Verifiable Mix-Nets and Distributed Decryption for Voting from Lattice-Based Assumptions*

Diego F. Aranha
dfaranha@cs.au.dk
Aarhus University
Aarhus, Denmark

Carsten Baum
cabau@dtu.dk
DTU Compute
Copenhagen, Denmark

Kristian Gjøsteen
kristian.gjosteen@ntnu.no
Norwegian University of Science and Technology
Trondheim, Norway

Tjerand Silde†
tjerand.silde@ntnu.no
Norwegian University of Science and Technology
Trondheim, Norway

ABSTRACT

Cryptographic voting protocols have recently seen much interest from practitioners due to their (planned) use in countries such as Estonia, Switzerland, France, and Australia. Practical protocols usually rely on tested designs such as the mixing-and-decryption paradigm. There, multiple servers verifiably shuffle encrypted ballots, which are then decrypted in a distributed manner. While several efficient protocols implementing this paradigm exist from discrete log-type assumptions, the situation is less clear for post-quantum alternatives such as lattices. This is because the design ideas of the discrete log-based voting protocols do not carry over easily to the lattice setting, due to specific problems such as noise growth and approximate relations.

This work proposes a new verifiable secret shuffle for BGV ciphertexts and a compatible verifiable distributed decryption protocol. The shuffle is based on an extension of a shuffle of commitments to known values which is combined with an amortized proof of correct re-randomization. The verifiable distributed decryption protocol uses noise drowning, proving the correctness of decryption steps in zero-knowledge. Both primitives are then used to instantiate the mixing-and-decryption electronic voting paradigm from lattice-based assumptions.

We give concrete parameters for our system, estimate the size of each component and provide implementations of all important subprotocols. Our experiments show that the shuffle and decryption protocol is suitable for use in real-world e-voting schemes.

CCS CONCEPTS

• **Security and privacy → Public key (asymmetric) techniques**; **Privacy-preserving protocols**.

*The full version of this paper is available at eprint.iacr.org/2022/422.pdf.
†Work done in part while visiting Aarhus University.

CCS '23, November 26–30, 2023, Copenhagen, Denmark

ACM ISBN 979-8-4007-0050-7/23/11.
https://doi.org/10.1145/3576915.3616683

KEYWORDS

Lattice-Based Cryptography, Electronic Voting, Implementation

ACM Reference Format:
Diego F. Aranha, Carsten Baum, Kristian Gjøsteen, and Tjerand Silde. 2023. Verifiable Mix-Nets and Distributed Decryption for Voting from Lattice-Based Assumptions. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23), November 26–30, 2023, Copenhagen, Denmark.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3576915.3616683

1 INTRODUCTION

Mix-nets were originally proposed for anonymous communication [14], but have since been used extensively for cryptographic voting systems. A mix-net is a multi-party protocol that gets as input a collection of ciphertexts and outputs another collection of ciphertexts whose decryption is the same set, up to order. It guarantees that the permutation between input and output ciphertexts is hidden if at least one party is honest, while none of the servers involved learns the plaintexts.

Mix-nets are commonly used in cryptographic voting. Here, encrypted ballots are submitted to a bulletin board or ballot box with identifying information attached. These ciphertexts are then sent through a mix-net before decryption, to break the identity-ballot correlation. In addition to hiding the permutation, the correctness of the mix-net output must be verifiable. In applications such as voting it is important that the mix-net provides a proof of correctness that can be verified by anyone at any later point in time, to ensure universal verifiability.

A *shuffle* of a set of ciphertexts is another set of ciphertexts whose decryption is the same as the original set, up to order. Compared to a mix-net, it is performed by one server only (which does not learn the plaintexts). As for mix-nets, a shuffle is *secret* if it is hard for any external party to correlate input and output ciphertexts, and *verifiable* if there is a proof that decryptions are the same.

If we have a verifiable secret shuffle for some cryptosystem, then this can be used to construct a mix-net: for this, the nodes of the mix-net receive a set of ciphertexts as input, shuffle them sequentially and each provides a proof of correctness. The mix-net proof then consists of the intermediate ciphertexts along with the shuffle proofs. If at least one node in this chain is honest, it is hard to correlate the inputs and outputs.

For applications in cryptographic voting, it must also be guaranteed that the correct result can be obtained from the mix-net output, while nobody has the secret decryption key. One strategy is to use verifiable threshold decryption, where the key is secret-shared among a committee of decryption parties. Each of them contributes to the decryption and proves that they did so honestly.

Verifiable shuffling and verifiable distributed decryption protocols are well-known for cryptosystems based on discrete logarithm-type assumptions. For example, Neff [31] proposed the first efficient verifiable secret shuffle for ElGamal-like cryptosystems. Verifiable decryption for ElGamal-like cryptosystems can be constructed using standard Verifiable Secret-Sharing and Σ-protocols.

While ample voting schemes have been constructed based on the aforementioned outline, they essentially all rely on assumptions that are not secure against quantum computers. Given the need for the long-term privacy of elections, it is important to construct verifiable shuffles and distributed decryption from quantum-safe computational problems such as lattice assumptions. NIST recently standardized post-quantum key-encapsulation mechanisms and digital signatures based on lattices [28, 33, 35]. Using shuffles and distributed decryption schemes based on the same assumptions, it seems well-motivated to build a plausibly post-quantum voting scheme following the aforementioned approach. The main obstacle to simply adopting the protocols for discrete logarithms to lattices is the (presumed) lack of suitable efficient techniques for verification, as well as the problem of noise growth.

1.1 Our contributions

In this work, we make progress in the direction of plausibly quantum-secure voting. We design a verifiable secret shuffle for BGV ciphertexts [13] that is suitable for cryptographic voting systems dealing with arbitrary vote structures. In addition, we construct a verifiable distributed decryption protocol by compiling previous passively-secure constructions with zero-knowledge proofs and show how to integrate these and other building blocks into a voting scheme. Finally, we implemented the main parts of the verifiable shuffle and distributed decryption protocols to demonstrate the viability and efficacy of our overall design.

Lattice-based shuffle. To construct a mix-net for BGV ciphertexts we extend the shuffle of commitments to known values by Aranha et al. [4]. Their construction only works for BDLOP [7] commitments of message length 1, while we generalize their construction to an arbitrary length. Given such a generalized verifiable shuffle of commitment openings, our verifiable shuffle for input ciphertexts $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_\tau$ then works as follows: We let the shuffler commit to BGV re-randomization ciphertexts $\hat{\boldsymbol{c}}_1, \ldots, \hat{\boldsymbol{c}}_\tau$ (encryptions of 0) using the linearly homomorphic BDLOP commitment scheme $\mathtt{Com}$. Together with efficient proofs of well-formedness for the committed re-randomization ciphertexts, this gives us a verifiable shuffle:

(1) The shuffler commits to the re-randomization ciphertexts $\hat{\boldsymbol{c}}_1, \ldots, \hat{\boldsymbol{c}}_\tau$ as $\mathtt{Com}(\hat{\boldsymbol{c}}_i)$ and shows that they are well-formed using zero-knowledge proofs.
(2) The shuffler computes $\boldsymbol{d}_i = \boldsymbol{c}_i + \hat{\boldsymbol{c}}_i$ and sends shuffled elements $L = (\boldsymbol{d}_{\pi(i)})_{i \in [\tau]}$ to the receiver.
(3) Finally, the prover shows that L is a list of openings of the commitments obtained from $\boldsymbol{c}_i + \mathtt{Com}(\hat{\boldsymbol{c}}_i)$ using the extended shuffle of commitments to known values.

To prove the well-formedness of the ciphertexts, we utilize proofs of shortness where the proof size is sublinear in the number of ciphertexts τ. For this, we use a version of recent amortized proofs of shortness [10]. Unfortunately their construction as-is is suboptimal for our setting, so we adapt and re-prove their protocol.

Verifiable distributed decryption. As explained, a verifiable secret shuffle on its own is usually not sufficient to build a cryptographic voting system. The ciphertexts must also be decrypted, without introducing correctness and privacy problems. Our solution is to distribute the decryption operation in a verifiable way. We hand out key shares of the secret decryption key to each decryption server, and all of them perform a partial decryption of each ciphertext. In addition, we publish commitments to the key shares. The decryption servers then add noise to the partial decryption to hide information about their shares, in a process called noise drowning. Finally, decryption servers publish the partial decryptions together with a proof of correctness of the decryption, and the plaintexts are computed in public by combining all the partial decryptions.

We use a decryption protocol for BGV ciphertexts that is similar to existing works such as [18]. Their construction is only passively secure. We, therefore, modify the protocol to be resistant to active attacks even if all decryption servers are malicious, and prove it secure. For this, we again utilize an (amortized) zero-knowledge proof of shortness that allows each decryption server to show that it behaved honestly during decryption.

Putting things together. Lattice-based cryptography is very delicate, and one has to be cautious when combining multiple sub-protocols into a larger (voting) construction. This is mainly due to *noise* in ciphertexts, which can lead to faulty decryptions, overly large parameters, or both.

In our construction, each shuffle adds noise to the ciphertexts, which means that to ensure the correctness of decryption we need to choose parameters based on the number of shuffles and the amount of noise added in each shuffle. Each partial decryption also adds noise to the ciphertexts to hide the secret key. Because of the noise drowning technique, the norm must be quite large, influencing the choice of parameters for the overall construction as well as the choice of zero-knowledge proof techniques involved.

In particular, it is important when measuring performance to use parameters suitable for the complete system, not parameters optimized for individual components only. In order to provide proper context for our contributions, we give a sketch of a full cryptographic voting protocol and provide example parameters. A simplified version could be used as a quantum-safe Helios [1] variant.

Implementation results. Our example parameters assume 4 mix-nodes and 4 decryption nodes. We have estimated the size of each component with respect to the parameters for the full protocol in addition to implementing all sub-protocols, showing that it can be used for large-scale real-world elections where ballots typically are counted and verified in batches of several thousand.

To summarize our implementation results, a ciphertext ballot is of size 80 KB (encoding a vote of size 4096 bits), each mixing proof is

of total size 370τ KB and each decryption proof is of total size 157τ KB, where τ is the number of total ciphertexts. It takes only 0.74 ms to encrypt a ballot, while the mixing proof takes 158.4τ ms and the decryption proof takes 138.11τ ms. Verification is much faster, with only 12.9τ ms for the mixing and 30.2τ ms for the decryption. These results improve on the state of the art considerably, see Section 7.

Quantum security. While our work constructs and implements a voting scheme from post-quantum assumptions, we do *not* claim that it is post-quantum secure. We discuss this in Appendix B.

1.2 Related work

Aranha et al. [4] provide a verifiable shuffle of known commitment openings together with concrete parameters and implementation of a complete voting protocol. However, their trust model has the limitation that the ballot box and the shuffle server must not collude to ensure the privacy of the ballots, which is too restrictive for most real-world settings. This is inherent for the protocol which can not easily be extended to several shuffles unless layered encryption is used, and this would heavily impact the performance.

Costa et al. [16] design a shuffle with a straight-forward approach similar to Neff [31] based on roots of polynomials. Their protocol requires committing to two evaluations of a polynomial and then proving the correctness of the evaluation using a sequence of multiplication proofs which are quite costly in practice. Farzaliyev et al. [22] implements the mix-net by Costa et al. [16] using the amortization techniques by Attema et al. [5] for the commitment scheme by Baum et al. [7]. Here, the proof size is approximately 14 MB per voter, a factor 40 larger than our shuffle proof, even for a smaller parameter set that does not take into account distributed decryption afterward. We expect our shuffle proof to be an additional factor 10 smaller than what we presented above with optimal parameters for the shuffle only ($q \approx 2^{32}$ and $N = 1024$ instead of $q \approx 2^{78}$ and $N = 4096$). Furthermore, their proof generation and verification respectively take 1.54 and 1.51 second per vote, which is approximately 17 times slower than it takes to produce and verify our shuffle proof in sequence (when normalizing for clock frequency), with parameters that do not take decryption into account.

Recently, Herranz et al. [26] gave a new proof of correct shuffle based on Beneš networks and sub-linear lattice-based proofs for arithmetic circuit satisfiability. However, the scheme is not implemented and the example parameters do not take the soundness slack of the amortized zero-knowledge proofs into account. Moreover, [26] does not consider the decryption of ballots, which would heavily impact the parameters of their protocol in practice.

A completely different approach to mix-nets is the so-called decryption mix-nets. The idea is that the input ciphertexts are actually nested encryptions. Each node in the mix-net is then responsible for decrypting one layer of each ciphertext. These can be made fully generic, relying only on public key encryption. Boyen et al. [11] carefully adapt these ideas to lattice-based encryption, resulting in a very fast scheme. Decryption mix-nets are well-suited to applications in anonymous communication. However, for voting applications, they are often less well-suited due to their trust requirements. An important goal for cryptographic voting is universal verifiability: after the election is done, anyone should be able to verify that the ballot decryption was done correctly without needing to trust anyone. This trust issue generalizes to any situation where it is necessary to convince someone that a shuffle has been performed correctly, but no auditor is available. Fast and generic decryption mix-nets such as Boyen et al. [11] need an auditor (potentially distributed) to verify the mix-net, but then it must be trusted during the operation. This conflicts with universal verifiability.

del Pino et al. [19] give a practical voting protocol based on homomorphic counting. They only support yes/no elections, and the total size depends directly on the number of candidates for larger elections. It was shown by Boyen et al. [12] that the protocol in [19] is not end-to-end verifiable unless all tallying authorities and all voters' voting devices are honest. This problem is solved by [12], but their construction still has the downside of only supporting homomorphic tallying. Strand [38] built a verifiable shuffle for the GSW cryptosystem, but this construction is too restrictive for practical use. Chillotti et al. [15] uses fully homomorphic encryption, which for the foreseeable future is most likely not efficient enough to be considered for practical deployment.

2 BUILDING BLOCKS

In this section, we define the building blocks that we use in our construction of the voting scheme. Then, in Section 3 we show how these can be put together.

Let κ be the computational and sec the statistical security parameter. We define the ring $R_q = \mathbb{Z}_q[X]/\langle X^N + 1\rangle$, its norms, the discrete Gaussian distribution $\mathcal{N}$, rejection sampling, and knapsack problems $\mathrm{SKS}^2_{n,k,\beta}$ and $\mathrm{DKS}^\infty_{n,k,\beta}$ in the full version of the paper.

We use $S_B \subseteq R_q$ to denote the subset of R_q where each coefficient is less or equal B.

2.1 PKE with Distributed Decryption

We first present a definition of a secure public key encryption (PKE) scheme with a distributed decryption protocol. Such a scheme works like a regular PKE scheme but allows the secret key to be shared among decryption servers. Then, for a given ciphertext, the decryption servers can compute decryption shares using their key shares which, when combined, reveal the plaintext. The goal here is that the decryption shares do not reveal information about the secret key shares.

Definition 1 (PKE with Distributed Decryption). *A PKE scheme with distributed decryption consists of five algorithms: key generation* (KGen)*, encryption* (Enc)*, decryption* (Dec)*, distributed decryption* (DDec)*, and combine* (Comb)*, where*

- KGen *On input security parameter* 1^κ *and number of key-shares* ξ_2*, outputs public parameters* pp*, a public key* pk*, a secret key* sk*, and key-shares* $\{\mathsf{sk}_j\}$*,*
- 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* $\{\mathsf{ds}_{i,j^*}\}$*,*
- Comb *On input ciphertexts* $\{c_i\}$ *and decryption shares* $\{\mathsf{ds}_{i,j}\}$*, outputs either messages* $\{m_i\}$ *or* $\perp$*,*

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

For such a scheme, we require multiple security properties. (Threshold) correctness and IND-CPA security are standard and we only

provide their definitions for completeness in the full version of the paper.

