{sec}}(B_{\text{Dec}}/p\xi_2)$
σ_{NTRU}	Standard deviation for encryption secret key	So that NTRU is hard
ν	Bound on encryption randomness	So that LWE is hard
σ_{Com}	Standard deviation in ZK proofs of linear relations	Chosen to be $\sigma_{\text{Com}} = \kappa \cdot B_{\text{Com}} \cdot \sqrt{kd}$
τ	Total number of messages/number of voters	For soundness we need $(\tau^\delta + 1)/ R_q < 2^{-\lambda}$
η	Reed-Solomon encoding randomness length	Make soundness $\geq 2^{-\lambda}$ in Π_{Small} and Π_{Bnd}
ℓ_{Small}	Proof batch size in Π_{Small}	Same secret length as in [ABGS23]
ℓ_{Bnd}	Proof batch size in Π_{Bnd}	Same secret length as in [ABGS23]
μ_{Small}	Reed-Solomon message length in Π_{Small}	$\mu_{\text{Small}} = (k + 2) \cdot d + \eta$
μ_{Bnd}	Reed-Solomon message length in Π_{Bnd}	$\mu_{\text{Bnd}} = (k + 1) \cdot d + \eta$
μ'_{Small}	Reed-Solomon message dimension in Π_{Small}	$\mu_{\text{Small}} \leq \mu'_{\text{Small}} \leq \gamma < q$
μ'_{Bnd}	Reed-Solomon message dimension in Π_{Small}	$\mu_{\text{Bnd}} \leq \mu'_{\text{Bnd}} \leq \gamma < q$
γ_{Small}	Reed-Solomon code length in Π_{Bnd}	$\mu_{\text{Small}} \leq \mu'_{\text{Small}} \leq \gamma < q$
γ_{Bnd}	Reed-Solomon code length in Π_{Bnd}	$\mu_{\text{Bnd}} \leq \mu'_{\text{Bnd}} \leq \gamma < q$