Error Correction and Ciphertext Quantization in Lattice Cryptography[*]

Daniele Micciancio and Mark Schultz

UC San Diego, {daniele,mdschultz}@eng.ucsd.edu

March 11, 2023

Abstract. Recent work in the design of rate $1 - o(1)$ lattice-based cryptosystems have used two distinct design paradigms, namely replacing the noise-tolerant encoding $m \mapsto (q/2)m$ present in many lattice-based cryptosystems with a more efficient encoding, and post-processing traditional lattice-based ciphertexts with a lossy compression algorithm, using a technique very similar to the technique of "vector quantization" within coding theory.

We introduce a framework for the design of lattice-based encryption that captures both of these paradigms, and prove information-theoretic rate bounds within this framework. These bounds separate the settings of trivial and non-trivial quantization, and show the impossibility of rate $1 - o(1)$ encryption using both trivial quantization and polynomial modulus. They furthermore put strong limits on the rate of constructions that utilize lattices built by tensoring a lattice of small dimension with $\mathbb{Z}^k$, which is ubiquitous in the literature. We additionally introduce a new cryptosystem, that matches the rate of the highest-rate currently known scheme, while encoding messages with a "gadget", which may be useful for constructions of Fully Homomorphic Encryption.

1 Introduction

Lattice-based cryptography has many advantages over traditional number-theoretic encryption, from conjectured security against quantum attacks, to the ability to perform arbitrary computations over encrypted data, while at the same time enjoying very fast (quasi-linear time) encryption and decryption operations. This is much better than the cubic running time of the modular exponentiation typically used in constructions based on number theory. However, there is one aspect for which lattice-based constructions have always lagged behind number-theoretic ones: key and ciphertext *sizes*. In fact, early proposals of encryption schemes based on lattices suffered from a very poor *rate*, meaning the ratio of the size of a plaintext to the size of a ciphertext was very small.

Improving the rate of encryption schemes is an important and well-studied problem, and a problem with a well-understood solution: hybrid encryption. By using public-key encryption on a fixed size, randomly chosen symmetric key, and then using this key to encrypt the actual message using a much more efficient block cipher, the cost of the public-key operation (both in terms of running time and rate) can be amortized over a large payload. However, by using hybrid encryption one loses one of the main attractions of lattice-based cryptography: the ability to compute on encrypted data, as data is now encrypted using

[*] Research supported in part by NSF Award 1936703. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

a block cipher with no useful homomorphic properties. Homomorphically decrypting AES or other “FHE-friendly” block ciphers [2,3], addresses this problem, but only partially: it allows one to move data from AES (or another symmetric encryption scheme) to lattice-based cryptography and then perform homomorphic computations on it. The reverse step, e.g. converting the FHE ciphertext back to a space-efficient symmetric ciphertext, is an open problem and would seem to require the symmetric cryptosystem to be fully homomorphic. This has motivated the study of lattice-based encryption schemes with better rate, leading to two constructions of lattice-based homomorphic encryption schemes with rate asymptotically close to 1 [18,6]. In this paper we present a unified study of high-rate lattice-based encryption schemes, presenting a general framework that parameterizes LWE-based (Learning With Error) encryption with two coding-theoretic objects we call *lattice codes*. The simplest lattice-based encryption scheme (originally proposed by Regev [32]), combines an LWE sample with simple scaling and rounding operations. Here, we replace these scalar operations with two arbitrary lattice codes, one used for error-correction (generalizing scaling), and one used for quantization (generalizing rounding). We then show that known constructions of rate $1 - o(1)$ encryption [18,6] can be described as instances of our general constructions for particular choices of lattice codes, and prove upper and lower bounds on the rate achievable in this framework. Analysis of these schemes in our framework highlights inefficiencies in many current constructions, which we fix to attain asymptotic (rate) improvements.

Organization The rest of this paper is organized as follows. In the rest of the introduction we provide more details on our technical contributions and related work. In Section 2 we present background information on error-correcting codes needed to describe and analyze our construction. In Section 3 we present our generalized encryption framework. In Section 4 we show how previous constructions can be obtained as special cases of our framework simply by properly choosing a pair of error correcting codes, and also present a construction combining the desirable properties of [18] and [6]. In Section 5 we present impossibility results that limit the rate achievable using common subcases of our generalized construction. In Section 6, we give concluding thoughts, and present some open problems.

1.1 Our Contributions

There is a well-known strategy for building (private-key) encryption from LWE, namely

- start with an LWE sample $(\mathbf{A}, \mathbf{b} := A\mathbf{s} + \mathbf{e})$, and
- add an encoding of the message $\mathsf{encode}(\mathbf{m})$ to the second component.

Provided one can later recover the message $\mathbf{m}$ from the noisy encoding $\mathsf{encode}(\mathbf{m}) + \mathbf{e}$, this suffices to build private-key encryption.

Given the ciphertext $(\mathbf{A}, \mathbf{b} := \mathbf{A}\mathbf{s}+\mathsf{encode}(\mathbf{m})+\mathbf{e})$, how might we compress it? The matrix $\mathbf{A}$ is itself uniformly random, and can be easily compressed using standard techniques[1].

[1] In theory, the same $\mathbf{A}$ can be reused with many different $\mathbf{s}_i$, making the amortized cost of $\mathbf{A}$ arbitrarily small. In practice, $\mathbf{A}$ is often replaced with a short seed that is deterministically expanded to $\mathbf{A}$. This is process is not fully justified theoretically, but it is easily proved secure in the random oracle model.

Therefore, we focus on compressing $\mathbf{b}$. This is pseudorandom under the LWE assumption, so we must appeal to some form of *lossy* compression. As the ciphertext already contains a form of error-correction, it can plausibly correct some additional noise.

We leverage a form of compression commonly known as *vector quantization*, where one maps a vector $\mathbf{v} \in \mathbb{R}^m$ to some discrete subset, say $\mathbb{Z}^m$, or more generally a lattice. We use this methodology to quantize $\mathbf{b}$ to a nearby lattice point $\lfloor \mathbf{b} \rceil_L \in L$, where $\lfloor \cdot \rceil_L : \mathbb{R}^m \to L$ is a generalized form of rounding, for example by solving the closest vector problem. Provided the sum of the quantization error $[\mathbf{b}]_Q := \mathbf{b} - \lfloor \mathbf{b} \rceil_Q$ and LWE error $\mathbf{e}$ can be corrected by the error-correcting code, our scheme will decrypt correctly, i.e. we will have successfully compressed an LWE ciphertext.

The above describes how our framework leverages two codes E, Q, for error-correction and quantization respectively. Concretely, the quantized LWE encryption scheme using E and Q (which we call $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$) encrypts by computing

$$\mathsf{Enc}_{\mathbf{s}}(\mathbf{m}) := (\mathbf{A}, \lfloor \mathbf{A}\mathbf{s} + \mathbf{e} + \mathsf{encode}_E(\mathbf{m}) \rceil_Q), \tag{1}$$

where $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, and $\mathbf{e} \leftarrow \chi_e^m$ for an error distribution χ_e. This is a mild modification of (standard) LWE-based encryption (see Definition 11 for details). Despite the simplicity of this approach, our framework is

- broad,
- modular, and
- necessary to achieve high rate.

We discuss all of these points next.

Breadth: Our framework includes all forms of error-correction and vector quantization that are expressible in terms of *lattice codes* (Definition 1), which are the reduction of a q-ary lattice L modulo q. Equivalently, they are discrete subgroups $L_q := (L \bmod q) \subseteq \mathbb{R}^m/q\mathbb{Z}^m$. For any such subgroup, there are (many) fundamental domains $\mathcal{V}_L$ such that $L_q + \mathcal{V}_L = \mathbb{R}^m/q\mathbb{Z}^m$ is a partition. A lattice code can be thought of as the choice of a pair $(L_q, \mathcal{V}_L)$, along with algorithms to efficiently decompose $\mathbb{R}^m/q\mathbb{Z}^m \to (L_q, \mathcal{V}_L)$. This includes most techniques of decoding a point $\mathbf{x} \in \mathbb{R}^m$ to $\lfloor x \rceil \in L$, say by solving the closest vector problem exactly, or approximately via techniques such as Babai's Nearest Planes [4].