Threshold verifiability and decryption simulatability are of more interest, which we define below.

Let $P_{\mathsf{sk}}(c)$ be an efficiently computable predicate that on input secret key sk and a ciphertext c outputs 1 or 0. Such a predicate signals that the ciphertext is reliably decryptable - which we need to consider as ciphertexts contain noise. We first define threshold verifiability, which models that distributed decryption is secure against active attacks.

Definition 2 (Threshold Verifiability). *A PKE scheme with distributed decryption is* threshold verifiable *with respect to* $P_{\mathsf{sk}}(\cdot)$ *if an adversary* Adv *corrupting* $J \subseteq [\xi_2]$ *secret key shares* $\{\mathsf{sk}_j\}_{j\in J}$ *cannot convince* Comb *to accept maliciously created decryption shares* $\{\mathsf{ds}_{i,j}\}_{i\in[\tau],j\in J}$. *More concretely, the following probability is bounded by a negligible* $\epsilon(\kappa)$*:*

$$\Pr\left[\begin{array}{c} \mathsf{Dec}(\mathsf{sk},\{c_i\}_{i\in[\tau]}) \\ \neq \\ \mathsf{Comb}(\{c_i\}_{i\in[\tau]},\{\mathsf{ds}_{i,j}\}_{i\in[\tau]}^{j\in[\xi_2]}) : \\ \neq \\ \bot \end{array}\middle|\begin{array}{c} (\mathsf{pp},\mathsf{pk},\mathsf{sk},\{\mathsf{sk}_j\}_{j\in[\xi_2]}) \leftarrow \mathsf{KGen}(1^\kappa,\xi_2) \\ (\{c_1,\ldots,c_\tau\}) \leftarrow \mathsf{Adv}(\mathsf{pp},\mathsf{pk},\{\mathsf{sk}_j\}_{j\in J}) \\ \forall i\in[\tau]:\ P_{\mathsf{sk}}(c_i)=1,\forall j\notin J: \\ \{\mathsf{ds}_{i,j}\}_{i\in[\tau]} \leftarrow \mathsf{DDec}(\mathsf{sk}_j,\{c_i\}_{i\in[\tau]}) \\ \{\mathsf{ds}_{i,j}\}_{i\in[\tau],j\in J} \leftarrow \mathsf{Adv}(\{\mathsf{ds}_{i,j}\}_{i\in[\tau],j\notin J}) \end{array}\right],$$

where the probability is taken over KGen *and* DDec.

We moreover define a simulation property, that shows that decryption shares do not leak any information about the secret key. This models security against passive attackers.

Definition 3 (Distributed Decryption Simulatability). *A PKE scheme with distributed decryption is* simulatable *with respect to* $P_{\mathsf{sk}}(\cdot)$ *if an adversary* Adv *corrupting* $J \subsetneq [\xi_2]$ *secret key shares* $\{\mathsf{sk}_j\}_{j\in J}$ *cannot distinguish the transcript of the decryption protocol from a simulation by a simulator* Sim *which only gets* $\{\mathsf{sk}_j\}_{j\in J}$ *as well as correct decryptions as input. More concretely, the following probability is bounded by a negligible* $\epsilon(\mathsf{sec})$*:*

$$|\Pr\left[b=b' \ :\ \begin{array}{c} (\mathsf{pp},\mathsf{pk},\mathsf{sk},\{\mathsf{sk}\}_{j\in[\xi_2]}) \leftarrow \mathsf{KGen}(1^\kappa,\xi_2) \\ (\{c_1,\ldots,c_\tau\}) \leftarrow \mathsf{Adv}(\mathsf{pp},\mathsf{pk},\{\mathsf{sk}_j\}_{j\in J}) \\ \forall i\in[\tau]:\ P_{\mathsf{sk}}(c_i)=1 \\ \{\mathsf{ds}^0_{i,j}\} \leftarrow \mathsf{DDec}(\{\mathsf{sk}_j\}_{j\in[\xi_2]},\{c_i\}_{i\in[\tau]}) \\ \{\mathsf{ds}^1_{i,j}\} \leftarrow \mathsf{Sim}(\mathsf{pp},\{\mathsf{sk}_j\}_{j\in J},\{c_i,\mathsf{Dec}(\mathsf{sk},c_i)\}_{i\in[\tau]}) \\ b \xleftarrow{\$} \{0,1\}, b' \leftarrow \mathsf{Adv}(\{\mathsf{ds}^b_{i,j}\}_{i\in[\tau],j\in[\xi_2]}) \end{array}\right] - \frac{1}{2}|,$$

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

2.1.1 Our instantiation. Let $p \ll q$ be primes, define R_q and R_p for a fixed N, let $B_{\mathsf{Key}}, B_{\mathsf{Err}} \in \mathbb{N}$ be bounds. We use the BGV [13] encryption scheme, which consists of three algorithms: key generation (KGen), encryption (Enc) and decryption (Dec), where:

KGen Samples a uniform element $a \xleftarrow{\$} R_q$, a short $s \xleftarrow{\$} S_{B_{\mathsf{Key}}}$ and noise $e \xleftarrow{\$} S_{B_{\mathsf{Err}}}$. The algorithm outputs the public key $\mathsf{pk} = (a,b) = (a, as+pe)$ and secret key $\mathsf{sk} = s$.

Enc On input the public key $\mathsf{pk} = (a,b)$ and a message $m \in R_p$, samples a uniform $r \xleftarrow{\$} S_{B_{\mathsf{Key}}}$, noise $e', e'' \xleftarrow{\$} S_{B_{\mathsf{Err}}}$ and outputs the ciphertext $c = (u,v) = (ar+pe', br+pe''+m)$.

Dec On input secret key $\mathsf{sk} = s$ and ciphertext $c = (u,v)$, outputs message $m = (v - su \mod q) \mod p$.

The following theorem follows from [13] and [30].

Theorem 1. *The BGV encryption scheme is* correct *if* $\|v - su\|_\infty \leq B_{\mathsf{Dec}} < \lfloor q/2 \rfloor$*, and IND-CPA secure if the* $\mathsf{DKS}^\infty_{N,2,\beta}$ *problem is hard for some* $\beta = \beta(N,q,B_{\mathsf{Key}},B_{\mathsf{Err}},p)$.

We use this theorem to define the predicate $P_{\mathsf{sk}}(u,v)$ to be 1 iff $\|v - su\|_\infty < B_{\mathsf{Dec}}$ and otherwise 0. Since each ciphertext consists of 2 elements from R_q, it can be represented using $2N\log_2 q$ bits.

2.1.2 Threshold decryption. We quickly recap the passively secure distributed decryption protocol by Damgård *et al.* [8, 17, 18]. Here, the KGen algorithm on input $\xi_2 \in \mathbb{N}$ additionally outputs uniformly random shares $\mathsf{sk}_j = s_j$ of the secret key $\mathsf{sk} = s$ such that $s = s_1 + \cdots + s_{\xi_2}$ in R_q. This defines a passively secure threshold decryption protocol by using the linearity of the decryption function:

DDec On input a secret key-share $\mathsf{sk}_j = s_j$ and a ciphertext $c = (u,v)$, does the following:

(1) Compute $m_j = s_j u$ and sample a uniformly random $E_j \xleftarrow{\$} R_q$ such that $\|E_j\|_\infty \leq 2^{\mathsf{sec}}(B_{\mathsf{Dec}}/p\xi_2)$ for statistical security parameter sec and noise-bound B_{Dec},

(2) Output $\mathsf{ds}_j = t_j = m_j + pE_j$.

Comb On input ciphertext $c = (u,v)$ and set of decryption shares $\{\mathsf{ds}_j = t_j\}_{j\in[\xi_2]}$, outputs message $m = (v - t \mod q) \mod p$, where $t = t_1 + \cdots + t_{\xi_2}$.

The following theorem follows from [17, 18].

Theorem 2. *Let* sec *be the statistical security parameter. The distributed BGV encryption scheme is* correct *for input ciphertexts with* $\|v - us\|_\infty \leq (1+2^{\mathsf{sec}})B_{\mathsf{Dec}} < \lfloor q/2 \rfloor$*, and is* decryption simulatable *against passive adversaries (i.e fulfills Definition 3).*

Each partial decryption consists of one element from R_q, namely the output of DDec, which means that the output from the passively secure protocol is of size $N\log_2 q$ bits per party.

This scheme is not secure against active adversaries, i.e. it does not have threshold verifiability. We, therefore, modify it in Section 6 to withstand active attacks.

2.2 Commitments

Commitment schemes were first introduced by Blum [9], and we use these at multiple points in this work to achieve verifiability.

Definition 4 (Commitment Scheme). *A commitment scheme consists of three algorithms: key generation* (Setup)*, commitment* (Com) *and opening* (Open)*, where*

Setup *On input security parameter* 1^κ*, outputs public params* pp,

Com *On input message* m*, outputs commitment* c *and opening* r,

Open *On input* m, c *and* r*, outputs either* 0 *or* 1,

and the public parameters pp *are implicit inputs to* Com *and* Open.

For the commitment scheme, we require that it is correct, binding, and hiding. *Correctness* means that an honestly generated commitment is always accepted by the opening algorithm. *Binding* requires that no PPT adversary can provide two different valid openings of a given commitment for different messages. Hiding means that the commitment itself does not reveal any information about the committed value. We provide these definitions for completeness in the full version of the paper.

2.2.1 Our instantiation. Our work uses the BDLOP [7] commitment scheme. Let R_q be defined as above and let $\mathcal{N}_{\sigma_{\mathsf{C}}}$ be a Gaussian distribution with standard deviation σ_{C} and B_{Com} be a noise bound. The algorithms are defined as follows:

Setup Outputs a pk which allows to commit to length-l_c messages from $R_q^{l_c}$ using length-k randomness from $S^k_{B_{\text{Com}}}$ outputting length-$(n+l_c)$ vectors. For this, we define

$$\boldsymbol{A}_{\text{C},1} = \begin{bmatrix} \boldsymbol{I}_n & \widehat{\boldsymbol{A}}_{\text{C},1} \end{bmatrix} \quad \text{where } \widehat{\boldsymbol{A}}_{\text{C},1} \xleftarrow{\$} R_q^{n\times(k-n)}$$

$$\boldsymbol{A}_{\text{C},2} = \begin{bmatrix} \boldsymbol{0}^{l_c\times n} & \boldsymbol{I}_{l_c} & \widehat{\boldsymbol{A}}_{\text{C},2} \end{bmatrix} \quad \text{where } \widehat{\boldsymbol{A}}_{\text{C},2} \xleftarrow{\$} R_q^{l\times(k-n-l_c)}.$$

Let $\text{pk} = \boldsymbol{A}_{\text{C}} = \begin{bmatrix} \boldsymbol{A}_{\text{C},1} \\ \boldsymbol{A}_{\text{C},2} \end{bmatrix}$. $\boldsymbol{A}_{\text{C}}$ has height $n+l_c$ and width k.

Com On input $\boldsymbol{m} \in R_q^{l_c}$ samples $\boldsymbol{r}_{\boldsymbol{m}} \xleftarrow{\$} S^k_{B_{\text{Com}}}$ and computes

$$\text{Com}_{\text{pk}}(\boldsymbol{m}; \boldsymbol{r}_{\boldsymbol{m}}) = \boldsymbol{A}_{\text{C}} \cdot \boldsymbol{r}_{\boldsymbol{m}} + \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{m} \end{bmatrix} = \begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix} = [\![\boldsymbol{m}]\!].$$

Com outputs $[\![\boldsymbol{m}]\!]$ and the opening $\boldsymbol{d} = (\boldsymbol{m}, \boldsymbol{r}_{\boldsymbol{m}}, 1)$.

Open Verifies whether an opening $(\boldsymbol{m}, \boldsymbol{r}_{\boldsymbol{m}}, f)$, with $f \in \bar{C}$, is a valid opening of $[\![\boldsymbol{m}]\!]$ by checking that $\|\boldsymbol{r}_{\boldsymbol{m}}[i]\| \le 4\sigma_{\text{C}}\sqrt{N}$, for $i \in [k]$, and if

$$f \cdot \begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix} \stackrel{?}{=} \boldsymbol{A}_{\text{C}} \cdot \boldsymbol{r}_{\boldsymbol{m}} + f \cdot \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{m} \end{bmatrix}.$$

It outputs 1 if all conditions hold, and 0 otherwise.

We define the set $\bar{C}$ in the full version of the paper.

The openings generated by Com form a subset of those accepted by Open, which is necessary for efficient zero-knowledge proofs on BDLOP commitments. Observe that Open always accepts honestly generated openings (except with negligible probability) by setting $f = 1$. The following theorem follows from Baum et al. [7].

Theorem 3. *The aforementioned commitment scheme is* computationally hiding *if the* $\text{DKS}^{\infty}_{n+l_c,k,B_{\text{Com}}}$ *problem is hard, and the scheme is* computationally binding *if the* $\text{SKS}^2_{n,k,16\sigma_{\text{C}}\sqrt{\nu N}}$ *problem is hard.*

Each commitment consists of $n+l_c$ elements from R_q and can hence be represented using $(n+l_c)N\log_2 q$ bits.

2.3 Zero-Knowledge Proofs

Zero-Knowledge (ZK) proofs were first introduced by Goldwasser et al. [25]. They are cryptographic protocols to show that a certain statement is true, without revealing the witness. We use ZK proofs in our constructions to achieve verifiability: protocol participants show that they indeed followed the protocol steps correctly, while not revealing any secret randomness that they used in the process.

Let L be a language, and let R be an NP-relation on L. Then, x is an element in L if there exists a witness w such that $(x,w) \in \text{R}$. We let $\mathcal{P}$, $\mathcal{P}^*$, $\mathcal{V}$ and $\mathcal{V}^*$ be polynomial time algorithms.

Definition 5 (Interactive Proofs). *An interactive proof protocol* Π *consists of two parties: a prover* $\mathcal{P}$ *and a verifier* $\mathcal{V}$*, and a setup algorithm* (Setup)*, where* Setup*, on input the security parameter* 1^κ*, outputs public setup parameters* sp*. The protocol consists of a transcript* T *of the communication between* $\mathcal{P}$ *and* $\mathcal{V}$*, with respect to* sp*, and the conversation terminates with* $\mathcal{V}$ *outputting either* 1 *or* 0*. Let* $\langle \mathcal{P}(\text{sp}, x, w), \mathcal{V}(\text{sp}, x)\rangle$ *denote the output of* $\mathcal{V}$ *on input* x *after its interaction with* $\mathcal{P}$*, who holds a witness* w*.*

We call an Interactive Proof a Zero-Knowledge proof[1] if it has the following three properties:

Completeness: If $\mathcal{P}$ has a valid witness w such that $(x,w) \in \text{R}$, then $\mathcal{V}$ accepts.

Knowledge Soundness: If $\mathcal{P}^*$ can make an honest verifier accept with large enough probability for statement x, then there exists a polynomial-time algorithm E that can, through black-box access to $\mathcal{P}^*$, extract w such that $(x,w) \in \text{R}$.

Honest Verifier Zero Knowledge: There exists a PPT algorithm S, called *simulator*, that given only x can create transcripts whose distribution is indistinguishable from those of an honest prover and verifier.

We give the formal definitions in the full version of the paper.

Note that an interactive honest-verifier zero-knowledge proof protocol can be made non-interactive using the Fiat-Shamir transform [23].

2.3.1 Linear relations among commitments. Assume that there are $\hat{n}$ BDLOP commitments

$$[\![\boldsymbol{m}_i]\!] = \begin{bmatrix} \boldsymbol{c}_{i,1} \\ \boldsymbol{c}_{i,2} \end{bmatrix}, \text{ for } 1 \le i \le \hat{n} \text{ where } \boldsymbol{c}_{i,2} \in R_q^{l_c}.$$

For the public scalar vector $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_{\hat{n}-1}) \in R_q^{\hat{n}-1}$ the prover wants to prove that the following relation holds:

$$\mathcal{R}_{\text{Lin}} = \left\{ (x,w) \;\middle|\; \begin{array}{c} x = (\text{pk}, \{[\![\boldsymbol{m}_i]\!]\}_{i\in[\hat{n}]}, \boldsymbol{\alpha}) \wedge \\ w = (f, \{\boldsymbol{m}_i, \boldsymbol{r}_i\}_{i\in[\hat{n}]}) \wedge \\ \forall i \in [\hat{n}] : \text{Open}_{\text{pk}}([\![\boldsymbol{m}_i]\!], \boldsymbol{m}_i, \boldsymbol{r}_i, f) = 1 \\ \wedge\, \boldsymbol{m}_{\hat{n}} = \sum_{i=1}^{\hat{n}-1} \alpha_i \boldsymbol{m}_i \end{array} \right\}.$$

We will require proof of this relation at multiple points in our constructions. In the full version of the paper

we provide a ZK proof Π_{Lin} for this relation, which is a directly extended version of the linearity proof in [7]. It works like a standard Σ protocol when adapted to lattices.

The relation $\mathcal{R}_{\text{Lin}}$ is relaxed because of the additional factor f in the opening, which appears in the soundness proof. It does not show up in protocol Π_{Lin}, because an honest prover uses $f = 1$. The bound is $B = 2\sigma_{\text{C}}\sqrt{N}$ and the protocol produces a proof transcript $\pi_{\text{Lin}} = ((\{\boldsymbol{t}_i\}_{i\in[\hat{n}]}, u), \beta, (\{\boldsymbol{z}_i\}_{i\in[\hat{n}]}))$. Using the standard Fiat-Shamir transform, we make Π_{Lin} non-interactive.

2.3.2 Amortized Proofs of Boundedness. It is well-known that polynomials in R_q can be represented as vectors in $\mathbb{Z}_q^N$ and multiplication by a polynomial $\hat{a}$ in R_q can be expressed as a matrix-vector product with a nega-cyclic matrix in $\mathbb{Z}_q^{N\times N}$. Let $\boldsymbol{A}$ be a publicly known $r \times v$ matrix over R_q, that is, a $rN \times vN$ matrix over $\mathbb{Z}_q$. We will now consider how to prove generically in zero-knowledge that $\boldsymbol{t}_i = \boldsymbol{A}\boldsymbol{s}_i$ for bounded $\boldsymbol{s}_i$ and known $\boldsymbol{t}_i$ over $\mathbb{Z}_q$. This is the same as proving correct multiplication over the ring R_q of the respective elements. We use proofs that are *amortized*, meaning that the proof size is sublinear in the number τ of individual statements that we prove. Both the BGV encryption and BDLOP commitment can be expressed in this form and require bounds on inputs for correctness, so this ZK proof can be used to show that encryptions or commitments were honestly made.

[1] More concretely, an Honest-Verifier Zero Knowledge Proof of Knowledge

Let $\boldsymbol{A}$ be a publicly known $r \times v$-matrix over R_q, let $\boldsymbol{s}_1, \boldsymbol{s}_2, \ldots, \boldsymbol{s}_\tau$ be bounded elements in R_q^v and let $\boldsymbol{A}\boldsymbol{s}_i = \boldsymbol{t}_i$ for $i \in [\tau]$. Letting $\boldsymbol{S}$ be the matrix whose columns are $\boldsymbol{s}_i$ and $\boldsymbol{T}$ be the same matrix for $\boldsymbol{t}_i$, but defined over $\mathbb{Z}_q^N$ instead of R_q, then [6] give an efficient amortized zero-knowledge proof of knowledge for the relation

$$\mathcal{R}_{\text{BND}} = \left\{ (x, w) \;\middle|\; \begin{array}{l} x = (\boldsymbol{A}, \boldsymbol{T}) \wedge w = \boldsymbol{S} \quad \wedge \forall i \in [\tau] : \\ \boldsymbol{t}_i = \boldsymbol{A}\boldsymbol{s}_i \wedge ||s_{i,j}||_2 \leq 2 \cdot B_{\text{BND}} \end{array} \right\}.$$

Let

$$\pi_{\text{BND}} \leftarrow \Pi_{\text{BND}}(\boldsymbol{S}; (\boldsymbol{A}, \boldsymbol{T}, \sigma_{\text{BND}})), 0 \vee 1 \leftarrow \Pi_{\text{BNDV}}((\boldsymbol{A}, \boldsymbol{T}, B_{\text{BND}}); \pi_{\text{BND}}),$$

denote the run of the proof and verification protocols, respectively, where the Π_{BND}-protocol, using Fiat-Shamir, produces a non-interactive proof of the form $\pi_{\text{BND}} = (\boldsymbol{C}, \boldsymbol{Z})$, where $\boldsymbol{C}$ is the output of a hash function, and the Π_{BNDV}-protocol verifies the NIZK. $\mathcal{N}_{\sigma_{\text{BND}}}$ is a Gaussian distribution over $\mathbb{Z}$ with standard deviation σ_{BND}, and $B_{\text{BND}} = \sqrt{2N}\sigma_{\text{BND}}$. See the full version of the paper for more details.

2.3.3 Exact Amortized Proofs of Shortness. As can be seen from $\mathcal{R}_{\text{BND}}$ the non-exact amortized proof has the disadvantage of introducing a "slack" factor $B_{\text{BND}} = \sqrt{2N}\sigma_{\text{BND}}$, meaning that the proven bound is substantially larger than what an honest party would generate. This ultimately leads to larger parameters for any application that uses Π_{BND}, as one always has to assume that dishonestly provided encryptions or commitments only fulfills the larger bound.

We will therefore also use a tighter ZK amortized proof of shortness which shows $\mathcal{R}_{\text{BND}}$ for the ℓ_∞-norm and with B_{BND} being 1. The disadvantage of this proof, over Π_{BND}, is that it does not scale as well with the number of statements that are proven as Π_{BND}.

For our exact amortized proof, we use a version of the protocol from Bootle et al. [10]. They give an efficient amortized sublinear zero-knowledge protocol for proving the knowledge of short vectors $\boldsymbol{s}_i$ and $\boldsymbol{e}_i$ over $\mathbb{Z}_q$ satisfying $\boldsymbol{A}\boldsymbol{s}_i + \boldsymbol{e}_i = \boldsymbol{t}_i$. We adapt their techniques for the case where $\boldsymbol{e}_i$ is zero and always prove that $\|\boldsymbol{s}_i\|_\infty \leq 1$. Our amortized protocol will be denoted throughout this work as $(\Pi_{\text{SMALL}}, \Pi_{\text{SMALLV}})$. These modifications are non-trivial and require us to re-prove that the construction is a ZK proof. We present more details in Section 4.

2.4 Verifiably Shuffling Ciphertexts

We construct a shuffle of BGV ciphertexts $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_\tau$ as follows:

(1) The server creates encryptions $\boldsymbol{c}'_1, \ldots, \boldsymbol{c}'_\tau$ of 0 and commits to each $\boldsymbol{c}'_i$ as $\text{Com}(\boldsymbol{c}'_i)$. Then, by homomorphically adding $\boldsymbol{c}_i$ to $\text{Com}(\boldsymbol{c}'_i)$ we obtain commitments $\text{Com}(\hat{\boldsymbol{c}}_i)$ to the same plaintexts as in $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_\tau$, with "fresh" randomness.
(2) The shuffle server reveals the openings $\hat{\boldsymbol{c}}_i$, but in random order. It then runs the verifiable shuffle protocol of [4] to prove that these openings are indeed the correct openings of the commitments.

To make the full construction verifiable, we use additional zero-knowledge proofs: the shuffle server will have to show that the $\text{Com}(\hat{\boldsymbol{c}}_i)$ are valid BDLOP commitments with bounded noise and contain well-formed encryptions of 0 (i.e. have small noise as well). For this, we use the ZK proofs introduced in the previous subsection. But this is insufficient because the protocol of [4] only supports BDLOP commitments of single elements from R_q, while BGV ciphertexts consist of two elements from R_q. We, therefore, extend the shuffle protocol by Aranha et al. to verifiably shuffle vectors in $R_q^{l_c}$. The full construction is described in Section 5.

2.5 Verifiable Decryption

We verifiably decrypt the BGV ciphertexts containing the votes in the voting scheme. In order to avoid a single party that has the secret decryption key (and could decrypt the inputs into the mix-net) we secret-share the key among multiple decryption servers.

The decryption algorithm introduced in Section 2.1 is only passively secure, but we assume that attackers may act maliciously in the voting scheme. We, therefore, modify the passively secure decryption protocol as follows:

- During key generation, a BDLOP commitment to each share is generated and published. The opening information is given to the shareholder.
- Each decryption share will additionally contain a proof that the decryption share is well-formed; the decryption algorithm proves that the decryption share is generated using the committed key share and that the randomness used is bounded. We will use the ZK proofs in Section 2.3.

We fully describe these transformations and prove them secure in Section 6. We do, however, not implement the (verifiable) key generation for our construction, which can e.g. be obtained by modifying a threshold key generation protocol such as [34].

3 THE VOTING SCHEME

The high-level architecture for the counting phase of our protocol is shown in Figure 1. As it follows a standard design [37], we do not describe its security properties further here, but refer the reader to the full version of the paper for a more formal treatment. We also have left out some aspects, such as voter authentication, to focus on the core building blocks of our construction.

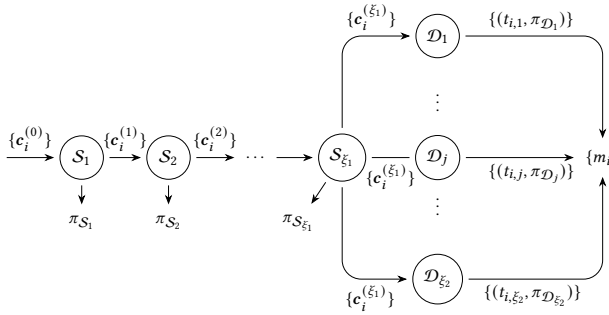

Figure 1: The high-level counting phase of our voting protocol. Each shuffle server $\mathcal{S}_k$ receives a set of ciphertexts $\{\boldsymbol{c}_i^{(k-1)}\}$, shuffles them, and outputs a new set of ciphertexts $\{\boldsymbol{c}_i^{(k)}\}$ and a proof $\pi_{\mathcal{S}_k}$. When all shuffle proofs are verified, each decryption server $\mathcal{D}_j$ partially decrypts every ciphertext and outputs the partial decryptions $\{t_{i,j}\}$ and a proof of correctness $\pi_{\mathcal{D}_j}$. Votes $\{m_i\}$ are reconstructed from the partial decryptions.

The voting protocol requires a *trusted set of players* to run the setup, a set of *voters* Voter_i and their *computers* Comp_i, a *ballot box* Ballot, a collection of *shuffle servers* $\mathcal{S}_k$, a collection of *decryption servers* $\mathcal{D}_j$ and one or more *auditors* Audit. We will assume that there are ξ_1 shuffle servers and ξ_2 decryption servers in total. The voting protocol consists of a *setup phase*, a *registration phase*, a *casting phase*, a *counting phase* as well as a verification algorithm to check casting and counting.

Setup Phase. A trusted set of players runs the key generation algorithm KGen of the PKE scheme with Distributed Decryption. The key generation can either be done in a trusted fashion or distributed using the protocol by Rotaru et al. [34]. The derived public parameters pk are given to every participant, while the decryption key shares sk_j are given to the decryption servers $\mathcal{D}_j$.

A key pair $(\mathsf{sk}_B, \mathsf{vk}_B)$ for a EUF-CMA-secure signature scheme is also generated and given to the ballot box. The verification key vk_B is given to every participant.

Casting phase. Each voter Voter_i instructs its computer Comp_i which ballot to cast. The computer encrypts the ballot under the public key pk and creates a *ballot proof*, sending the encrypted ballot and proof to the box Ballot. The ballot proof is tied to the voter's identity and is supposed to stop copy-and-paste attacks against privacy. In the security proof, the ballot proof must allow us to extract ballots from adversarially generated encryptions. Either we can use an argument of knowledge, but to simplify the security proof we often encrypt the ballot under two distinct keys and use an argument of equality. The ballot box will check the proof and signs the encrypted ballot and the proof using sk_B. This signature σ_i is sent to the voter's computer. The computer verifies the signature σ_i from Ballot using vk_B and only accepts if it is valid. It then shows the voter the encrypted ballot, proof, and signature, which constitutes the voter's *receipt*. The voter Voter_i accepts the ballot as cast if and only if the computer accepts it with a receipt.

Counting phase. The ballot box Ballot sends the encrypted ballots and ballot proofs that it has seen to the auditor Audit as well as every decryption server $\mathcal{D}_j$. Ballot then sorts the list of encrypted ballots $\{\boldsymbol{c}_i^{(0)}\}$ and sends this to the first shuffle server $\mathcal{S}_1$ and every decryption server. If some voter has cast more than one ballot, only the encrypted ballot seen last is included in this list.

The ξ_1 shuffle servers $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_{\xi_1}$ consecutively use the shuffle algorithm on the input encrypted ballots $\{\boldsymbol{c}_i^{(k-1)}\}$, passing the shuffled and re-encrypted ballots $\{\boldsymbol{c}_i^{(k)}\}$ to the next shuffle server. They also pass the shuffled re-encrypted ballots $\{\boldsymbol{c}_i^{(k)}\}$ and the shuffle proof $\pi_{\mathcal{S}_k}$ to Audit and every decryption server.

Each decryption server verifies that the data from Ballot as well as each shuffle server is consistent (input-output wise), and that every shuffle proof $\pi_{\mathcal{S}_k}$ verifies for the respective ciphertexts. Only then will they run the distributed decryption algorithm DDec with their decryption key share sk_j and send their partial decryption shares $t_{i,j}$ of each ballot $\boldsymbol{c}_i^{(\xi_1)}$ to the auditor as well as each recipient of the output. To obtain the result, each recipient can then run Comb on the partial decryption shares $t_{i,j}$.

Verification. The auditor verifies the data from Ballot (it checks that the ballot proofs of knowledge verify), that the encrypted ballots received by the first shuffle are consistent with the data from Ballot, that every shuffle proof verifies, and then runs the combining algorithm Comb on the received partial decryption shares $t_{i,j}$ from each $\mathcal{D}_j$. If all checks pass then the auditor accepts, otherwise, it rejects. Finally, Audit outputs the list of messages, including public key material, as its transcript.