In Section 4, we instantiate our framework with many different non-trivial LWE-based encryption schemes. In particular, we show that all existing rate $1 - o(1)$ encryption schemes [18,6] fit into our framework. Beside the schemes that we explicitly analyze, our framework additionally includes any scheme that encodes messages into a lattice for error correction (of which there are many [31,32,33,18]). All known cryptosystems which quantize ciphertexts are expressible in our framework, although this is a much shorter list (containing solely [6][2], and schemes which quantize via rounding each coordinate independently, which are common in practice [14,12]).

[2] We defer discussion of how one can realize this work in our framework to Section 4.3.

Moreover, we demonstrate the ease of working in our framework by "quantizing" several pre-existing cryptosystems. One such construction combines the desirable properties of [6,18], namely it encodes messages under a "gadget" lattice (similar to [18]), but attains the same (quasi-optimal) rate as [6].

Modular: Our framework separates the coding-theoretic analysis from the cryptographic analysis of encryption schemes. The cryptographic analysis of schemes in our framework is somewhat basic. We establish in Theorem 3 that $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,Q]$ is $\mathsf{RND\text{-}CPA}$-secure[3] (but potentially incorrect) for any choice of E, Q via a simple argument.

The coding-theoretic analysis is similarly straightforward. We express the rate of our cryptosystem in terms of a simple function of the LWE modulus q, dimension m, and volumes $\det E$ and $\det Q$ of the fundamental domains of E, Q.

Correctness analysis requires some knowledge about the shape of these fundamental domains, although we find that it is enough to know their packing and covering radii in the ℓ_2 and ℓ_∞ norms. This analysis frequently highlights inefficiencies in the choice E, Q of codes a cryptosystem (implicitly) uses. Most commonly, the quantizer Q can be replaced with a sparser quantizer Q' without (asymptotically) impacting the correctness of the cryptosystem. We make this modification in several cases, and often find asymptotic improvements. We summarize the results of our analysis in Table 1. Our optimizations tend to improve constructions from rate $1 - f(m)$ to $1 - \frac{f(m)}{\log_2 m}$, i.e. improve on known constructions by a logarithmic factor in the dimension. We discuss the reason for these small improvements shortly.

Name	E	Q	Rate	Gadget Quality of E	Source
Regev	$(q/p)\mathbb{Z}^m$	$\mathbb{Z}^m$	$1 - O(1)$	N/A	[32]
Quantized Regev	$(q/p)\mathbb{Z}^m$	$k\mathbb{Z}^m$	$1 - O\left(\frac{1}{\log_2 \frac{q}{k}}\right)$	N/A	Corollary 3
GH	$\Lambda_q^\perp(\mathbf{g}_p^t) \otimes \mathbb{Z}^{m/\ell}$	$\mathbb{Z}^m$	$1 - O(1)$	$O(q/p)$	[18]
Quantized GH	$\Lambda_q^\perp(\mathbf{g}_p^t) \otimes \mathbb{Z}^{m/\ell}$	$k\mathbb{Z}^m$	$1 - O\left(\frac{1}{\log_2 \frac{q}{k}}\right)$	$O(q/p)$	Section 4
BDGM	$(q/p)\mathbb{Z}^m$	$\Lambda_{q/p}(\mathbf{u}_m^t)$	$1 - O\left(\frac{\log_2(m\sigma)}{m\log_2 p}\right)$	N/A	[6]
Gadget	$\Lambda_q(\mathbf{g}_p^t) \otimes \mathbb{Z}^{m/\ell}$	$\mathbb{Z}^m$	$1 - O(1)$	$O(p)$	[28]
Quantized Gadget	$\Lambda_q(\mathbf{g}_p^t) \otimes \mathbb{Z}^{m/\ell}$	$\Lambda_{q/p}(\mathbf{u}_m^t)$	$1 - O\left(\frac{\log_2(m\sigma)}{m\log_2 p}\right)$	$O(p)$	Corollary 6

Table 1. The E, Q that we study the Quantized Encryption schemes $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,Q]$ in Section 4, where $\mathbf{g}_p^t = (1, p, p^2, \ldots, p^{\ell-1})$ for $\ell = \lceil \log_p q \rceil$, $\mathbf{u}_m^t = (1, 1, \ldots, 1)^t \in \mathbb{R}^m$, and k is a free parameter, typically set to some small polynomial in n. Note that the various parameters p, q, m, k may be required to satisfy certain divisibility constraints, see details in Section 4. The rates are computed assuming Gaussian parameter $\sigma = \Theta(\sqrt{n})$, secret key length $n = \Theta(m)$, ciphertext dimension m, and decryption failure rate $\delta = \exp(-n)$. The quality of a gadget (defined in [17]) directly controls noise growth of scalar multiplications (and any operations that use scalar multiplication as a sub-routine) in "Gadget-based" FHE constructions, i.e. smaller quality parameter leads to lower noise growth FHE constructions. Note that gadget encryption is also closely related to GSW-based encryption, see [28].

[3] This is a stronger notion of security than $\mathsf{IND\text{-}CPA}$-security, where one requires ciphertexts be pseudorandom, see Definition 9.

Necessary: Our framework allows us to derive (strong) coding-theoretic bounds on the rate of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,Q]$, for broad classes of E,Q. Our bounds are on the *rate* of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,Q]$, namely we show it can be at most $1-f(n,q,m,\sigma,\delta)$ for explicit functions $f(\cdot)$ of the scheme parameters. Under the assumption that $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,Q]$ meets some notion of correctness (described next), we show universal rate bounds of the above form in the settings of

- **Trivial Quantization**: Arbitrary E, with $Q=\mathbb{Z}^m$, and
- **Small Quantization**: Arbitrary E, with $\sqrt[m]{\det Q}\leq O(\sigma)$ of the same size as the LWE error.

We investigate two correctness notions, namely

- **Bounded Noise**: decryption failure rate $\delta=0$, with respect to bounded noise of the same size (with high probability) as Gaussian noise of parameter σ (in an ℓ_2 ball of radius $\sqrt{m}\sigma$), and
- **Unbounded Noise**: decryption failure rate $\delta>0$, with respect to arbitrary (concentrated) noise of variance σ^2.

For the first correctness notion, we proceed via "packing bounds", while in the second we proceed via "anti-concentration bounds". Throughout, we state the interesting consequences of our bounds for the case of q polynomially large, see Section 5 for full statements.

Our first set of bounds are in the bounded noise model. In this setting, the assumption $\delta=0$ implies that E_q is a *packing* of $\mathbb{R}^m_q$, meaning that for $\mathcal{S}$ the support of the noise (either solely the LWE error, or the sum of the LWE and quantization error), the sets $\{v+\mathcal{S}\}_{v\in E_q}$ are all disjoint, i.e. one can always (uniquely) decode the noisy encoded points $v+\mathcal{S}$ back to $v\in E_q$.

Under the assumption E_q is a packing, we follow a standard volume-based argument (called the *sphere packing* or *Hamming* bound, depending on the context) to obtain an inequality between our parameters of interest. Instantiating this argument in the setting of trivial quantization $Q=\mathbb{Z}^m$ leads to the following bound (Theorem 4).

Bound 1. For any lattice code E, $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,\mathbb{Z}^m]$ has rate at most $1-\Omega(1)$, i.e. rate $1-o(1)$ encryption is impossible.

To handle non-trivial quantization, we require a heuristic assumption (Heuristic 1) that the LWE noise and quantization noise are independent, though we can remove this heuristic for a mild modification of our framework (Section 3.2). Our next bound (Theorem 5) then proceeds in essentially the same way, albeit in the case of small quantization, where the set $\mathcal{S}$ is more complicated.