One can easily design a verification algorithm that takes as input a transcript, a result, and optionally a receipt, and either accepts or rejects. The verification algorithm simply runs the auditor with the public key material and the messages listed in the transcript and checks if the auditor's result matches the input. If a receipt is present, it also verifies the signature σ_i using the ballot box' verification key vk_B, checks that the encrypted ballot and ballot proof are present in the ballot box data set, and that the encrypted ballots are present in the first shuffle server's input.

Note that there are many variations of this protocol. It can be used with so-called return codes, which allow human verification of the vote cast and detect a cheating computer Comp_i of the voter.

Many comparable schemes are phrased in terms of an ideal bulletin board, where every player posts their messages. Implementing a bulletin board is tricky in practice, so instead, we have described the scheme as a conventional cryptographic protocol passing messages via a network.

It is worth noting that for our concrete scheme, anyone can redo the auditor's work (since no secret key material is involved) by running the verification algorithm (and parts of the code algorithm) on the public data, making the voting protocol (universally) verifiable.

4 EXACT AMORTIZED ZK PROOFS

Bootle et al. [10] give an efficient amortized sublinear zero-knowledge protocol for proving the knowledge of short vectors $\boldsymbol{s}_i$ and $\boldsymbol{e}_i$ over $\mathbb{Z}_q$ satisfying $\boldsymbol{A}\boldsymbol{s}_i + \boldsymbol{e}_i = \boldsymbol{t}_i$. For our setting, we adapt their techniques for the case where $\boldsymbol{e}_i$ is zero, and prove that $\|\boldsymbol{s}_i\|_\infty \leq 1$[2].

We explain the main idea of [10] for proving knowledge of a preimage $\boldsymbol{s}$ of $\boldsymbol{t} = \boldsymbol{A}\boldsymbol{s}$ and then generalize to an amortized proof for τ elements with sublinear communication.

The approach follows an ideal linear commitments-technique with vector commitments $\mathsf{Com}_{\mathsf{L}}(\cdot)$ over $\mathbb{Z}_q$. The prover initially commits to the vector $\boldsymbol{s}$ as well as an auxiliary vector $\boldsymbol{s}_0$ of equal length. Implicitly, this defines a vector of polynomials $\boldsymbol{f}(X) = \boldsymbol{s}_0(X) + \boldsymbol{s}$ for the prover. Now consider the vector of polynomials $\boldsymbol{f}(X) \circ (\boldsymbol{f}(X) - \boldsymbol{1}) \circ (\boldsymbol{f}(X) + \boldsymbol{1})$, where $\circ$ denote the coordinate-wise product, then the coefficients of X^0 are exactly $\boldsymbol{s} \circ (\boldsymbol{s} - \boldsymbol{1}) \circ (\boldsymbol{s} + \boldsymbol{1})$ and therefore $\boldsymbol{0}$ if and only if the aforementioned bound on $\boldsymbol{s}$ holds. In that case, each aforementioned polynomial in $\boldsymbol{f}(X) \circ (\boldsymbol{f}(X) - \boldsymbol{1}) \circ (\boldsymbol{f}(X) + \boldsymbol{1})$ is divisible by X. Therefore, the prover computes the coefficient vectors

$$1/X \cdot \boldsymbol{f}(X) \circ (\boldsymbol{f}(X) - \boldsymbol{1}) \circ (\boldsymbol{f}(X) + \boldsymbol{1}) = \boldsymbol{v}_2 X^2 + \boldsymbol{v}_1 X + \boldsymbol{v}_0$$

and commits to these. Additionally, define the value $\boldsymbol{d} = \boldsymbol{t} - \boldsymbol{A}\boldsymbol{f} = -\boldsymbol{A}\boldsymbol{s}_0$, which the prover also commits to.

[2]The authors of [10] mention that this optimization is possible, but neither present the modified protocol nor a proof.

The verifier now sends a challenge x, for which the prover responds with $\overline{f} = f(x)$. The prover also uses the linear property of the commitment scheme to show that:

(1) $\mathsf{Com}_{\mathsf{L}}(s_0) \cdot x + \mathsf{Com}_{\mathsf{L}}(s)$ opens to $\overline{f}$.
(2) $\mathsf{Com}_{\mathsf{L}}(v_2) \cdot x^2 + \mathsf{Com}_{\mathsf{L}}(v_1) \cdot x + \mathsf{Com}_{\mathsf{L}}(v_0)$ opens to the value $\frac{1}{x} \cdot \overline{f} \circ (\overline{f} + 1) \circ (\overline{f} - 1)$.

The prover additionally opens the commitment to $\boldsymbol{d}$ and the verifier checks that it opens to $\frac{1}{x} \cdot (\boldsymbol{t} - \boldsymbol{A}\overline{\boldsymbol{f}})$. Here, the first two commitment openings allow us to deduce that the correct $\overline{f}$ is sent by the prover and that the values committed as $\boldsymbol{s}$ are indeed commitments to $\{-1, 0, 1\}$. Then, from opening $\boldsymbol{d}$ we get that the committed $\boldsymbol{s}$ is the preimage of $\boldsymbol{t}$ under $\boldsymbol{A}$.

The ideal linear commitments in [10] get realized using an Encode-then-Hash scheme. In this commitment scheme, the prover commits to vectors $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n \in \mathbb{Z}_q^{l_{\text{msg}}}$:

(1) Sample n random vectors $\boldsymbol{r}_1, \ldots, \boldsymbol{r}_n \in \mathbb{Z}_q^{\eta}$
(2) Let Encode be the encoding function of an $[l, l_{\text{msg}}+\eta, d]$ Reed-Solomon Code with code-length l, message length $l_{\text{msg}} + \eta$ and minimal distance d. Compute $\boldsymbol{e}_i \leftarrow \mathsf{Encode}(\boldsymbol{x}_i \| \boldsymbol{r}_i)$ for each $i \in [n]$.
(3) Construct matrix $\boldsymbol{E} = \mathsf{RowsToMatrix}(\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n)$ where $\boldsymbol{e}_i$ is row i.
(4) Commit to each *column* of $\boldsymbol{E}$ using a hash, then compress all commitments to Merkle root M.
(5) Send M to the verifier.

For the prover to show to the verifier that $\boldsymbol{x}$ is an opening of the linear combination $\sum_{i=1}^{n} \gamma_i \boldsymbol{x}_i$:

(1) It computes $\boldsymbol{r} = \sum_{i=1}^{n} \gamma_i \boldsymbol{r}_i$ and sends $\boldsymbol{r}$ to the verifier.
(2) The verifier chooses a subset I of size η from $[l]$.
(3) The prover opens the commitment for each column $i \in I$ of $\boldsymbol{E}$ and proves that it lies in the Merkle tree M by revealing the path.
(4) The verifier checks that $\mathsf{Encode}(\boldsymbol{x} \| \boldsymbol{r})$ coincides at position i with the respective linear combination of all n opened values in column i of $\boldsymbol{E}$.

This is a proof of the respective statement due to the random choice of the set I. Intuitively, if each row of $\boldsymbol{E}$ is in the code[3], but they do not sum up to $\boldsymbol{x}$, then the linear combination of the codewords in $\boldsymbol{E}$ must differ from $\mathsf{Encode}(\boldsymbol{x} \| \boldsymbol{r})$ in at least d positions, which is the minimum distance of the code. By the random choice of I and by setting η appropriately, the verifier would notice such a disagreeing entry with high probability. At the same time, because only η columns of $\boldsymbol{E}$ are opened, this leaks no information about the vectors $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$ if the evaluation points of the output of Encode are different from those of the input, i.e. if the code is not systematic.

For the case of more than one secret, the prover wants to show that $\boldsymbol{t}_i = \boldsymbol{A}\boldsymbol{s}_i$ for τ values $\boldsymbol{t}_i$ known to the verifier, subject to $\boldsymbol{s}_i$ again being ternary vectors. The goal is to establish the latter for all $\boldsymbol{t}_i$ simultaneously while verifying only one equation and sending only one vector $\overline{\boldsymbol{f}}$. Then the prover commits to $\boldsymbol{s}_i$ as well as an additional blinding value $\boldsymbol{s}_0$. Let $a_1, \ldots, a_\tau \in \mathbb{Z}_q$ be distinct interpolation points and define the ith Lagrange polynomial

$$\ell_i(X) = \prod_{i \neq j} \frac{X - a_j}{a_i - a_j}.$$

Additionally, let $\ell_0(X) = \prod_{i=1}^{\tau}(X - a_i)$. Then every $f \in \mathbb{Z}_q[X]/\ell_0(X)$ can be written uniquely as $f(X) = \sum_{i=1}^{\tau} \lambda_i \ell_i(X)$ and any $g \in \mathbb{Z}_q[X]/\ell_0(X)^b$ as a linear combination of $\{\ell_i(X)\ell_0(X)^j\}_{j=0}^{b-1}$. Define the polynomial

$$\boldsymbol{f}(X) = \sum_{i=0}^{\tau} \boldsymbol{s}_i \ell_i(X),$$

and observe that $\boldsymbol{f}(X) \circ (\boldsymbol{f}(X) - \boldsymbol{1}) \circ (\boldsymbol{f}(X) + \boldsymbol{1})$ is divisible by $\ell_0(X)$ if and only if all $\ell_i(X)$-coefficients of $\boldsymbol{f}(X)$ for $i \in [\tau]$ are 0. Additionally, since $\ell_i(X) \cdot \ell_j(X) = 0 \bmod \ell_0(X)$ if $i, j \in [n], i \neq j$ this then also implies that the $\boldsymbol{s}_i$ are ternary. Moreover, we only have to commit to additional $3 \cdot \tau$ coefficients of $\{\ell_i(X)\ell_0(X)^j\}_{j=0}^{b-1}$ to prove well-formedness of any evaluation of $\boldsymbol{f}(X)$ sent by the prover.

The protocol is described in detail in the full version of the paper.

As our construction substantially deviates from that of [10] we show that the protocol indeed is a ZKPoK. In the full version of the paper we show that the following holds:

Theorem 4. *The amortized zero-knowledge proof of exact openings is* complete *when the secrets $\boldsymbol{s}_i$ has ternary coefficients, it is* special sound *if the $\mathsf{SKS}^2_{r,v,1}$ problem is hard and the hash-function is collision-resistant, and it is statistically* honest-verifier zero-knowledge.

Towards defining the size of the proof, we see that the proof size is dominated by the sending of the openings of the homomorphic commitments (step 9 in Figure 4 in the full version) and the opening of the column-wise commitments of $\boldsymbol{E}$ via Merkle tree paths (step 11). More concretely:

- In step 9, prover sends polynomials which are openings to the homomorphic commitments of total size $3vN \log_2 q$ and additional randomness of total size $3\eta \log_2 q$.
- In step 11, the preimages of the hash column commitments $(\boldsymbol{E}|_I)$ have length $(3\tau+2)\eta \log_2 q$ while the Merkle tree paths add another $2\kappa\eta(1 + \log_2 l)$ bits.

This leads to a proof of size

$$(3vN + (3\tau + 2)\eta) \log_2 q + 2\kappa\eta(1 + \log_2 l) \text{ bits} \tag{1}$$

in total. The second part is essentially independent of τ, which decides how good the proof amortizes after fixing the lattice components. We get the optimal result by setting $3vN \approx (3\tau + 2)\eta$.

5 VERIFIABLE SHUFFLE OF CIPHERTEXTS

The recent work by Aranha et al. [4] presents an efficient protocol Π_{SHUF} for a shuffle of openings of the lattice-based commitments from Section 2.2 using proofs of linear relations. The protocol of [4] only supports committed secrets coming from R_q. We now extend their protocol to verifiably shuffle vectors in $R_q^{l_c}$.

[3] For the proof to work, the verifier additionally has to verify this claim or rather, that all rows are close to actual codewords. One mechanism to achieve this is to commit to an additional auxiliary row and also open a random linear combination of all rows, including the auxiliary row.

5.1 The Extended Shuffle for Commitments

To prove a shuffle, both the prover and verifier are given a list of commitments $[\![\boldsymbol{m}_1]\!], \ldots, [\![\boldsymbol{m}_\tau]\!]$ as well as potential messages $(\hat{\boldsymbol{m}}_1, \ldots, \hat{\boldsymbol{m}}_\tau)$ from $R_q^{l_c}$. The prover additionally obtains openings $\boldsymbol{m}_i, \boldsymbol{r}_i, f_i$ and wants to prove that the set of plaintext elements is the same set as the underlying elements of the commitments for some secret permutation π of the indices in the lists. More formally, our goal is to prove the following relation

$$\mathcal{R}_{\Pi_{\text{SHUF}}^{l_c}} = \left\{ (x, w) \;\middle|\; \begin{array}{l} x = ([\![\boldsymbol{m}_1]\!], \ldots, [\![\boldsymbol{m}_\tau]\!], \hat{\boldsymbol{m}}_1, \ldots, \hat{\boldsymbol{m}}_\tau), \\ w = (\pi, f_1, \ldots, f_\tau, \boldsymbol{r}_1, \ldots, \boldsymbol{r}_\tau), \pi \in S_\tau, \\ \forall i \in [\tau] : \; f_i \cdot [\![\boldsymbol{m}_{\pi^{-1}(i)}]\!] = f_i \cdot \begin{bmatrix} \boldsymbol{c}_{1,\pi^{-1}(i)} \\ \boldsymbol{c}_{2,\pi^{-1}(i)} \end{bmatrix} \\ = \boldsymbol{A}_{\text{C}} \boldsymbol{r}_i + f_i \cdot \begin{bmatrix} \boldsymbol{0} \\ \hat{\boldsymbol{m}}_i \end{bmatrix} \wedge \;\; ||\boldsymbol{r}_i[j]|| \leq 4\sigma_C \sqrt{N} \end{array} \right\}.$$

Towards proving this relation, we observe that it is sufficient to let the verifier choose a random element $h \xleftarrow{\$} R_q$. Then instead of proving a shuffle on $\boldsymbol{m}_1, \ldots, \boldsymbol{m}_\tau$, the prover instead performs the same proof on $\langle \boldsymbol{m}_1, \rho \rangle, \ldots, \langle \boldsymbol{m}_\tau, \rho \rangle$ where $\rho = (1, h, \ldots, h^{l_c - 1})^\top$. The problem with this approach is that we must also be able to apply ρ to the commitments $[\![\boldsymbol{m}_1]\!], \ldots, [\![\boldsymbol{m}_\tau]\!]$, without re-committing to the inner product and proving correctness in zero-knowledge.

Since each commitment $[\![\boldsymbol{m}]\!]$ can be written as

$$\begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix} = \boldsymbol{A}_{\text{C}} \boldsymbol{r} + \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{m} \end{bmatrix}$$

we can write $\boldsymbol{c}_1 = \boldsymbol{A}_{\text{C},1} \boldsymbol{r}$ and $\boldsymbol{c}_2 = \boldsymbol{A}_{\text{C},2} \boldsymbol{r} + \boldsymbol{m}$. From this we can create a new commitment $[\![\langle \boldsymbol{\rho}, \boldsymbol{m} \rangle]\!]$ under the new commitment key $\text{pk}' = (\boldsymbol{A}_{\text{C},1}, \rho \boldsymbol{A}_{\text{C},2})$ where $\boldsymbol{c}'_1 = \boldsymbol{c}_1$ remains the same, while we set $c'_2 = \langle \rho, \boldsymbol{c}_2 \rangle$. This does not increase the bound of the randomness of the commitment. Since

$$\boldsymbol{A}_{\text{C},2} = \begin{bmatrix} \boldsymbol{0}^{l_c \times n} & \boldsymbol{I}_{l_c} & \widehat{\boldsymbol{A}}_2 \end{bmatrix} \text{ where } \widehat{\boldsymbol{A}}_2 \in R_q^{l \times (k-n-l_c)},$$

it holds that

$$\boldsymbol{a}'_2 = \boldsymbol{\rho} \boldsymbol{A}_{\text{C},2} = \begin{bmatrix} \boldsymbol{0}^n & \boldsymbol{\rho}^\top & \boldsymbol{\rho} \widehat{\boldsymbol{A}}_2 \end{bmatrix}.$$

It is easy to see that breaking the binding property for pk' is no easier than breaking the binding property for pk.

Proposition 1. *If there exists an efficient attacker* Adv *that breaks the binding property on commitments under the key* pk' *with probability* ϵ*, then there exists an efficient algorithm* Adv' *that breaks the binding property on* pk *with the same probability.*

We can now construct the protocol $\Pi_{\text{SHUF}}^{l_c}$:

(1) Initially, prover $\mathcal{P}$ and verifier $\mathcal{V}$ hold $\{[\![\boldsymbol{m}_i]\!], \hat{\boldsymbol{m}}_i\}_{i \in [\tau]}$ for a public key $\text{pk} = (\boldsymbol{A}_{\text{C},1}, \boldsymbol{A}_{\text{C},2})$ while the prover additionally hold secrets $\{\boldsymbol{m}_i, \boldsymbol{r}_i\}_{i \in [\tau]}, \pi \in S_\tau$.
(2) $\mathcal{V}$ chooses $h \xleftarrow{\$} R_q$ and sends it to $\mathcal{P}$. Both parties compute $\boldsymbol{\rho} \leftarrow (1, h, \ldots, h^{l_c - 1})^\top$.
(3) $\mathcal{P}$ and $\mathcal{V}$ for each $[\![\boldsymbol{m}_i]\!] = (\boldsymbol{c}_{1,i}, \boldsymbol{c}_{2,i})$ compute $[\![\langle \boldsymbol{\rho}, \boldsymbol{m}_i \rangle]\!] = (\boldsymbol{c}_{1,i}, \langle \boldsymbol{\rho}, \boldsymbol{c}_{2,i} \rangle) = (\boldsymbol{c}_{1,i}, \boldsymbol{c}'_{2,i})$.
(4) $\mathcal{P}$ and $\mathcal{V}$ run Π_{SHUF} on input commitments $\{[\![\langle \boldsymbol{\rho}, \boldsymbol{m}_i \rangle]\!]\}_{i \in [\tau]}$ and messages $\langle \boldsymbol{\rho}, \hat{\boldsymbol{m}}_i \rangle$. $\mathcal{P}$ uses same permutation π, randomness $\boldsymbol{r}_i$ as before. The commitment key $\text{pk}' = (\boldsymbol{A}'_{\text{C},1}, \boldsymbol{A}'_{\text{C},2})$ is used by both.
(5) If Π_{SHUF} accepts then $\mathcal{V}$ accepts in $\Pi_{\text{SHUF}}^{l_c}$, otherwise rejects.

We show the following in the full version of the paper:

Lemma 1. *Assuming that the commitment scheme is binding and that* $\tau \leq \frac{1}{2 - 2\exp(-N/q)}$*. Then the protocol* $\Pi_{SHUF}^{l_c}$ *is a HVZK PoK for the relation* $\mathcal{R}_{\Pi_{SHUF}^{l_c}}$ *with soundness error* $\kappa_\Gamma = 2/q^N + 2\tau/|C|^\tau + 4(\tau+1) \cdot (\tau \cdot (l_c - 1) + 1)/q^{2N}$*, where* C *is the challenge space of the proof of linearity employed in* Π_{SHUF}*.*

5.2 Verifiable Shuffle of BGV Ciphertexts

We now implement the verifiable shuffle for ciphertexts that we outlined in Section 2.4. To recap quickly, the idea behind the shuffle of BGV ciphertexts $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_\tau$ is as follows:

(1) The shuffle server creates encryptions $\boldsymbol{c}'_1, \ldots, \boldsymbol{c}'_\tau$ of 0 and commits to each $\boldsymbol{c}'_i$ as $\text{Com}\boldsymbol{c}'_i$. Then, by homomorphically adding $\boldsymbol{c}_i$ to $\text{Com}\boldsymbol{c}'_i$ we obtain commitments $\text{Com}\hat{\boldsymbol{c}}_i$ to the same plaintexts as in $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_\tau$, with "fresh" randomness.
(2) The shuffle server reveals the openings $\hat{\boldsymbol{c}}_i$, but in random order. It then runs the verifiable shuffle protocol from the previous subsection to prove that these openings are indeed the commitments' correct (permuted) openings.

In the following, we describe the resulting approach in more detail.

Public parameters. Let $p \ll q$ be primes, let R_q and R_p be defined as above for a fixed N, and let $B_{\text{Key}}, B_{\text{Err}} \in \mathbb{N}$ be bounds for an instance of our chosen PKE scheme. We assume properly generated keys and ciphertexts according to the KeyGen and Enc algorithms in Section 2.1.

The shuffle server $\mathcal{S}$ takes as input a set of τ publicly known BGV ciphertexts $\{\boldsymbol{c}_i\}_{i=1}^\tau$, where the total noise in each ciphertext is bounded by B_{Dec}, i.e. each ciphertexts fulfills $P_{\text{sk}}(\cdot)$.

Randomizing. First, $\mathcal{S}$ randomizes all the received ciphertexts. Towards this it creates a new set of ciphertexts $\{\boldsymbol{c}'_i\}_{i=1}^\tau$:

$$\boldsymbol{c}'_i = (u'_i, v'_i) = (ar'_i + pe'_{i,1}, br'_i + pe'_{i,2}),$$

where $r'_i \xleftarrow{\$} S_{B_{\text{Key}}}$ and $e'_{i,1}, e'_{i,2} \xleftarrow{\$} S_{B_{\text{Err}}}$ as in fresh ciphertexts. This corresponds to creating fresh, independent encryptions of 0. $\mathcal{S}$ will *not* publish these $\boldsymbol{c}'_i$.

Committing. $\mathcal{S}$ now commits to the $\boldsymbol{c}'_i$. Towards this, we re-write the commitment matrix from Section 2.2 for $l_c = 2$ and add the public key of the encryption scheme to get a $(n+2) \times (k+3)$ matrix $\boldsymbol{A}_{\text{M}}$, where $\boldsymbol{0}^n$ are row-vectors of length n, $\boldsymbol{a}_{1,1}, \boldsymbol{a}_{1,2}$ are column vectors of length n, $\boldsymbol{a}_{2,3}, \boldsymbol{a}_{3,3}$ are row vectors of length $k-n-2$ and $\boldsymbol{A}_{1,3}$ is of size $n \times (k-n-2)$. Then,

$$\text{Com}(u'_i, v'_i) = [\![(ar'_i + pe'_{i,1}, br'_i + pe'_{i,2})]\!] = \boldsymbol{A}_{\text{M}} \boldsymbol{r}'_i$$

$$= \begin{bmatrix} \boldsymbol{I}_n & \boldsymbol{a}_{1,1} & \boldsymbol{a}_{1,2} & \boldsymbol{A}_{1,3} & 0 & 0 & 0 \\ \boldsymbol{0}^n & 1 & 0 & \boldsymbol{a}_{2,3} & a & p & 0 \\ \boldsymbol{0}^n & 0 & 1 & \boldsymbol{a}_{3,3} & b & 0 & p \end{bmatrix} \begin{bmatrix} \boldsymbol{r}_i \\ r'_i \\ e'_{i,1} \\ e'_{i,2} \end{bmatrix},$$

where $\boldsymbol{r}_i \in R_q^k$ is the randomness used in the commitment. Further, let $[\![(u_i, v_i)]\!]_{\boldsymbol{0}}$ be the trivial commitment to (u_i, v_i) with no randomness. Then, given the commitment $[\![(u'_i, v'_i)]\!]$ and $[\![(u_i, v_i)]\!]_{\boldsymbol{0}}$ we can be compute a commitment

$$[\![(\hat{u}_i, \hat{v}_i)]\!] = [\![(u_i, v_i)]\!]_{\boldsymbol{0}} + [\![(u'_i, v'_i)]\!].$$

Thus, the commitments $[\![(\hat{u}_i, \hat{v}_i)]\!]$ contain re-randomized encryptions of the original ciphertexts. $\mathcal{S}$ can therefore open a permutation of the $(\hat{u}_i, \hat{v}_i)$ and prove correctness of the shuffled opening using algorithm $\Pi^{l_c}_{\text{SHUF}}$. To ensure correctness, we have to additionally show that each u'_i, v'_i was created so that decryption is correct, i.e., it has small enough noise.

Proving correctness of commitments. Let $\boldsymbol{A}_{\mathsf{M}}$ be the $(n+2)\times(k+3)$ matrix defined above. Then $\mathcal{S}$ needs to prove that, for all i, it knows secret short vectors $\boldsymbol{r}'_i$ of length $k+3$ that are solutions to the following equations:

$$\boldsymbol{t}_i = \boldsymbol{A}_{\mathsf{M}}\boldsymbol{r}'_i = [\![(ar'_i + pe'_{i,1}, br'_i + pe'_{i,2})]\!], \left\|\boldsymbol{r}'_i\right\|_\infty \leq B_\infty.$$

To show this, $\mathcal{S}$ runs the Π_{SMALL}-protocol on $\boldsymbol{A}_{\mathsf{M}}, \{\boldsymbol{r}'_i\}_{i=1}^{\tau}, \{\boldsymbol{t}_i\}_{i=1}^{\tau}$. $\mathcal{S}$ uses Fiat-Shamir to ensure the non-interactivity of the proof.

The full protocol. We summarize this protocol as Π_{MIX}:

(1) $\mathcal{S}$ obtains ciphertexts $\{\boldsymbol{c}_i\}_{i\in[\tau]} = \{(u_i, v_i)\}_{i\in[\tau]}$.
(2) $\mathcal{S}$ for each $i \in [\tau]$ samples $r'_i, e'_{i,1}, e'_{i,2}$ as above. It then creates commitments $\{[\![u'_i, v'_i]\!] = [\![ar'_i + pe'_{i,1}, br'_i + pe'_{i,2}]\!]\}_{i\in[\tau]}$ using randomness $\boldsymbol{r}_i$ for each such commitment.
(3) Let $\boldsymbol{t}_i = [\![(u'_i, v'_i)]\!]$ and $\boldsymbol{r}'_i = [\boldsymbol{r}_i^\top, r'_i, e_{i,1}, e_{i,2}]^\top$. Then $\mathcal{S}$ computes $\pi_{\text{SMALL}} \leftarrow \Pi_{\text{SMALL}}$ for matrix $\boldsymbol{A}_{\mathsf{M}}$, input vectors $\{\boldsymbol{r}'_i\}$, target vectors $\{\boldsymbol{t}_i\}$ and bound B_∞.
(4) Let $\hat{\boldsymbol{c}}_i = (u_i + u'_i, v_i + v'_i)$ and L be a random permutation of $\{\hat{\boldsymbol{c}}_i\}_{i\in[\tau]}$. Then $\mathcal{S}$ computes $\pi_{\text{SHUF}} \leftarrow \Pi^{l_c}_{\text{SHUF}}$ with input commitments $\{[\![(\hat{u}_i, \hat{v}_i)]\!]\}_{i\in[\tau]}$, commitment messages $\{\hat{\boldsymbol{c}}_i\}_{i\in[\tau]}$, randomness $\{\boldsymbol{r}_i\}_{i\in[\tau]}$ and ciphertexts L.
(5) $\mathcal{S}$ outputs $(\{\boldsymbol{t}_i\}_{i\in[\tau]}, \pi_{\text{SMALL}}, L, \pi_{\text{SHUF}})$.

Given such a string $(\{\boldsymbol{t}_i\}_{i\in[\tau]}, \pi_{\text{SMALL}}, L, \pi_{\text{SHUF}})$ from $\mathcal{S}$ as well as ciphertext vector $\{\boldsymbol{c}_i\}_{i\in[\tau]}$ any third party $\mathcal{V}$ can now run the following algorithm Π_{MIXV} to verify the mix:

(1) Run the verification algorithm of Π_{SMALLV} for π_{SMALL} on inputs $\boldsymbol{A}_{\mathsf{M}}, \{\boldsymbol{t}_i\}_{i\in[\tau]}$ and B_∞. If verification fails: output 0.
(2) For $\forall i \in [\tau]$ set $[\![\hat{\boldsymbol{c}}_i]\!] = [\![\boldsymbol{c}_i]\!]_{\boldsymbol{0}} + \boldsymbol{t}_i$.
(3) Run the verification algorithm of $\Pi^{l_c}_{\text{SHUFV}}$ for π_{SHUF} on input $\{[\![\hat{\boldsymbol{c}}_i]\!]\}_{i\in[\tau]}, L$. If the verification fails, then output 0. Otherwise, output 1.

In the following, define noise bound B_{Mix} to be the maximum level of noise in ciphertexts $\boldsymbol{c}'_i$, i.e. the maximal noise of the randomness $r'_i, e'_{i,1}, e'_{i,2}$ used to create the ciphertexts.

We want that the outputs of the mixing protocol fulfill the following relation $\mathcal{R}_{\text{MIX}}$:

$$\left\{ (x, w) \;\middle|\; \begin{array}{l} x = (\boldsymbol{c}_1, \ldots, \boldsymbol{c}_\tau, \hat{\boldsymbol{c}}_1, \ldots, \hat{\boldsymbol{c}}_\tau, [\![\boldsymbol{c}'_1]\!], \ldots, [\![\boldsymbol{c}'_\tau]\!]), \\ w = (\pi, \boldsymbol{r}'_1, \ldots, \boldsymbol{r}'_\tau,), \pi \in S_\tau, \forall i \in [\tau] : \\ [\![\boldsymbol{c}'_i]\!] = \boldsymbol{A}_{\mathsf{M}}\boldsymbol{r}'_i, \left\|\boldsymbol{r}'_i\right\|_\infty \leq B_{\mathsf{Mix}}, \hat{\boldsymbol{c}}_{\pi(i)} = \boldsymbol{c}_i + \boldsymbol{c}'_i \end{array} \right\}.$$

If the noise-levels in all $\boldsymbol{c}_i$ and $\boldsymbol{c}'_i$ are bounded by B_{Dec} and B_{Mix} respectively, and $(B_{\mathsf{Dec}} + B_{\mathsf{Mix}}) < \lfloor q/2 \rfloor$, then all $\boldsymbol{c}_i$ and $\hat{\boldsymbol{c}}_{\pi(i)}$ will, for some permutation π, decrypt to the same message m_i under sk. In the full version of the paper we analyze the guarantees of Π_{MIX} in more detail.