Bound 2. Under a heuristic assumption, for any lattice codes E,Q, if there exists $\epsilon>0$ such that $\sqrt[m]{\det Q}=\sigma^{1-\epsilon}$, then $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,Q]$ has rate $1-\Omega(1)$, i.e. rate $1-o(1)$ encryption is impossible. If instead $\sqrt[m]{\det Q}\leq O(\sigma)$, then rate $1-o\left(\frac{1}{\log_2 q}\right)$ encryption is impossible.

Therefore, in the bounded error model, quantization is necessary to achieve rate $1-o(1)$ encryption from polynomial modulus. One can further show the aforementioned bounds are tight by repeating the analysis of Corollary 3 in this noise model, though we omit this analysis for brevity.

Our remaining bounds are in the more general setting of δ-correct encryption (for $\delta > 0$) with respect to what is known as *log-concave* noise. We include a brief primer on these random variables in Section 2.5, but for now simply state they include (continuous variants of) all of the noise distributions relevant to public-key lattice-based cryptography, and admit anti-concentration bounds of the form we will require.

The anti-concentration techniques yield bounds with more technical caveats (so *weaker* than the bounded noise model), although one of the bounds is "dimension dependent", which we leverage to give a *stronger* bound than any of our results in the bounded noise model.

Recall that to prove correctness of cryptographic constructions, one often upper bounds the decryption failure rate using concentration inequalities. To prove impossibility results in this noise model, we *lower bound* the decryption failure rate using *anti-concentration* inequalities (Proposition 6), i.e. upper bounds (rather than lower) on how likely it is for a random variable to be close to any particular point (such as its mean).

Our first bound is again for the case of no trivial quantization.

Bound 3. For any lattice code E, either

- the rate of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, \mathbb{Z}^m]$ is $1 - \Omega(1)$, i.e. not rate $1 - o(1)$, or
- the normalized covering radius satisfies $\overline{R}_E = \Omega(m)$.

While this bound is weaker than its analogue in the bounded noise model, we expect this to be a proof artifact — it would be quite peculiar if the way to achieve rate $1 - o(1)$ encryption was to use codes E for error-correction that are very bad *quantizers*[4]. Note that this result does suffice to rule out rate $1-o(1)$ encryption from a class of *a priori* interesting codes (Corollary 8), namely codes E that are nearly optimal for *both* error-correction and quantization. Such codes are known to exist via randomized constructions, and are nearly optimal in many (non-cryptographic) settings.

Our next bound (Theorem 7) again extends our prior bound to the case of $\sqrt[m]{\det Q} \leq O(\sigma)$.

Bound 4. Under a heuristic assumption, for any lattice codes E, Q with $\sqrt[m]{\det Q} \leq O(\sigma)$, the rate of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ is at most

$$1 - \Omega\left(\frac{1}{m \log_2(q/\sigma)}\right).$$

This bound is tight up to the $\log_2(q/\sigma)$ factor. Note that this bound explicitly depends on the dimension m, instead of solely σ, q. This is significant, due to a simple result (Lemma 7) showing that the rates of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ and $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E \otimes \mathbb{Z}^k, Q \otimes \mathbb{Z}^k]$ are equal[5] for any k.

[4] For an indication of how bad $\overline{R}_E = \Omega(m)$ is, the most trivial lattice $\overline{R}_{\mathbb{Z}^m} = \Theta(\sqrt{m})$ is within a constant factor of being an optimal quantizer.

[5] There is a mild caveat that various parameters q, n, δ, σ may (implicitly) depend on $m = \dim E = \dim Q$, and these must be taken to be the same size for both instantiations. This will not impact the conclusions we draw from this bound.

As one can see from Table 1, lattices of this form (for large $k = O(m/\log_2 m)$) are incredibly common in practice. All constructions we are aware of (except for [6]) can be instantiated in our framework using lattices of this type. As a result, one gets a refinement of Bound 4 in this exceedingly common setting.

Bound 5. Under a heuristic assumption, for any lattice codes $E = E' \otimes \mathbb{Z}^{m/\log_2 m}, Q = Q' \otimes \mathbb{Z}^{m/\log_2 m}$ with $\sqrt[m]{\det Q} \leq O(\sigma)$, the rate of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ is at most

$$1 - \Omega\left(\frac{1}{(\log_2 m)\log_2(q/\sigma)}\right).$$

While this is still theoretically rate $1 - o(1)$, practically (for cryptographically relevant dimensions) the convergence is slow. This can be readily observed via concrete comparisons (Figure 1), where we find a practical gap between cryptosystems that satisfy the preconditions of Bound 5 (all of which are rate ≤ 0.9) and those that do not (of rate ≈ 1).

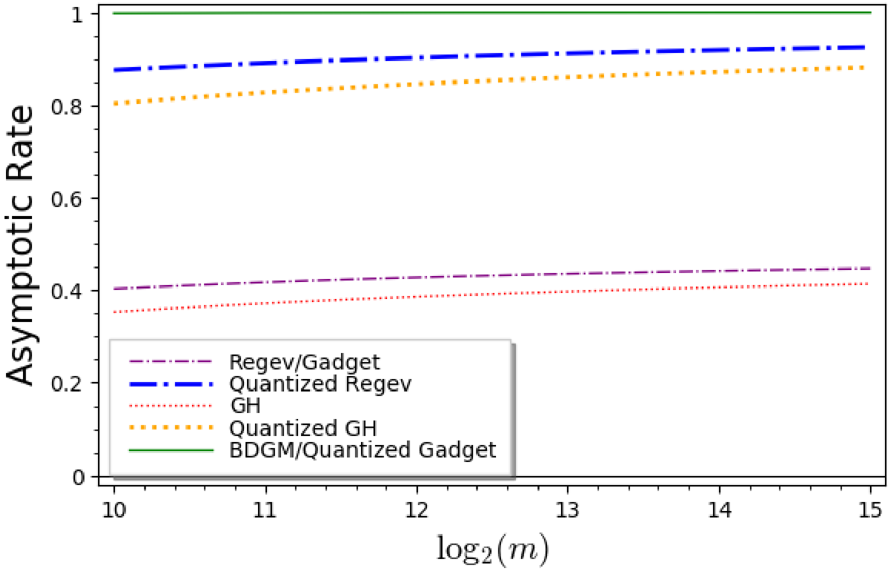

Fig. 1. The rate of the various cryptosystems $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$, for the codes E, Q in Table 1. Throughout, we assume that $q \leq m^2$, $m = n$, $\delta = \exp(-128)$, $\sigma = 2\sqrt{n}$, and then optimize p and k to attain as high rate as possible for $m \in [2^{10}, 2^{15}]$, the range of dimensions included in the Homomorphic Encryption Standard [1].

Fortunately, one can get around this (exponentially) stronger bound by appealing to lattices without this special structure, such as the quantizer $\Lambda_{q/p}(\mathbf{u}^t_m)$ of [6]. As already summarized in Table 1, the we find that the pre-existing scheme of [6] is within an $O(\log_2 m)$ factor of optimal, i.e. beats Bound 5 by a significant margin. We then reuse the quantizer $\Lambda_{q/p}(\mathbf{u}^t_m)$ to quantize messages encoded with a "gadget" $\Lambda_q(\mathbf{g}^t_p) \otimes \mathbb{Z}^{m/\ell}$ (similarly to [18], though with a different "gadget"), while attaining the same (much higher) rate as [6]. We view this construction as simultaneously achieving the best properties of both of [18,6] at no cost[6].