5.3 Communication of a BGV Shuffle

The mixing phase transcript contains τ new ciphertexts generated by the server, which are of size $2\tau N \log_2 q$ bits.

The server must provide a proof of shuffle and an amortized proof of shortness for $r'_i, e'_{i,1}, e'_{i,2}$. Both proofs prove a relation about commitments to the randomization factors u'_i, v'_i added to the old ciphertexts to get the new ciphertexts. Each commitment to u'_i, v'_i is of size $(n+2)N \log_2 q$ bits. We denote the proof by π_{SMALL}.

The shuffle proof consists of τ commitments of size $(n+1)N \log_2 q$ bits, τ R_q-elements of size $N \log_2 q$ bits and a proof of linearity per ciphertext. This adds up to an overall size of $((n+2)N \log_2 q + 2(k-n)N \log_2(6\sigma_{\mathrm{C}}))\tau$ bits for the proof of shuffle, and in total

$$((2n+6)N \log_2 q + 2(k-n)N \log_2(6\sigma_{\mathrm{C}}))\tau + |\pi_{\text{SMALL}}| \text{ bits.}$$

6 VERIFIABLE DISTRIBUTED DECRYPTION

In this section, we provide a construction for a PKE scheme with distributed decryption which is secure against active attacks. We combine the distributed decryption protocol from Section 2.1 with zero-knowledge proofs to achieve this. In a nutshell, the DDec algorithm that we introduced in Section 2.1 first requires that each decryption server chooses a uniformly random E_j from a bounded distribution. Next, it outputs a linear combination t_j involving a ciphertext element u, the decryption key share s_j, and E_j. To make this actively secure, we will let the key generation algorithm output a commitment to s_j. Then, to show that it computed t_j correctly from u, the decryption server will reveal a commitment to E_j as well as two zero-knowledge proofs: i) it will show that E_j is bounded as required, and ii) it will show that t_j is indeed computed using a linear combination.

6.1 The Actively Secure Protocol

Let the ring R_q, the statistical security parameter sec, and bounds $B_{\mathsf{Err}}, B_{\mathsf{Com}}, B_{\mathsf{Dec}}$ be public information, together with the plaintext modulus p for the PKE scheme. Let $\boldsymbol{A}_{\mathsf{C}}$ be the public commitment matrix of a BDLOP instance for message size $l_c = 1$.

- $\mathsf{KGen}_A(1^\kappa, \xi_2)$:
  (1) Get $(\mathsf{pk}, \mathsf{sk}, s_1, \ldots, s_{\xi_2}) \leftarrow \mathsf{KGen}(1^\kappa, \xi_2)$ as in the passive distributed encryption protocol.
  (2) $\forall j \in [\xi_2]$ compute $([\![s_j]\!], \boldsymbol{d}_j) \leftarrow \mathsf{Com}(s_j)$.
  (3) Output $\mathsf{pk}_A = (\mathsf{pk}, [\![s_1]\!], \ldots, [\![s_{\xi_2}]\!])$ and finally $\mathsf{sk}_A = \mathsf{sk}$ and $\mathsf{sk}_{A,j} = (s_j, \boldsymbol{d}_j)$ for all $j \in [\xi_2]$.
- Enc_A and Dec_A work just like the original Enc and Dec in the passively secure threshold encryption scheme, ignoring additional information in pk_A.
- $\mathsf{DDec}(\mathsf{sk}_{A,j}, \{c_i\}_{i\in[\tau]})$ where $c_i = (u_i, v_i)$:
  (1) For each $i \in [\tau]$ compute $m_{i,j} = s_j u_i$, sample uniform noise $E_{i,j} \leftarrow R_q$ such that $\left\|E_{i,j}\right\|_\infty \leq 2^{\mathsf{sec}}(B_{\mathsf{Dec}}/p\xi_2)$ and compute the decryption share $t_{i,j} = m_{i,j} + pE_{i,j}$.
  (2) For each $i \in [\tau]$ compute $([\![E_{i,j}]\!], \boldsymbol{r}''_{i,j}) \leftarrow \mathsf{Com}(E_{i,j})$ and use the Π_{LIN}-protocol to compute a proof for the linear relation $t_{i,j} = s_j u_i + pE_{i,j}$ from

$$\pi_{L_{i,j}} \leftarrow \Pi_{\text{LIN}}(((s_j, \boldsymbol{r}_j), (E_{i,j}, \boldsymbol{r}''_{i,j})); \\ ([\![s_j]\!], [\![E_{i,j}]\!], t_{i,j}), (u_i, p)).$$

(3) Each commitment $[\![E_{i,j}]\!]$ is of the form

$$\begin{bmatrix} \boldsymbol{I}_n & \boldsymbol{a}_{1,1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{0}^n & 1 & \boldsymbol{a}_{2,2} \end{bmatrix} \cdot \boldsymbol{r}''_{i,j} + \begin{bmatrix} 0 \\ E_{i,j} \end{bmatrix} = \underbrace{\begin{bmatrix} \boldsymbol{I}_n & \boldsymbol{a}_{1,1} & \boldsymbol{A}_{1,2} & 0 \\ \boldsymbol{0}^n & 1 & \boldsymbol{a}_{2,2} & 1 \end{bmatrix}}_{\boldsymbol{A}_{\mathrm{D}}} \begin{bmatrix} \boldsymbol{r}''_{i,j} \\ E_{i,j} \end{bmatrix},$$

where $\left\|\boldsymbol{r}''_{i,j}\right\|_\infty \leq B_{\mathsf{Com}}$ is the randomness used in the commitments. Run the zero-knowledge protocol Π_{BND} on $(\{(E_{i,j}, \boldsymbol{r}''_{i,j}\}_{i\in[\tau]}); (\boldsymbol{A}_{\mathrm{D}}, \{[\![E_{i,j}]\!]\}_{i\in[\tau]}))$ to obtain the amortized zero-knowledge PoK π_{BND_j}.

(4) Output $\mathsf{ds}_j = (\{t_{i,j}\}_{i=1}^{\tau}, \pi_{\mathcal{D}_j})$ with the decryption proof $\pi_{\mathcal{D}_j} = (\{[\![E_{i,j}]\!]\}_{i=1}^{\tau}, \{\pi_{L_{i,j}}\}_{i=1}^{\tau}, \pi_{\mathrm{BND}_j})$.

- $\mathsf{Comb}_A(\{c_i\}_{i=1}^{\tau}, \{\mathsf{ds}_j\}_{j\in[\xi_2]})$:

(1) Parse ds_j as $(\{t_{i,j}\}_{i=1}^{\tau}, \pi_{\mathcal{D}_j})$.
(2) Verify the proofs $\pi_{L_{i,j}}$.
(3) Verify the proofs π_{BND_j}.
(4) If any verification protocol returned 0 then output $\perp$. Otherwise, compute

$$m_i = (v_i - t_i \mod q) \mod p, \text{ where}$$
$$t_i = t_{i,1} + \cdots + t_{i,\xi_2} \text{ for } i = 1, \ldots, \tau,$$

and output the set of messages $m_1, \ldots, m_\tau$.

The randomness $\boldsymbol{r}''_{i,j}$ has much smaller ℓ_∞ norm than $E_{i,j}$, and hence, we will run the Π_{BND} protocol with small standard deviation σ_{BND} for rows 1 to k, while row $k+1$ will have large $\hat{\sigma}_{\mathrm{BND}}$.This trivially works for Π_{BND} as all operations, also in the extractor for the soundness-proof, are coordinate-wise.

The following theorems refer to definitions of threshold correctness, threshold verifiability, and distributed decryption simulatability given in Section 2.1. In the following theorems, let the noise bounds B_{Dec} and $\hat{B}_{\mathrm{BND}}$ satisfy $(1 + B_{\mathrm{Dec}}) \cdot 2^{\mathsf{sec}} < 2\hat{B}_{\mathrm{BND}} < \lfloor q/2 \rfloor$.

Theorem 5. *Let ciphertext-noise be bounded by B_{Dec}, and let the noise added in* DDec *be bounded by $2^{\mathsf{sec}} B_{\mathrm{Dec}}$. Suppose the passively secure protocol is threshold correct and the protocols Π_{LIN} and Π_{BND} are complete. Then the actively secure protocol is threshold correct.*

Informally, since $B_{\mathrm{Dec}} + 2^{\mathsf{sec}} B_{\mathrm{Dec}} < q/2$, it follows that decryption is correct. Furthermore, since $(1 + B_{\mathrm{Dec}}) \cdot 2^{\mathsf{sec}} < 2\hat{B}_{\mathrm{BND}} < q/2$ and Π_{LIN} and Π_{BND} are complete, the arguments will be accepted, which means that the decryption proof will be accepted.

Theorem 6. *Let Adv_0 be an adversary against threshold verifiability for the actively secure protocol 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 $\epsilon_0 \leq \epsilon_1 + \epsilon_2$. The runtime of Adv_1 and Adv_2 are essentially the same as of Adv_0.*

We sketch the argument. We only consider ciphertexts with noise bounded by B_{Dec}, so we may assume that the noise in any particular ciphertext is bounded by B_{Dec}.

If the decryption is incorrect for a particular ciphertext, then for some j no relation $t_{i,j} = s_j u_i + p E_{i,j}$ holds for an $E_{i,j}$ of the norm at most $2\hat{B}_{\mathrm{BND}}$. This can happen in two ways: Either the argument for the linear combination of the commitments to $E_{i,j}$ and s_j is incorrect or the bound on $E_{i,j}$ is incorrect. In the former case, we trivially get an adversary Adv_1 against soundness for Π_{LIN}. Similar for the case of Π_{BND}.

Theorem 7. *Suppose the passively secure protocol is simulatable and Π_{LIN} and Π_{BND} are honest-verifier zero-knowledge. Then there exists a simulator for the actively secure protocol such that for any distinguisher Adv_0 for this simulator with advantage ϵ_0, there exists an adversary Adv_4 against hiding for the commitment scheme[4], with advantage ϵ_4, and distinguishers Adv_1, Adv_2 and Adv_3 for the simulators for the passively secure protocol, Π_{LIN} and Π_{BND}, respectively, with advantages $\epsilon_1, \epsilon_2, \epsilon_3$, such that $\epsilon_0 \leq \epsilon_1 + \epsilon_2 + \epsilon_3 + \epsilon_4$. The runtime of Adv_1, Adv_2, Adv_3 and Adv_4 are essentially the same as of Adv_0.*

We sketch the argument. Using appropriate simulators, the simulator simulates the arguments and the passively secure distributed decryption algorithm. It replaces the commitment to the noise $E_{i,j}$ by commitments to zero.

The claim about the simulator follows from a straightforward hybrid argument. We begin with distributed decryption.

First, we replace the Π_{LIN} arguments with simulated arguments, which gives us a distinguisher Adv_2 for the Π_{LIN} honest verifier simulator. Second, we replace the Π_{BND} arguments with simulated arguments, which gives us a distinguisher Adv_3 for the Π_{BND} honest verifier simulator. Third, we replace the commitments to the noise $E_{i,j}$ with random commitments, giving us an adversary Adv_4 against hiding for the commitment scheme. Fourth, we replace the passively secure distributed decryption algorithm with its simulator, which gives us a distinguisher Adv_1 for the simulator.

After four changes, we are left with the claimed simulator for the actively secure protocol, and the claim follows.

6.2 Communication Complexity of DistDec

Each partial decryption consists of one element from R_q, namely the output of DDec, which means that the output from the passively secure protocol is of size $\xi_2 \tau N \log_2 q$ bits.

Each decryption server outputs a commitment $[\![E_{i,j}]\!]$ to the added noise and proof of linearity per ciphertext, and an amortized proof of shortness for all the added noise values. Each server has a public commitment of their decryption key-share to be used in the proof of linearity, but we neglect this as it is constant.

Each commitment $[\![\cdot]\!]$ is of size $(n+1)N \log_2 q$ bits, and each proof of linearity is of size $(k-n)N(\log_2(6\sigma_{\mathrm{C}}) + \log_2(6\hat{\sigma}_{\mathrm{C}})$ bits because the partial decryption is given in the clear and one commitment is re-used in all equations. Finally, each of the amortized proofs is of size $k\hat{n}N \log_2(6\sigma_{\mathrm{BND}}) + \hat{n} \log_2(6\hat{\sigma}_{\mathrm{BND}})$ bits because of the different norms of the secret values as noted earlier. As the bounds in the amortized proof depend on the number of commitments in the statement, each amortized proof is for a batch of N equations at once to control the growth of parameters.

The total size of the distributed decryption is

$$\xi_2((n+2)N \log_2 q + (k-n)N(\log_2(6\sigma_{\mathrm{C}}) + \log_2(6\hat{\sigma}_{\mathrm{C}})) + k\hat{n} \log_2(6\sigma_{\mathrm{BND}}) + \hat{n} \log_2(6\hat{\sigma}_{\mathrm{BND}}))\tau \text{ bits.}$$

7 PERFORMANCE

We provide an overview of parameters and descriptions in Table 1.

[4] A more careful argument could allow us to dispense with this adversary. We have opted for a simpler argument since the commitment scheme is also used elsewhere.

Parameter	Explanation	Constraints
κ	Computational security parameter	At least 128 bits
sec	Statistical security parameter	40 bits
N	Degree of polynomial X^N+1 in R_p, R_q	N a power of two
p	Plaintext modulus	p a small prime
q	Ciphertext and commitment modulus	Prime $q = 1 \mod 2N$ s.t. $\max\|v - su\| \ll q/2$
k	Portion of homomorphic commitment vector dedicated to binding	
n	Length of commitment vector	
C	Challenge space for Linear ZK proofs of commitments	$C = \{c \in R_p \mid \|c\|_\infty = 1, \|c\|_1 = \nu\}$
ν	Maximum ℓ_1-norm of elements in C	
S_B	Set of elements of ∞-norm at most B	$S_B = \{x \in R_p \mid \|x\|_\infty \leq B\}$
B_{Com}	Bound on the commitment noise	–
B_{Key}	Bound for secret key in encryption scheme	Chosen as 1
B_{Err}	Bound for noise in ciphertexts	Chosen as 1
σ_{C}	Standard deviation in linear ZK proofs for one-time commitments	Chosen to be $\sigma_{\text{C}} = 0.954 \cdot \nu \cdot B_{\text{Com}} \cdot \sqrt{kN}$
$\hat{\sigma}_{\text{C}}$	Standard deviation in linear ZK proofs for reusable commitments	Chosen to be $\hat{\sigma}_{\text{C}} = 22 \cdot \nu \cdot B_{\text{Com}} \cdot \sqrt{kN}$
σ_{BND}	Standard deviation for the one-time amortized proof in mixing	Chosen to hide the commitment randomness $\boldsymbol{r}''_{i,j}$
$\hat{\sigma}_{\text{BND}}$	Standard deviation for the one-time amortized proof in mixing	Chosen to hide the decryption noise $E_{i,j}$
$\hat{n}$	Dimension of proof in Π_{BND}	$\hat{n} \geq \kappa + 2$
ξ_1, ξ_2	Number of shuffle- and decryption-servers	
τ	Total number of messages/number of voters	For soundness we need $(\tau^\delta + 1)/\|R_q\| < 2^{-\kappa}$
l	Encoding length in Π_{SMALL}	–
l_c	Length of the committed message in Π_{SMALL}	–
η	Randomness of encodings in Π_{SMALL}	–
g	Dimension of Reed-Solomon Code in Π_{SMALL}	–

Table 1: System parameters and constraints.