Optimal Decoding for the Quantizer of [6]: Independently of the rest of our work, we give an (optimal) $O(m \log_2 q)$ complexity algorithm (Corollary 1) to solve the closest vector problem on the lattice $\Lambda_q(\mathbf{u}_m^t)$, via a simple reduction to a $O(m \log_2 q)$-time CVP algorithm for the scaled root lattice qA_{m-1}^* [26]. We expect this CVP algorithm to be broadly applicable, due to this quantizer leading to constructions that do not satisfy the preconditions of Bound 5. While $\Lambda_q(\mathbf{u}_m^t)$ is used for quantization in [6], a formal decoding algorithm was not given (instead they focused on bounding the ℓ_∞ covering radius of $\mathcal{V}_{\Lambda_{q/p}(\mathbf{u}_m^t)}$). From the description in [6], there is an obvious sorting-based algorithm of complexity $\Theta(m(\log_2 m)(\log_2 q))$, i.e. slightly slower than our optimal algorithm. Our algorithm also has the benefit of having a simple to analyze distribution of quantization errors, namely for many distributions of random inputs[7] it is uniform over an explicit convex body[8].

Log-Concavity of Distributions Relevant to Lattice-based Cryptography As mentioned before, we leverage the class of *log-concave* random variables. Much of our analysis can be done by simply quoting standard references regarding this topic (for example [34]). To justify the claim that our lower-bounds apply to all noise distributions one encounters in public-key (algebraically-unstructured) lattice-based cryptography, we additionally require that $\langle \mathbf{e}, \mathbf{e}' \rangle$ is log-concave (for $\mathbf{e}, \mathbf{e}'$ independent Gaussians) as well as $\langle \mathbf{e}, \mathbf{e}_K \rangle$ is log-concave (for $\mathbf{e}$ Gaussian, $\mathbf{e}_K$ uniform over a convex body K). We establish these results in Section 2.5, though for simplicity of presentation we focus on the case of private-key encryption in the main body of our paper.

1.2 Related Work

Our framework is similar to those of [33], which parametrizes the design of lattice-based KEMs via two nested[9] (lattice-based) error-correcting codes. Despite these similarities, [33] does not include bounds on constructions built within their framework, and moreover only considers instantiations with $E = E' \otimes \mathbb{Z}^k \subseteq Q' \otimes \mathbb{Z}^k = Q$ sharing a common low-dimensional structure with $\dim E' = \dim Q' = 8$, which by Bound 5 leads to constructions of severely limited rate.

The framework that has the most similar methods to ours is the framework for the construction of lattice-based KEMs of [23]. They parameterize the construction of lattice-based KEMs via novel primitives they call *Key Consensus* and *Asymmetric Key Consensus (AKC)*, and prove inequalities similar to our rate bounds in this setting. In comparison to our work, they require the assumption of perfect correctness ($\delta = 0$), and solely prove impossibility results in the setting of *single dimension* lattices. This leads them to suggest

[6] There may be some poly-logarithmic encoding/decoding cost, but in practice this seems small compared to computing the matrix-vector multiplication as part of LWE-based encryption.

[7] In particular, this holds for what are known as *modulo uniform* distributions, see Chapter 4 of [35].

[8] This is $\mathcal{V}_{\Lambda_{q/p}(\mathbf{u}_m^\ell)}$, which by Lemma 2 is the Minkowski sum of a (scaled) permutahedron and an interval.

[9] Note that our framework does not require a nesting assumption.

lattices of the form $Q = Q' \otimes \mathbb{Z}^k$ for $\dim Q' = O(1)$ as "optimal", which (again by Bound 5) is the opposite of what we find.

There is a relatively large body of work that (essentially) quantizes with $Q = c\mathbb{Z}^m$ a scaled integer lattice, dating back to Peikert's work quantizing LWE-based encryption [29], as well as cryptosystems based on the Learning with Rounding problem [5,12]. Additionally, the "modulus switching" technique [8,7] used in the Fully Homomorphic Encryption literature can be viewed from this perspective.

The work of [19] similarly obtains bounds on (public-key) constructions achievable from LWE with polynomially-large modulus, although they show the impossibility of *non-interactive* key exchange, rather than bounds on the rate of constructions.

Finally, our work is closely related to the currently-known rate $1 - o(1)$ lattice-based encryption schemes [18,6], as a large motivation for our work was to find a way to formally compare the techniques underlying their design.

2 Preliminaries

We write $x \leftarrow \chi$ for the operation of choosing x at random with distribution χ. If S is a finite set, then $x \leftarrow S$ chooses x at random from S with uniform distribution. We write $[\mathbf{A}, \mathbf{B}]$ for horizontal concatenation of matrices, and $(\mathbf{A}, \mathbf{B}) = [\mathbf{A}^t, \mathbf{B}^t]^t$ for vertical concatenation. We write $f(X) = \{f(x) \colon x \in X\}$ for the image of a set $X \subseteq A$ under a function $f \colon A \to B$, and $X + Y = \{x + y \mid x \in X, y \in Y\}$ for the (Minkowski) sum of two subsets $X, Y \subseteq A$ of an abelian group $(A, +, 0)$. We will write $r \cdot \mathcal{B}_n$ for the Euclidean ball of radius r, centered at 0, and $r \cdot \mathcal{B}_n^{(\infty)} = [-r, r)^n$ for the ℓ_∞ ball of radius r. We will write $r \cdot \mathcal{B}_n^{(p)}$ to uniformly refer to either of these objects (but omit p for the more common Euclidean case).

2.1 Lattices

A *lattice* is a discrete subgroup $L \subseteq \mathbb{R}^n$. The *rank* of a lattice is the dimension of the $\mathbb{R}$-subspace that it spans. Any rank k lattice can be written as $\mathbf{B}\mathbb{Z}^k$, where $\mathbf{B} \in \mathbb{R}^{n \times k}$ is a *basis* of its linear span. A lattice is called *full-rank* if its rank equals its dimension. Associated with any lattice L is its *dual lattice* $L^* = \{\mathbf{x} \in \mathsf{span}_{\mathbb{R}}(L) \mid \forall \mathbf{v} \in L, \langle \mathbf{x}, \mathbf{v} \rangle \in \mathbb{Z}\}$. The *determinant* of a lattice $L = \mathbf{B}\mathbb{Z}^k$ is the k-dimensional volume of its fundamental region $\mathbf{B}[0, 1)^k$. The determinant does not depend on the choice of the basis $\mathbf{B}$, and can be efficiently computed as $\det(L) = \sqrt{\det \mathbf{B}^t\mathbf{B}}$, where $\det \mathbf{B}^t\mathbf{B}$ is the matrix determinant of $\mathbf{B}^t\mathbf{B} \in \mathbb{R}^{k \times k}$.

We say that L is a q-ary lattice if $q\mathbb{Z}^m \subseteq L$, i.e., L is periodic modulo q. Notice that q-ary lattices are always full rank, and the vectors of a q-ary lattice do not necessarily have integer coordinates. There are two standard q-ary integer lattices associated to any matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$:

$$\begin{aligned} \Lambda_q^\perp(\mathbf{A}) &= \{\mathbf{x} \in \mathbb{Z}^m \mid \mathbf{A}\mathbf{x} \equiv 0 \bmod q\}, \\ \Lambda_q(\mathbf{A}) &= \{\mathbf{y} \in \mathbb{Z}^m \mid \exists \mathbf{x} \in \mathbb{Z}_q^m \text{ s.t. } \mathbf{A}^t\mathbf{x} = \mathbf{y} \bmod q\}. \end{aligned}$$

These lattices are scaled duals of each other, meaning $\Lambda_q(\mathbf{A})^* = \frac{1}{q}\Lambda_q^\perp(\mathbf{A})$. For a q-ary lattice, we define the scaled dual as $L^\perp = qL^*$, which is such that $\Lambda_q(\mathbf{A})^\perp = \Lambda_q^\perp(\mathbf{A})$. We say that a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ is *primitive* if $\mathbf{A}\mathbb{Z}_q^m = \mathbb{Z}_q^n$, i.e. it is a surjection. For primitive matrices $\mathbf{A}$, $\det(\Lambda_q(\mathbf{A})) = q^{m-n}$ and $\det(\Lambda_q^\perp(\mathbf{A})) = q^n$.

We say that two full-rank lattices L, L' are *nested* if $L' \subseteq L$. Given nested lattices $L' \subseteq L$, the quotient L/L' forms a group of size $\frac{\det L'}{\det L} \in \mathbb{N}$, and therefore $\det(L)$ divides $\det(L')$. Any two lattices $L \subset \mathbb{R}^n$ and $L' \subset \mathbb{R}^{n'}$, can be combined into the *direct sum* $L \oplus L' \subset \mathbb{R}^{n+n'}$, and the *tensor product* $L \otimes L' \subset \mathbb{R}^{n\cdot n'}$. The direct sum is simply the Cartesian product of the two lattices $L \oplus L' = L \times L'$, obtained by concatenating vectors from L and L'. If $\mathbf{A}$ and $\mathbf{B}$ are bases of L and L', then the tensor product $L \otimes L'$ is the lattice with basis $\mathbf{A} \otimes \mathbf{B}$ given by the Kronecker product of $\mathbf{A}$ and $\mathbf{B}$, i.e., the block matrix obtained replacing each entry $a_{i,j}$ of $\mathbf{A}$ with the block $a_{i,j} \cdot \mathbf{B}$. The tensor product $L \otimes L'$ satisfies $\det(L \otimes L') = \det(L')^n \cdot \det(L)^{n'}$. The k-fold direct sum of a lattice $L^{\oplus k} = \oplus_{i=1}^k L$ can be equivalently expressed as the tensor product $L^{\oplus k} = \mathbb{Z}^k \otimes L$.

2.2 Convex Bodies

We say a set $K \subseteq \mathbb{R}^n$ is *convex* if, for any $\mathbf{x}, \mathbf{y} \in K$, and $t \in [0,1]$, $(1-t)\mathbf{x} + t\mathbf{y} \in K$. We furthermore say K is *symmetric* if $\mathbf{x} \in K \iff -\mathbf{x} \in K$. Associated with any convex symmetric set K is a *norm* $\|\mathbf{x}\|_K = \inf\{t > 0 \mid \mathbf{x}/t \in K\}$. For such K, we define the *ℓ_p-packing radius* $r_K^{(p)}$ to be the maximal r such that $r \cdot \mathcal{B}_n^{(p)} \subseteq K$. Similarly, we define the *ℓ_p-covering radius* $R_K^{(p)}$ to be the minimal R such that $K \subseteq R \cdot \mathcal{B}_n^{(p)}$. Again, when p is omitted, we mean $p = 2$. For a pair of convex symmetric sets K, K', we write $\|K'\|_K := \sup_{\mathbf{x}\in K'} \|\mathbf{x}\|_K$. We will need the following bounds, which are straightforward to derive.

Lemma 1. *Let K, K' be convex symmetric sets in $\mathbb{R}^n$. Then*

1. *if $K \subseteq K'$, then for all $\mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\|_K \geq \|\mathbf{x}\|_{K'}$,*
2. *if $s > 0$, then for all $\mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\|_{sK} = \frac{1}{s}\|\mathbf{x}\|_K$, and*
3. $\|K'\|_K \in \left[\frac{R_{K'}^{(p)}}{R_K^{(p)}}, \frac{R_{K'}^{(p)}}{r_K^{(p)}}\right]$.

2.3 Lattice Codes

Applications of lattices often require not only a lattice L, but also an efficient algorithm to map arbitrary vectors $\mathbf{x} \in \mathbb{R}^n$ to a nearby lattice point.

Definition 1. *A* lattice code $(L, \lfloor\cdot\rceil)$ *is a lattice $L \subset \mathbb{R}^n$ together with a rounding algorithm $\lfloor\cdot\rceil : \mathsf{span}_{\mathbb{R}}(L) \to L$ such that $\lfloor\mathbf{0}\rceil = \mathbf{0}$ and $\lfloor\mathbf{x} + \mathbf{v}\rceil = \lfloor\mathbf{x}\rceil + \mathbf{v}$ for all $\mathbf{x} \in \mathsf{span}_{\mathbb{R}}(L)$ and $\mathbf{v} \in L$.*

We will be primarily interested in q-ary lattice codes, i.e., lattice codes $(L, \lfloor\cdot\rceil)$ such that L is a q-ary (but not necessarily integer) lattice. For any q-ary lattice code $L \subset \mathbb{R}^n$, we can take the quotients of L and $\mathbb{R}^n$ modulo the additive subgroup $q\mathbb{Z}^n$, and define the *codebook*

$L_q = L/q\mathbb{Z}^n$, and *ambient torus* $\mathbb{R}^n_q = (\mathbb{R}/q\mathbb{Z})^n \equiv \mathbb{R}^n/q\mathbb{Z}^n$. Elements of the codebook L_q are called *codewords*, and can be represented as vectors $L \cap [0,q)^n$ with (not necessarily integer) coordinates in the range $[0,q)$. Given a $\mathbb{Z}$-basis of the lattice $\mathbf{B}$, one can moreover represent these codewords as integer via the encoding function $\mathsf{encode}_L(\mathbf{m}) := \mathbf{Bm} \bmod q$, and decoding function $\mathsf{decode}_L(\mathbf{c}) := \mathbf{B}^{-1}\lfloor \mathbf{c} \rceil_L \bmod q$. The codebook L_q is a subgroup of the ambient torus $\mathbb{R}^n_q$, and the rounding function $\lfloor\cdot\rceil : \mathbb{R}^n \to L$ induces a well-defined map $\mathbb{R}^n_q \to L_q$ from the ambient torus to the codebook. Notice that the codebook L_q is a finite set of size $|L_q| = \frac{q^n}{\det(L)}$, so codewords can be represented with $\lceil \log_2 |L_q| \rceil \approx n \log q - \log\det(L)$ bits.

For any lattice code $(L, \lfloor\cdot\rceil_L)$, we define the fundamental decoding region $\mathcal{V}_L = \{\mathbf{x} \in \mathbb{R}^n \colon \lfloor \mathbf{x} \rceil_L = \mathbf{0}\}$, i.e., the set of all points that decode to $\mathbf{0}$. When $\lfloor\cdot\rceil$ is the CVP rounding function, $\mathcal{V}_{\mathsf{CVP}_L}$ is called the *Voronoi cell* of the lattice. The reduction of a point $\mathbf{x} \in \mathbb{R}^n$ modulo a lattice code $(L, \lfloor\cdot\rceil_L)$ is defined as $[\mathbf{x}]_L = \mathbf{x} - \lfloor \mathbf{x} \rceil_L$, so that every point in space can be (uniquely) written as the sum $\mathbf{x} = \lfloor \mathbf{x} \rceil_L + [\mathbf{x}]_L$ of a lattice point $\lfloor \mathbf{x} \rceil_L \in L$ and a rounding error $[\mathbf{x}]_L \in \mathcal{V}_L$ in the fundamental decoding region. Notice that the rounding error depends not only on the lattice L but also on the rounding function $\lfloor\cdot\rceil$ of the lattice code.

Throughout, we will assume that $\mathcal{V}_L$ is a convex symmetric set. When the choice of $\lfloor\cdot\rceil$ is unambiguous, we will refer to the norm $\|\cdot\|_L := \|\cdot\|_{\mathcal{V}_L}$, packing radius $r^{(p)}_L := r^{(p)}_{\mathcal{V}_L}$, and covering radius $R^{(p)}_L := R^{(p)}_{\mathcal{V}_L}$ of L. Note that when $\lfloor\cdot\rceil$ solves CVP on L, the parameters r_L and R_L are the familiar lattice parameters $\lambda_1(L)/2$ and $\rho(L)$. When discussing bounds on the packing/covering radii, we will find it useful to work with normalized (to be invariant to scaling $L \mapsto cL$) versions of these quantities $\overline{r} = (\det L)^{-1/n} r$ and $\overline{R} = (\det L)^{-1/n} R$.

Some Explicit Lattice Codes We briefly summarize some explicit lattice codes we will use in our work, namely the lattice codes (implicitly) used in previous high-rate constructions of LWE-based encryption [18,6] (we justify this claim in Section 4).