7.1 Concrete Parameters and Total Size

We begin by fixing the rejection-sampling parameter as $M = 3$, leading to a general abort probability of $2/3$ for each rejection sampling proof. This allows us to define the standard deviations in all proof instances.

We pick the noise in the BGV ciphertexts as well as in commitments to come from ternary distributions, as this gives tight control on the noise growth during the protocols.

To be able to choose concrete parameters for the mix-net, we need to estimate how much noise is added to the ciphertexts through the two stages of the protocol: 1) the shuffle phase, and 2) the decryption phase. This follows from a standard analysis that incorporates the slacks of the ZK proofs involved in the protocols and will be one lower bound on choosing q as the noise should not wrap around computations mod q.

For our example, we let the number of shuffle and decryption servers be $\xi_1 = \xi_2 = 4$. We fix the plaintext modulus to be $p = 2$, statistical security parameter sec $= 40$ (a common choice in the MPC literature), and need $N = 4096$ when q is chosen as outlined above in order for the underlying lattice problems to be hard, see details in Table 5. This allows for votes of size 4096 bits, which is a feasible size for real-world elections representing a wide range of voter options.

Finally, we must decide on parameters for the exact proof of shortness from Section 4. The soundness of the protocol depends on the ratio between the number of equations and the size of the modulus. We choose to compute the proof in batches of size N instead of computing the proof for all τ commitments at once and will have to run each proof twice to achieve negligible soundness error. After choosing appropriate parameters for code length and the number of tested rows η, the total size of π_{SMALL}, by instantiating equation 1, is $\approx 20\tau$ KB.

We summarize the concrete sizes of each part of the protocol in Table 2. Each voter submits a ciphertext size of approximately 80 KB. The size of the mix-net, including ciphertexts, commitments, shuffle proof, and proof of shortness, is approximately 370τ KB per mixing node $\mathcal{S}_k$. The size of the decryption phase, including partial decryptions, commitments, proofs of linearity, and proofs of boundedness, is approximately 157τ KB per decryption node $\mathcal{D}_j$.

$\boldsymbol{c}_i^{(k)}$	$[\![R_q^{l_c}]\!]$	π_{SHUF}	$\pi_{L_{i,j}}$	π_{SMALL}	π_{BND}	$\pi_{\mathcal{S}_i}$	$\pi_{\mathcal{D}_j}$
80 KB	$40(l_c+1)$ KB	150τ KB	35 KB	20τ KB	2τ KB	370τ KB	157τ KB

Table 2: Size of the ciphertexts, commitments, and proofs.

See Appendix A for more details on the choice of parameters.

7.2 Implementation

We developed a proof-of-concept implementation to compare with previous results in the literature. Our performance figures were collected on an Intel Kaby Lake Core i7-7700 CPU machine with 64GB of RAM running single-threaded at 3.6GHz, with Turbo Boost disabled to reduce measurement variability. The results can be found in Tables 3 and 4. Our research prototype is publicly available at https://github.com/dfaranha/lattice-verifiable-mixnet.

First, we compare the performance of the main building blocks with an implementation of the shuffle-proof protocol proposed in [4]. That work used the FLINT library to implement arithmetic involving polynomials of degree $N = 1024$ with 32-bit coefficients, fitting a single machine word. Their parameters were not compatible with the fast Number Theoretic Transform (NTT) for polynomial multiplication, so a CRT decomposition to two half-degree polynomials was used instead. The code was made available, so a direct comparison is possible.

In this work, the degree is much larger ($N = 4096$) and coefficients are multi-word ($q \approx 2^{78}$), but the parameters are compatible with the NTT. We implemented polynomial arithmetic with the efficient NFLlib [2] library using the RNS representation for coefficients and accelerated with AVX2 instructions. We observed that our polynomial multiplication is around 19 times faster than [4] (61,314 cycles instead of 1,165,997), despite parameters being considerably larger. We also employed the FLINT library for arithmetic routines not supported in NFLlib, such as polynomial inversion, but that incurred some non-trivial costs to convert representations between two libraries. We adapted [40] and [39] for Gaussian sampling and adjusting the standard deviation σ accordingly. We employ BLAKE3 [32] for fast hashing inside the various proofs.

Computing a commitment takes 0.45 ms on the target machine, which is 2x faster than [4]. Opening a commitment is slower due to conversions between libraries for performing the norm-test. Our implementation of BGV encryption at 0.74 ms is much faster than the 69 ms reported for verifiable encryption in [4], while decryption is improved by a factor of 10. Distributed decryption with passive security costs an additional 1.56 ms per party, but the zero-knowledge proofs for active security increase the cost. The shuffle proof performance is 44.9 ms per vote, thus slower than the 27 ms reported in [4] due to slower rejection sampling.

For the other sub-protocols, we benchmarked executions with $\tau = 1000$ and report the execution time amortized per vote for both prover and verifier in Table 4. In the case of Π_{SMALL}, we implement only the performance-critical polynomial arithmetic and encoding scheme, since this is already representative of the overall performance. From the table, we can compute the cost of distributed

Primitive	Commit	Open	Encrypt	Decrypt	DistDec
Time	0.45 ms	2.76 ms	0.74 ms	0.64 ms	1.56 ms

Table 3: Timings for cryptographic operations. Numbers were obtained by computing the average of 10^4 executions measured using the cycle counter available on the platform.

Protocol	$\Pi_{\text{LIN}} + \Pi_{\text{LINV}}$	$\Pi^{l_c}_{\text{SHUF}} + \Pi^{l_c}_{\text{SHUFV}}$
Time	$(43.4 + 6.4)\tau$ ms	$(44.9 + 7.9)\tau$ ms
Protocol	$\Pi_{\text{BND}} + \Pi_{\text{BNDV}}$	$\Pi_{\text{SMALL}} + \Pi_{\text{SMALLV}}$
Time	$(92.7 + 23.9)\tau$ ms	$(112.3 + 5.0)\tau$ ms

Table 4: Timings for cryptographic protocols, obtained by computing the average of 100 executions with $\tau = 1000$.

decryption Π_{DEC} with active security as $(1.56 + 0.45 + 43.4 + 92.7) = 138.11$ ms per vote, the cost of verification Π_{DECV} as $(6.4 + 23.9) = 30.2$ ms per vote, the cost of Π_{MIX} as $(0.74+0.45+112.3+44.9) = 158.4$ ms and the cost of Π_{MIXV} as $(5.0 + 7.9) = 12.9$ ms per vote. This result compares quite favorably with the costs of 1.54 s and 1.51 s per vote to respectively generate/verify a proof in the lattice-based shuffle proof of [22] in a Kaby Lake processor running at a similar frequency. Our total numbers are around 17 times faster after adjusting for clock frequency, while storage overhead is much lower.

8 CONCLUDING REMARKS

We have proposed a verifiable secret shuffle of BGV ciphertexts and a verifiable distributed decryption protocol. Together, these two novel constructions are practical and solve a long-standing problem in the design of quantum-safe cryptographic voting systems.

Verifiable secret shuffles for discrete logarithm-based cryptography has seen a long sequence of incremental designs follow Neff's breakthrough construction. While individual published improvements were often fairly small, the overall improvement in performance over time was significant. We expect that our designs can be improved in a similar fashion. In particular, we expect that the size of the proofs can be significantly reduced. While it is certainly straight-forward to download a few hundred gigabytes today (compare with high-quality video streaming), many voters will be discouraged and this limits the universality of verification in practice. It, therefore, seems reasonable to focus further effort on reducing the size of the proofs.

The distributed decryption protocol does not have an adjustable threshold. In practice, this is not much of a problem, since the keys will be shared among many key holders. Only when counting starts is the key material given to the decryption servers. Key reconstruction can then be combined with a key distribution protocol.

Shuffles followed by distributed decryption is one paradigm for the design of cryptographic voting systems. Another possible paradigm is to use key shifting in the shuffles. This would then allow us to use a single party for decryption (though it must still be verifiable, e.g., using the protocols [24, 36]). Key shifting can be done with many of the same techniques that we use for distributed decryption, but there seems to be difficulties in amortizing the proofs. This means that key shifting with just the techniques we use will be significantly slower and of increased size, as we would need additional proofs of linearity for each ciphertext in each shuffle.

Follow-up work by Høgåsen and Silde [27] shows how our voting protocol can be combined with a return-code mechanism to achieve individual voter verifiability against a cheating ballot box.

Finally, we note that our scheme and concrete instantiation using the NTT is optimized for speed, and that it is possible to slightly decrease the parameters by instantiating the encryption scheme based on the SKS^2 and DKS^∞ problems in higher dimensions k using a smaller, but still a power of 2, ring-dimension N. We leave this as future work. We also remark that lattice-based cryptography, and especially lattice-based zero-knowledge proofs such as the recent work by Lyubashevsky et al. [29], continuously improves the state-of-the-art, and we expect future works to improve the concrete efficiency of our protocol.

REFERENCES

[1] Ben Adida. 2008. Helios: Web-based Open-Audit Voting. In *USENIX Security 2008*, Paul C. van Oorschot (Ed.). USENIX Association, 335–348.

[2] Carlos Aguilar Melchor, Joris Barrier, Serge Guelton, Adrien Guinet, Marc-Olivier Killijian, and Tancrède Lepoint. 2016. NFLlib: NTT-Based Fast Lattice Library. In *CT-RSA 2016 (LNCS, Vol. 9610)*, Kazue Sako (Ed.). Springer, Heidelberg, 341–356. https://doi.org/10.1007/978-3-319-29485-8_20

[3] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. 2014. Quantum Attacks on Classical Proof Systems: The Hardness of Quantum Rewinding. In *55th FOCS*. IEEE Computer Society Press, 474–483. https://doi.org/10.1109/FOCS.2014.57

[4] Diego F. Aranha, Carsten Baum, Kristian Gjøsteen, Tjerand Silde, and Thor Tunge. 2021. Lattice-Based Proof of Shuffle and Applications to Electronic Voting. In *CT-RSA 2021 (LNCS, Vol. 12704)*, Kenneth G. Paterson (Ed.). Springer, Heidelberg, 227–251. https://doi.org/10.1007/978-3-030-75539-3_10

[5] Thomas Attema, Vadim Lyubashevsky, and Gregor Seiler. 2020. Practical Product Proofs for Lattice Commitments. In *CRYPTO 2020, Part II (LNCS, Vol. 12171)*, Daniele Micciancio and Thomas Ristenpart (Eds.). Springer, Heidelberg, 470–499. https://doi.org/10.1007/978-3-030-56880-1_17

[6] Carsten Baum, Jonathan Bootle, Andrea Cerulli, Rafaël del Pino, Jens Groth, and Vadim Lyubashevsky. 2018. Sub-linear Lattice-Based Zero-Knowledge Arguments for Arithmetic Circuits. In *CRYPTO 2018, Part II (LNCS, Vol. 10992)*, Hovav Shacham and Alexandra Boldyreva (Eds.). Springer, Heidelberg, 669–699. https://doi.org/10.1007/978-3-319-96881-0_23

[7] Carsten Baum, Ivan Damgård, Vadim Lyubashevsky, Sabine Oechsner, and Chris Peikert. 2018. More Efficient Commitments from Structured Lattice Assumptions. In *SCN 18 (LNCS, Vol. 11035)*, Dario Catalano and Roberto De Prisco (Eds.). Springer, Heidelberg, 368–385. https://doi.org/10.1007/978-3-319-98113-0_20

[8] Rikke Bendlin and Ivan Damgård. 2010. Threshold Decryption and Zero-Knowledge Proofs for Lattice-Based Cryptosystems. In *TCC 2010 (LNCS, Vol. 5978)*, Daniele Micciancio (Ed.). Springer, Heidelberg, 201–218. https://doi.org/10.1007/978-3-642-11799-2_13

[9] Manuel Blum. 1984. How to Exchange (Secret) Keys. *ACM Transactions on Computer Systems* 1 (1984), 175–193.

[10] Jonathan Bootle, Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. 2021. More Efficient Amortization of Exact Zero-Knowledge Proofs for LWE. In *ESORICS 2021, Part II (LNCS, Vol. 12973)*, Elisa Bertino, Haya Shulman, and Michael Waidner (Eds.). Springer, Heidelberg, 608–627. https://doi.org/10.1007/978-3-030-88428-4_30

[11] Xavier Boyen, Thomas Haines, and Johannes Müller. 2020. A Verifiable and Practical Lattice-Based Decryption Mix Net with External Auditing. In *ESORICS 2020, Part II (LNCS, Vol. 12309)*, Liqun Chen, Ninghui Li, Kaitai Liang, and Steve A. Schneider (Eds.). Springer, Heidelberg, 336–356. https://doi.org/10.1007/978-3-030-59013-0_17

[12] Xavier Boyen, Thomas Haines, and Johannes Müller. 2021. Epoque: Practical End-to-End Verifiable Post-Quantum-Secure E-Voting. In *IEEE European Symposium on Security and Privacy, EuroS&P 2021, Vienna, Austria, September 6-10, 2021*. IEEE, 272–291. https://doi.org/10.1109/EuroSP51992.2021.00027

[13] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. 2012. (Leveled) fully homomorphic encryption without bootstrapping. In *ITCS 2012*, Shafi Goldwasser (Ed.). ACM, 309–325. https://doi.org/10.1145/2090236.2090262

[14] David Chaum. 1981. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Commun. ACM* 24, 2 (1981), 84–88. https://doi.org/10.1145/358549.358563

[15] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. 2016. A Homomorphic LWE Based E-voting Scheme. In *Post-Quantum Cryptography - 7th International Workshop, PQCrypto 2016*, Tsuyoshi Takagi (Ed.). Springer, Heidelberg, 245–265. https://doi.org/10.1007/978-3-319-29360-8_16

[16] Núria Costa, Ramiro Martínez, and Paz Morillo. 2019. Lattice-Based Proof of a Shuffle. In *FC 2019 Workshops (LNCS, Vol. 11599)*, Andrea Bracciali, Jeremy Clark, Federico Pintore, Peter B. Rønne, and Massimiliano Sala (Eds.). Springer, Heidelberg, 330–346. https://doi.org/10.1007/978-3-030-43725-1_23

[17] Ivan Damgård, Marcel Keller, Enrique Larraia, Valerio Pastro, Peter Scholl, and Nigel P. Smart. 2013. Practical Covertly Secure MPC for Dishonest Majority - Or: Breaking the SPDZ Limits. In *ESORICS 2013 (LNCS, Vol. 8134)*, Jason Crampton, Sushil Jajodia, and Keith Mayes (Eds.). Springer, Heidelberg, 1–18. https://doi.org/10.1007/978-3-642-40203-6_1

[18] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. 2012. Multiparty Computation from Somewhat Homomorphic Encryption. In *CRYPTO 2012 (LNCS, Vol. 7417)*, Reihaneh Safavi-Naini and Ran Canetti (Eds.). Springer, Heidelberg, 643–662. https://doi.org/10.1007/978-3-642-32009-5_38

[19] Rafaël del Pino, Vadim Lyubashevsky, Gregory Neven, and Gregor Seiler. 2017. Practical Quantum-Safe Voting from Lattices. In *ACM CCS 2017*, Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM Press, 1565–1581. https://doi.org/10.1145/3133956.3134101

[20] Jelle Don, Serge Fehr, and Christian Majenz. 2020. The Measure-and-Reprogram Technique 2.0: Multi-round Fiat-Shamir and More. In *CRYPTO 2020, Part III (LNCS, Vol. 12172)*, Daniele Micciancio and Thomas Ristenpart (Eds.). Springer, Heidelberg, 602–631. https://doi.org/10.1007/978-3-030-56877-1_21

[21] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. 2022. Efficient NIZKs and Signatures from Commit-and-Open Protocols in the QROM. In *CRYPTO 2022, Part II (LNCS, Vol. 13508)*, Yevgeniy Dodis and Thomas Shrimpton (Eds.). Springer, Heidelberg, 729–757. https://doi.org/10.1007/978-3-031-15979-4_25

[22] Valeh Farzaliyev, Jan Willemson, and Jaan Kristjan Kaasik. 2023. Improved lattice-based mix-nets for electronic voting. *IET Inf. Secur.* 17, 1 (2023), 18–34.