Definition 2 (Primal Gadget Lattice). *For $p, q \in \mathbb{N}$, let $\mathbf{g}_p = (1, p, p^2, \ldots, p^{\lceil \log_p q \rceil - 1})$ be the base-p "gadget vector". The* primal gadget lattice *is the lattice $\Lambda_q(\mathbf{g}_p^t)$.*

Proposition 1. *Let $q = p^\ell$. Then the fundamental region when decoding with Babai's nearest planes $\mathcal{V}^{\mathsf{babai}}_{\Lambda_q^\perp(\mathbf{g}_p^t)} = \frac{q}{2p} \cdot \mathcal{B}^{(\infty)}_\ell$, and $\det \Lambda_q(\mathbf{g}_p^t) = q^{\ell-1}$. Moreover, $\det \Lambda_q(\mathbf{g}_p^t) \otimes \mathbb{Z}^{m/\ell} = \det((q/p)\mathbb{Z}^m)$.*

Proof. The fundamental region statement is from [27, Section 4], and the determinant calculation is straightforward. □

Definition 3 (Dual Gadget Lattice). *For $p, q \in \mathbb{N}$, let $\mathbf{g}_p = (1, p, p^2, \ldots, p^{\lceil \log_p q \rceil - 1})$ be the base-p "gadget vector". The* dual gadget lattice *is the lattice $\Lambda_q^\perp(\mathbf{g}_p^t)$.*

Proposition 2. *Let $p < q$, and let $\ell = \lceil \log_p q \rceil$. Then there exists a decoding algorithms for $\Lambda_q^\perp(\mathbf{g}_p^t)$ that satisfy*

- *when $q = p^\ell$, $r^{(\infty)}_K \geq p/2$,*

– *when* $q = p^\ell - 1$, $r_K^{(\infty)} \geq (p-1)/2$,
– *when* $q \in \mathbb{N}$, $r_K^{(\infty)} \geq \frac{(p-1)}{2}\frac{q}{p^\ell}$.

Proof. The case of $q = p^\ell$ follows from [27]. The case of $q = p^\ell - 1$ follows from [18] (we show that their "nearly square gadget matrix" is the dual gadget lattice in Section 4.2). The case of arbitrary q is implicit in [17] (it follows from standard analysis of a decoding algorithm they suggest). We provide this standard analysis below.

Let $S_q = [\mathbf{b}_0, \ldots, \mathbf{b}_{\ell-2}, \mathbf{q}]$ be the typical basis of $\Lambda_q^\perp(\mathbf{g}_p^t)$, where $\mathbf{b}_i = p\mathbf{e}_i - \mathbf{e}_{i+1}$, and $\mathbf{q} = (q_0, \ldots, q_{\ell-1})$ are the base-p digits of q. We abuse notation and state that $q = p^\ell$ has base-p decomposition of $(0, 0, \ldots, 0, p)$. The authors of [17] note that S_q admits a factorization as $S_q = S_{p^\ell} D_q$ where $D_q = [\mathbf{e}_0, \ldots, \mathbf{e}_{\ell-2}, \mathbf{d}_{p,q}]$ for the vector $\mathbf{d}_{p,q}$ with coefficients $\langle \mathbf{e}_i, \mathbf{d}_{p,q}\rangle = \frac{q \bmod p^{i+1}}{p^{i+1}}$. They then suggest using the decoder

$$\mathsf{decode}(x) = S_{p^\ell}\mathsf{decode}_{D_q\mathbb{Z}^\ell}(S_{p^\ell}^{-1}x), \tag{2}$$

where one decodes $D_q\mathbb{Z}^\ell$ using Babai's Nearest Planes. This has fundamental region that contains $\frac{q}{p^\ell}[-1/2, 1/2)^\ell$, and therefore the decoder of Eq. (2) has fundamental region that contains $S_{p^\ell}\frac{q}{p^\ell}[-1/2, 1/2)^\ell$. One can readily compute that this set contains $(p-1)\frac{q}{p^\ell}[-1/2, 1/2)^\ell$.

We omit the computation of the determinant, as it is straightforward. □

The next lattice belongs to parameterized family of lattices (for $\mathbf{u}_m^t = (1, 1, \ldots, 1) \in \mathbb{R}^m$) $\Lambda_q(\mathbf{u}_m^t)$ that we call the *Dual of Davenport's Lattice*. Well-known special cases are

– $q = 1$, where it is simply $\mathbb{Z}^m$, and
– $q = 2$, where it is a scaling of D_m^*, the dual of the standard $D_m = \Lambda_2(\mathbf{u}_m^t)$ root lattice.

The generalization to $m > 2$ has been implicit in many works, namely constructing explicit efficient coverings of $\mathbb{R}^m$ [13][11, Chapter 2, Section 1.3], constructing efficient decoding algorithms for certain lattices [15], and constructing rate $1 - o(1)$ fully homomorphic encryption [6].

Definition 4 (Scaled Dual of Davenport's Lattice). *Let* $m, q \in \mathbb{N}$. *The scaled dual of Davenport's lattice* $\Lambda_q(\mathbf{u}_m^t)$ *is the lattice* $\Lambda_q(\mathbf{u}_m^t) = q\mathbb{Z}^m + \mathbb{Z} \cdot \mathbf{u}_m$, *where* $\mathbf{u}_m$ *is the all-ones vector of length* m.

Definition 5 (A_{m-1}^* Lattice). *For any* $m \in \mathbb{N}$, *the* A_{m-1}^* *lattice is defined to be the projection of* $\mathbb{Z}^m$ *perpendicular to the vector* $\mathbf{u}_m$.

When $m \mid q$, this lattice admits a simple orthogonal decomposition in terms of the root lattice A_{m-1}^*, which admits an $O(m)$-arithmetic operation CVP algorithm [26].

Lemma 2. *Provided* $m \mid q$, $\Lambda_q(\mathbf{u}_m^t) = qA_{m-1}^* + \mathbb{Z} \cdot \mathbf{u}_m$, *where* $\langle qA_{m-1}^*, \mathbb{Z} \cdot \mathbf{u}_m\rangle = \{0\}$.

Proof. As A_{m-1}^* is defined to be a projection orthogonal to $\mathbf{u}_m$, the last condition is immediate. One can check that $qA_{m-1}^*, \mathbf{u}_m$ are the projections of $\Lambda_q(\mathbf{u}_m^t)$ onto the subspaces perpendicular to and parallel to $\mathbf{u}_m$, respectively, so $qA_{m-1}^* + \mathbf{u}_m \supseteq \Lambda_q(\mathbf{u}_m^t)$. For the other direction, note that $\mathbf{u}_m \subseteq \Lambda_q(\mathbf{u}_m^t)$, as $\Lambda_q(\mathbf{u}_m^t) = q\mathbb{Z}^n + \mathbb{Z} \cdot \mathbf{u}_m$. The equality then immediately follows by [25, Proposition 1.1.6], which implies that the indicies of

- the intersection of $\Lambda_q(\mathbf{u}_m^t) \cap \mathbb{R} \cdot \mathbf{u}_m$ within the projection of $\Lambda_q(\mathbf{u}_m^t)$ onto $\mathbb{R} \cdot \mathbf{u}_m$, and
- the index of $\Lambda_q(\mathbf{u}_m^t)$ within $qA_{m-1}^* + \mathbb{Z} \cdot \mathbf{u}_m$,

are equal. As it is clear that the first index is 1, we have that $\Lambda_q(\mathbf{u}_m^t) = qA_{m-1}^* + \mathbb{Z} \cdot \mathbf{u}_m$. □

This same argument works for $m \nmid q$, though the indices mentioned in the proof are not all equal to 1. It is fairly straightforward to verify that they are instead equal to $\frac{m}{\gcd(q,m)}$, so one gets an $O(m^2)$-arithmetic operation CVP algorithm for $\Lambda_q(\mathbf{u}_m^t)$ in general. This parameter setting does not appear to be useful for our setting though, as it is unclear how to get any useful information about the shape of the Voronoi cell of the lattice in general.

Proposition 3. *Let $m \mid q$. Then $R^{(\infty)}_{\Lambda_q(\mathbf{u}_m^t)} = \frac{q}{2}\left(1 - \frac{1}{m}\right) + \frac{1}{2}$, and $\det \Lambda_q(\mathbf{u}_m^t) = q^{m-1}$.*

Proof. The orthogonal decomposition implies that $\mathcal{V}_{\Lambda_q(\mathbf{u}_m^t)} = \mathcal{V}_{qA_{m-1}^*} + \mathcal{V}_{\mathbb{Z}\cdot\mathbf{u}_m}$. Applying triangle inequality, we can reduce computing $R^{(\infty)}_{\Lambda_q(\mathbf{u}_m^t)}$ to computing both $R^{(\infty)}_{qA_{m-1}^*}$ and $R^{(\infty)}_{\mathbb{Z}\cdot\mathbf{u}_m}$. The first is straightforward to compute given the explicit expression (found in [11, Chapter 4, Section 6.6]) for $\mathcal{V}_{A_{m-1}^*}$, namely as the convex hull of all coordinate permutations of the explicit vector $\mathbf{v} = \frac{1}{2m}(-m+1, -m+3, \ldots, m-3, m-1)$. The second is straightforward to compute as $1/2$.

Finally, to compute the determinant, note that the lattice may be generated by the $m+1$ vectors $[\mathbf{u}_m, q\mathbf{e}_1, \ldots, q\mathbf{e}_m]$, and that any single vector $q\mathbf{e}_i$ can easily be written as a linear combination of the other vectors in this generating set. It follows that $[q\mathbf{e}_1, q\mathbf{e}_2, \ldots, q\mathbf{e}_{m-1}, \mathbf{u}_m]$ is a triangular basis, and $\det \Lambda_q(\mathbf{u}_m^t) = q^{m-1}$.

□

Corollary 1. *If $m \mid q$, one can solve CVP on $\Lambda_q(\mathbf{u}_m^t)$ in $O(m \log_2 q)$ time.*

Proof. Project parallel/perpendicular to $\mathbf{u}_n$, then use the known $O(m)$-arithmetic operation CVP algorithms on qA_{m-1}^* [26] and $\mathbb{Z} \cdot \mathbf{u}_m$. There is an additional $O(\log_2 q)$ overhead as the algorithm of [26] costs arithmetic operations at unit cost. □

2.4 Bounds on Lattice Parameters

For any lattice L, the best normalized packing and covering radii are achieved by the CVP rounding algorithm, giving $\overline{r}_L$ and $\overline{R}_L$. For any m, let $\overline{r}_m = \sup_L \overline{r}_L$ and $\overline{R}_m = \inf_L \overline{R}_L$ be the optimal normalized radii over all lattices L of rank m. It is known that $\overline{r}_m = \Theta(\sqrt{m})$, and $\overline{R}_m = \Theta(\sqrt{m})$ (see Chapters 1 and 2 of [11]). It is additionally known that in each dimension m, there are lattices $L \subseteq \mathbb{R}^m$ that (nearly) simultaneously achieve these bounds, meaning such that $\overline{R}_L/\overline{r}_L \leq 2 + o(1)$, see [9].

2.5 Probability

We define the Gaussian kernel to be $\rho_\sigma(x) = \exp(-x^2/2\sigma^2)$. Let $\mathcal{N}(0, \sigma^2 I_n)$ be the multivariate (continuous) Gaussian, with probability density function $f_\sigma(\mathbf{x}) = \frac{1}{\sqrt{(2\pi\sigma^2)^n}}\rho_\sigma(\|\mathbf{x}\|_2)$. We say that a random vector is *isotropic* if it is mean zero and has identity covariance matrix.

We will require the class of *log-concave* random variables.

Definition 6. *Let X be a random variable with pdf $p(x)$. We say that X is* log-concave *if $p(x) = \exp(-V(x))$ for $V(x)$ a convex function.*

We briefly summarize (from [34]) the properties this class of random variables satisfies.

Proposition 4. *Let $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^n$ be log-concave and independent. Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ be any linear transformation. Then $\mathbf{x}+\mathbf{x}'$ and $\mathbf{A}\mathbf{x}$ are log-concave.*

Standard examples of log-concave random variables are Gaussians, and uniform random variables on convex sets K. We establish log concavity of a few other distributions relevant to lattice-based cryptography at the end of this sub-section.

Log-concave random variables are known to have strong concentration properties (they are "sub-exponential"). We use the following concentration bound mostly for simplicity of exposition — one can obtain tighter bounds by treating the cases of the $\|\cdot\|_2$ and $\|\cdot\|_\infty$ norms separately, though as we mention later (Section 4) this never impacts our (asymptotic) results.

Proposition 5 (Theorem 11 of [24]). *For any L-Lipschitz function $g \in \mathbb{R}^n$, if $\mathbf{x}$ is an isotropic log-concave random variable, then $\Pr[|g(\mathbf{x}) - \mathbb{E}[g(\mathbf{x})] > Lt|] \leq \exp(-\Omega(t\psi_n^{-1}))$.*

Here, ψ_n is *KLS constant*, which is (under the celebrated *KLS conjecture*) $O(1)$ as $n \to \infty$. The current best bound known is $\psi_n = O((\log n)^c)$ for $c = 3.2226$ [22]. In the rest of our work we will write $\exp(-\tilde{\Omega}(t))$, where this is understood to mean $\exp(-\Omega(t/(\log n)^c))$.

Corollary 2. *If $\mathbf{x}$ is a log-concave random variable in $\mathbb{R}^n$ with covariance matrix Σ, then for $p \in \{2, \infty\}$*

$$\Pr[\|\mathbf{x}\|_p > \sqrt{\mathsf{Tr}(\Sigma)}(t+\sqrt{n})] \leq \exp(-\tilde{\Omega}(t)).$$

Proof. Note that $\Sigma^{-1/2}\mathbf{x}$ is isotropic, so we will apply the previous proposition to this random variable and $g(\mathbf{x}) = \|\Sigma^{1/2}\mathbf{x}\|_p$. For the ℓ_2 norm, the Lipschitz constant is the ℓ_2 to ℓ_2 operator norm, i.e. the maximum singular value of $\Sigma^{1/2}$, which is at most $\sqrt{\mathsf{Tr}(\Sigma)}$. For the ℓ_∞ norm, the Lipschitz constant is the ℓ_∞-ℓ_2 operator norm, i.e. the maximum ℓ_2 norm of a column of $\Sigma^{1/2}$. Note that each element of the main diagonal of Σ is the (squared) ℓ_2 norm of a column of Σ, so again we get that $\sqrt{\mathsf{Tr}(\Sigma)}$ bounds the Lipschitz constant.

We therefore have reduced to bounding $\mathbb{E}[g(\mathbf{x})]$ in both cases. For the ℓ_2 norm, by Jenson's inequality, we have that $\mathbb{E}[\|\mathbf{x}\|_2]^2 \leq \mathbb{E}[\|\mathbf{x}\|_2^2] = \mathsf{Tr}(\Sigma)$. For the ℓ_∞ norm, we apply the bound $\mathbb{E}[\|\mathbf{x}\|_\infty] \leq \mathbb{E}[\|\mathbf{x}\|_2] \leq \sqrt{\mathsf{Tr}(\Sigma)}$. □

We next introduce our *anti-concentration* inequality, which (in a general form) holds for arbitrary polynomials in log-concave random variables. For $t \in \mathbb{R}$ we apply it to the degree-2 polynomial $\|\mathbf{x}\|_2^2 - t$.

Proposition 6 (Theorem 8 of [10]). *If $\mathbf{x}$ is a log-concave random variable on $\mathbb{R}^n$ with covariance matrix Σ, then for every $\epsilon > 0$,*

$$\Pr[|\|X\|_2 - t| \leq \epsilon] \leq O\left(\frac{\epsilon}{\sqrt{\mathsf{Tr}(\Sigma)}}\right).$$

We end the sub-section by establishing log-concavity of some distributions of cryptographic interest.