[23] Amos Fiat and Adi Shamir. 1987. How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In *CRYPTO'86 (LNCS, Vol. 263)*, Andrew M. Odlyzko (Ed.). Springer, Heidelberg, 186–194. https://doi.org/10.1007/3-540-47721-7_12

[24] Kristian Gjøsteen, Thomas Haines, Johannes Müller, Peter Rønne, and Tjerand Silde. 2022. Verifiable Decryption In The Head. In *Information Security and Privacy: 27th Australasian Conference, ACISP 2022, Wollongong, NSW, Australia, November 28–30, 2022, Proceedings* (Wollongong, NSW, Australia). Springer-Verlag, Berlin, Heidelberg, 355–374. https://doi.org/10.1007/978-3-031-22301-3_18

[25] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. 1985. The Knowledge Complexity of Interactive Proof-Systems (Extended Abstract). In *17th ACM STOC*. ACM Press, 291–304. https://doi.org/10.1145/22145.22178

[26] Javier Herranz, Ramiro Martínez, and Manuel Sánchez. 2021. Shorter lattice-based zero-knowledge proofs for the correctness of a shuffle. In *Financial Cryptography and Data Security. FC 2021 International Workshops: CoDecFin, DeFi, VOTING, and WTSC, Virtual Event, March 5, 2021, Revised Selected Papers 25*. Springer, 315–329.

[27] Audhild Høgåsen and Tjerand Silde. 2022. Return Codes from Lattice Assumptions. *E-VOTE-ID* (2022). https://doi.org/10.15157/diss/025

[28] Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. 2020. *CRYSTALS-DILITHIUM*. Technical Report. National Institute of Standards and Technology. available at https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions.

[29] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. 2022. Lattice-Based Zero-Knowledge Proofs and Applications: Shorter, Simpler, and More General. In *CRYPTO 2022, Part II (LNCS, Vol. 13508)*, Yevgeniy Dodis and Thomas Shrimpton (Eds.). Springer, Heidelberg, 71–101. https://doi.org/10.1007/978-3-031-15979-4_3

[30] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. 2013. A Toolkit for Ring-LWE Cryptography. In *EUROCRYPT 2013 (LNCS, Vol. 7881)*, Thomas Johansson and Phong Q. Nguyen (Eds.). Springer, Heidelberg, 35–54. https://doi.org/10.1007/978-3-642-38348-9_3

[31] C. Andrew Neff. 2001. A Verifiable Secret Shuffle and Its Application to e-Voting. In *ACM CCS 2001*, Michael K. Reiter and Pierangela Samarati (Eds.). ACM Press, 116–125. https://doi.org/10.1145/501983.502000

[32] Jack O'Connor, Jean-Philippe Aumasson, Samuel Neves, and Zooko Wilcox-O'Hearn. 2020. BLAKE3: one function, fast everywhere. https://www.blake3.io. , 31 pages.

[33] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. 2020. *FALCON*. Technical Report. National Institute of Standards and Technology. available at https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions.

[34] Dragos Rotaru, Nigel P. Smart, Titouan Tanguy, Frederik Vercauteren, and Tim Wood. 2022. Actively Secure Setup for SPDZ. *J. Cryptol.* 35, 1 (jan 2022), 32 pages. https://doi.org/10.1007/s00145-021-09416-w

[35] Peter Schwabe, Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Gregor Seiler, and Damien Stehlé. 2020. *CRYSTALS-KYBER*. Technical Report. National Institute of Standards and Technology. available at https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions.

[36] Tjerand Silde. 2022. Short Paper: Verifiable Decryption for BGV. In *Financial Cryptography and Data Security. FC 2022 International Workshops - CoDecFin, DeFi, Voting, WTSC, Grenada, May 6, 2022, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 13412)*, Shin'ichiro Matsuo, Lewis Gudgeon, Ariah Klages-Mundt, Daniel Perez Hernandez, Sam Werner, Thomas Haines, Aleksander Essex, Andrea Bracciali, and Massimiliano Sala (Eds.). Springer, 381–390. https://doi.org/10.1007/978-3-031-32415-4_26

[37] Kristian Gjøsteen. 2022. *Practical Mathematical Cryptography*. CRC Press.

[38] Martin Strand. 2019. A Verifiable Shuffle for the GSW Cryptosystem. In *FC 2018 Workshops (LNCS, Vol. 10958)*, Aviv Zohar, Ittay Eyal, Vanessa Teague, Jeremy Clark, Andrea Bracciali, Federico Pintore, and Massimiliano Sala (Eds.). Springer, Heidelberg, 165–180. https://doi.org/10.1007/978-3-662-58820-8_12

[39] Shuo Sun, Yongbin Zhou, Yunfeng Ji, Rui Zhang, and Yang Tao. 2022. Generic, efficient and isochronous Gaussian sampling over the integers. *Cybersecur.* 5, 1 (2022), 10.

[40] Raymond K. Zhao, Ron Steinfeld, and Amin Sakzad. 2020. COSAC: COmpact and Scalable Arbitrary-Centered Discrete Gaussian Sampling over Integers. In *Post-Quantum Cryptography - 11th International Conference, PQCrypto 2020*, Jintai Ding and Jean-Pierre Tillich (Eds.). Springer, Heidelberg, 284–303. https://doi.org/10.1007/978-3-030-44223-1_16

A CHOOSING PARAMETERS CONCRETELY

We let the success probability of each of the zero-knowledge protocols be $1/M \approx 1/3$. We will use the following parameters, where we note that the commitments used in the shuffle and in the amortized proofs are only used once, while the proof of linearity in the decryption protocol depends on a commitment to the secret key share each time. However, that is the only part that is reused, and we can use a smaller standard deviation for the other commitment.

The proofs of linearity have two terms, and each of them must have a success probability of $1/\sqrt{3}$. This gives $\sigma_{\mathrm{C}} = 0.954\nu B_{\mathsf{Com}}\sqrt{kN}$. For the re-usable commitments we get $\hat{\sigma}_{\mathrm{C}} = 22\nu B_{\mathsf{Com}}\sqrt{kN}$. The amortized proof also has two checks, and we get a standard deviation $0.954\|\boldsymbol{S}'\boldsymbol{C}'\|_2$, where σ_{BND} and $\hat{\sigma}_{\mathrm{BND}}$ are depending on the norm of the elements in the rows of $\boldsymbol{S}'$.

We let the noise bounds $B_{\mathsf{Key}} = B_{\mathsf{Err}} = 1$ for the encryption.

To be able to choose concrete parameters for the mix-net, we need to estimate how much noise is added to the ciphertexts through the two stages of the protocol: 1) the shuffle phase, and 2) the decryption phase. Each part of the system contributes the following amount of noise to the ciphertexts:

- Fresh ciphertext: $B_{\mathrm{START}} = p\left\|er + e_{i,2} - e_{i,1}s\right\|_\infty + \|m\|_\infty$.
- Noise per shuffle: $B_{\mathrm{SHUF}} = p(\|er'\|_\infty + \left\|e'_{i,2}\right\|_\infty + \left\|-e'_{i,1}s\right\|_\infty)$.
- Noise in partial decryption: $B_{\mathrm{DDec}} = p\xi_2\left\|E'_{i,j}\right\|_\infty \leq 2^{\mathsf{sec}} B_{\mathrm{Dec}}$,

where $B_{\mathrm{Dec}} = B_{\mathrm{START}} + \xi_1 B_{\mathrm{SHUF}}$ is the upper bound of the noise added before the decryption phase. This means that we have the following bounds on each of the noise terms above, when using ternary noise:

$$\|e\|_1 \leq N, \quad \|r\|_\infty \leq 1, \quad \left\|e_{i,2}\right\|_\infty \leq 1, \quad \left\|e_{i,1}\right\|_1 \leq N,$$
$$\|s\|_\infty \leq 1, \quad \left\|r'\right\|_\infty \leq 1, \quad \left\|e'_{i,2}\right\|_\infty \leq 1, \quad \left\|e'_{i,1}\right\|_1 \leq N.$$

We get the following upper bounds:

$$B_{\mathrm{START}} = p(2N+1) + \lceil (p-1)/2 \rceil, \quad B_{\mathrm{SHUF}} = p(2N+1),$$

N	p	q	sec	ξ_1	ξ_2	n	k
4096	2	$\approx 2^{78}$	40	4	4	1	$l_c + 2$

ν	B_{Com}	$\hat{n}$	σ_{C}	$\hat{\sigma}_{\text{C}}$	σ_{BND}	$\hat{\sigma}_{\text{BND}}$	$\hat{B}_{\text{BND}}$
36	1	130	$\approx 2^{12}$	$\approx 2^{16.5}$	$\approx 2^{13.5}$	$\approx 2^{66}$	$\approx 2^{72.5}$

Table 5: Concrete parameters estimated for $\kappa \approx 168$ bits of DKS^{∞} security using the LWE-estimator and $\kappa \approx 262$ bits of SKS^2 security (by computing the Hermite root value to be 1.00225 from the dimension, modulus, and 2-norm of the secret vector).

which for ξ_1 shuffles gives us

$$B_{\text{Dec}} = (\xi_1 + 1)p(2N + 1) + \lceil (p-1)/2 \rceil.$$

Finally, we need to make sure that $B_{\text{Dec}} + B_{\text{DDec}} < q/2$, where $B_{\text{DDec}} = 2p\xi_2\hat{B}_{\text{BND}}$ because of the soundness slack of the amortized proof of bounded values. A honestly generated value $E_{i,j}$ is bounded by $2^{\text{sec}}(B_{\text{Dec}}/p\xi_2)$, but the proof can only guarantee that the values are shorter than some larger bound $2\hat{B}_{\text{BND}}$ (following Baum et al. [6, Lemma 3]) that depends on the number of equations in the statement. Define $S'_{1,k}$ to be the first k rows of S' and define S'_{k+1} to be the last row of S'. For batches of N equations, we then get that:

$$\begin{aligned} B_{\text{BND}} &\leq \sqrt{2N} \cdot \sigma_{\text{BND}} \leq \sqrt{2N} \cdot 0.954 \cdot \max \left\| S'_{1,k} C' \right\|_2 \\ &\leq 1.35 \cdot \sqrt{N} \cdot \max \left\| S'_{1,k} \right\|_1 \cdot \max \left\| C' \right\|_\infty \\ &\leq 1.35 \cdot k \cdot \sqrt{N} \cdot N \cdot B_{\text{Com}}, \end{aligned}$$

and, similarly,

$$\hat{B}_{\text{BND}} \leq \sqrt{2N} \cdot \hat{\sigma}_{\text{BND}} \leq 1.35 \cdot \sqrt{N} \cdot N \cdot \left\| E_{i,j} \right\|_\infty,$$

with B_{BND} for rows 1 to k of Z and $\hat{B}_{\text{BND}}$ for the last.

We fix plaintext modulus $p = 2$, statistical security parameter sec $= 40$, and need $N = 4096$ when q is large to provide proper security. This allows for votes of size 4096 bits, which should be a feasible size for real-world elections. We let the number of shuffle and decryption servers be $\xi_1 = 4$. It follows that $B_{\text{Dec}} < 2^{17}$ and $B_{\text{DDec}} < 2^{76.5}$. We then set $q \approx 2^{78}$, and verify that

$$\max_{i \in [\tau]} \|v_i - s u_i\| < 2 \cdot (2^{17} + 2^{76.5}) < q.$$

Finally, we must decide on parameters for the exact proof of shortness. The soundness of the protocol depends on the ratio between the number of equations and the size of the modulus. We choose to compute the proof in batches of size N instead of computing the proof for all τ commitments at once. Then we get $18N/(q-N) \approx 2^{-62}$, and hence, we must compute each proof twice in parallel to achieve negligible soundness. Furthermore, we choose $g \approx 2^{20}, l \approx 2^{20.3}, \eta = 325$ to keep the soundness $\approx 2^{-62}$. The total size of π_{SMALL}, by instantiating 1, is $\approx 20\tau$ KB.

We give a complete set of parameters in Table 5, and the concrete sizes of each part of the protocol in Table 2. Each voter submits a ciphertext size of approximately 80 KB. The size of the mix-net, including ciphertexts, commitments, shuffle proof, and proof of shortness, is approximately 370τ KB per mixing node $\mathcal{S}_i$. The size of the decryption phase, including partial decryptions, commitments, proofs of linearity, and proofs of boundedness, is approximately 157τ KB per decryption node $\mathcal{D}_j$.

B SECURITY IN THE QUANTUM RANDOM ORACLE MODEL

In this work, we have chosen parameters for all primitives such as to make our voting protocol secure against all known classical attacks. Since we only use assumptions that are assumed to be post-quantum secure, it is obvious to ask if our construction is also post-quantum secure. We cannot answer this within this work, due to the complexity of proving such a statement.

As a "second-best" approach, we can alternatively look at the post-quantum security of the individual building blocks. Here, of particular importance are the NIZKs that this work uses. We use two different types of proofs, namely those exploiting the homomorphism of an underlying OWF (such as $\Pi_{\text{LIN}}, \Pi_{\text{BND}}$) and those that rely only on commitments and a combinatorial argument (Π_{SMALL}). Both of these are made non-interactive in the ROM using the Fiat-Shamir transform, which becomes the QROM in the quantum setting. Here, the recent work of [21] could be used to show that Π_{SMALL} is online-extractable in the QROM and therefore still secure, for adjusted parameters.

Unfortunately, the situation is a bit more problematic for the homomorphism-based proofs. There, the most efficient QROM Fiat-Shamir approach that we are aware of is [20], which applies to Σ-protocols. Their work implies a large loss in parameters that they show to be inherent, and this loss grows with the number of rounds of the protocol. Even worse, new techniques would have to be developed to prove the security of Π_{BND} as it seems unlikely that [20] applies to it. To achieve provable security of all these NIZKs in the QROM, it would be better to replace the homomorphic OWF-based protocols with Commit-and-Open-based proofs following Π_{SMALL}. We expect that this would come at a significant cost in proof size as well as prover runtime, impacting the practicality of our construction.

A more optimistic view, which we share, is that known counterexamples in the QROM on NIZKs such as [3] are contrived and that there are no known attacks (beyond Grover's algorithm) for the NIZKs that we use. One could therefore argue that our construction is *plausibly post-quantum*. We leave a more detailed post-quantum security analysis, which also includes parameter choices to withstand attacks based on Grover's algorithm, for future work.