Lemma 3. *Let* $\mathbf{e}_i \sim \mathcal{N}(0, \sigma_i^2 I_n)$ *for* $i \in \{0,1\}$*. Then the distribution of* $\langle \mathbf{e}_0, \mathbf{e}_1\rangle$ *is log-concave if* $n \geq 2$*.*

Proof. By [16, Eq. 2.15], one has that $\langle \mathbf{e}_0, \mathbf{e}_1\rangle = \frac{\sigma_0\sigma_1}{2}(V - V')$ as distributions, where V, V' are independent $\chi^2_{(n)}$ random variables. One can easily verify (by directly examining the pdf) that a $\chi^2_{(n)}$ random variable is log-concave if $n \geq 2$. By closure of log concavity under independent sums, the claimed result follows. □

Theorem 1. *Let* $n \geq 8$*, and let* K *be a bounded measurable subset of* $\mathbb{R}^n$*. Let* $\mathbf{x} \sim \mathcal{N}(0, \sigma^2 I_n)$*, and let* $\mathbf{y} \sim K$ *be independent from* $\mathbf{x}$*. Then* $\langle \mathbf{x}, \mathbf{y}\rangle$ *is log-concave.*

Note that by applying orthogonal transformations to both $\mathbf{x}, \mathbf{y}$, this implies log concavity in the more general case of $\mathbf{x} \sim \mathcal{N}(0, \Sigma)$.

Proof. One can verify that univariate $p(x)$ is log-concave if

$$\forall x : p(x)p''(x) \leq (p'(x))^2.$$

We will explicitly compute the pdf of $\langle \mathbf{x}, \mathbf{y}\rangle$, and show that it satisfies this inequality. Note that by the law of total probability

$$\begin{aligned}\mu(A) := \Pr[\langle \mathbf{x}, \mathbf{y}\rangle \in A] &= \int_{\mathbb{R}^n} \Pr[\langle \mathbf{x}, \mathbf{y}\rangle \in A \mid \mathbf{y} = \mathbf{z}] \Pr[\mathbf{y} = \mathbf{z}] d\mathbf{z}\\ &= \int_{\mathbb{R}^n} \left(\int_A \sqrt{2\pi\sigma^2\|\mathbf{z}\|_2^2}^{-1} \exp\left(-\frac{x^2}{2\sigma^2\|\mathbf{z}\|_2^2}\right) dx\right) \mathsf{vol}(K)^{-1}\chi_K(\mathbf{z}) d\mathbf{z}\\ &= \int_A \left(\mathsf{vol}(K)^{-1} \int_K \sqrt{2\pi\sigma^2\|\mathbf{z}\|_2^2}^{-1} \exp\left(-\frac{x^2}{2\sigma^2\|\mathbf{z}\|_2^2}\right)\right) dx,\end{aligned}$$

where in the last step we applied Fubini's theorem and simplified. If we define $f(x, y^2) = \mathsf{vol}(K)^{-1}\sqrt{2\pi\sigma^2 y^2}^{-1} \exp\left(-\frac{x^2}{2\sigma^2 y^2}\right)$, it follows that the density is $p(x) = \int_K f(x, \|\mathbf{z}\|_2^2) d\mathbf{z}$.

To compute the derivatives $p'(x), p''(x)$, we need to interchange differentiation and integration a few times, which we do via the (measure-theoretic) Leibniz Integral rule. Before discussing this, we compute that $\partial_x f(x, y^2) = -\frac{x}{\sigma^2 y^2} f(x, y^2)$, $\partial_x^2 f(x, y^2) = \left(\frac{x^2}{\sigma^4 y^4} - \frac{1}{\sigma^2 y^2}\right) f(x, y^2)$, $\partial_x^3 f(x, y^2) = -\left(\frac{x^3}{\sigma^6 y^6} - \frac{3x}{\sigma^4 y^4}\right) f(x, y^2)$. Our applications of the Leibniz integral rule will require all of these functions (as well as $f(x, y^2)$ itself) to be integrable for all x. The largest singularity at this (or any) point is y^{-7}. As switching to spherical coordinates introduces a multiplicative factor y^{n-1}, provided $n \geq 8$ we can switch to spherical coordinates to get an integrand with no singularity, and show convergence. Note that this step is additionally where we require K to be bounded and measurable, as otherwise $\int_{\mathbb{R}^n} f(x, \|\mathbf{z}\|_2^2) d\mathbf{z} = \infty$ for the same reason that $\int_{\mathbb{R}^n} \|\mathbf{z}\|_2^2 d\mathbf{z} = \infty$. As the other preconditions of Leibniz are straightforward to verify, we omit them.

We next note that one can write

$$\mathbb{E}_{\mathbf{x}}[f(\mathbf{x})]\mathbb{E}_{\mathbf{x}}[g(\mathbf{x})] = \mathbb{E}_{\mathbf{x},\mathbf{y}}\left[\frac{f(\mathbf{x})g(\mathbf{y}) + f(\mathbf{y})g(\mathbf{x})}{2}\right],$$

where $\mathbf{y}$ is an i.i.d. copy of $\mathbf{x}$. It follows that

$$(p'(x))^2 = \mathbb{E}_{\mathbf{z},\mathbf{z}'}\left[\frac{x^2}{\sigma^4}\|\mathbf{z}\|_2^{-2}\|\mathbf{z}'\|_2^{-2}f(x,\|\mathbf{z}\|_2^2)f(x,\|\mathbf{z}'\|_2^2)\right],$$

and

$$p(x)p''(y) = \mathbb{E}_{\mathbf{z},\mathbf{z}'}\left[\left(\frac{x^2}{\sigma^4}\left(\frac{\|\mathbf{z}\|_2^{-4} + \|\mathbf{z}'\|_2^{-4}}{2}\right) - \frac{1}{\sigma^2}\left(\frac{\|\mathbf{z}\|_2^{-2} + \|\mathbf{z}'\|_2^{-2}}{2}\right)\right)f(x,\|\mathbf{z}\|_2^2)f(x,\|\mathbf{z}'\|_2^2)\right].$$

Therefore establishing the inequality $(p'(x))^2 \geq p''(x)p(x)$ reduces to showing that some explicit integral is non-negative. Note that $f(x, \|\mathbf{z}\|_2^2) \geq 0$ by inspection. We therefore reduce to showing that the integrand

$$\frac{x^2}{\sigma^4}\|\mathbf{z}\|_2^{-2}\|\mathbf{z}'\|_2^{-2} - \left(\frac{x^2}{\sigma^4}\left(\frac{\|\mathbf{z}\|_2^{-4} + \|\mathbf{z}'\|_2^{-4}}{2}\right) - \frac{1}{\sigma^2}\left(\frac{\|\mathbf{z}\|_2^{-2} + \|\mathbf{z}'\|_2^{-2}}{2}\right)\right) \geq 0.$$

This itself follows from the bound $x^{-2}y^{-2} \geq \frac{x^{-4}+y^{-4}}{2}$, valid for any positive x, y, which in the more familiar form $\left(\frac{x^{-4}+y^{-4}}{2}\right)^{-1} \leq x^2y^2$ is simply the inequality between the Harmonic and Geometric means, applied to (x^4, y^4).

2.6 The Learning with Errors Problem

Much of lattice cryptography relies on the hardness of the *learning with errors* problem.

Definition 7 (LWE problem). *Let $m = n^{O(1)}$, and let $q \in [n^{O(1)}, 2^{O(n)}]$. Let χ_{sk} be a distribution on $\mathbb{Z}_q$, and χ_e be a distribution on $\mathbb{R}_q$. The Learning with Errors problem $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}$ is to distinguish the distribution $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ from $(\mathbf{A}, \mathbf{u})$, where $\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}$, $\mathbf{s} \leftarrow \chi_{\mathsf{sk}}^n$, and $\mathbf{e} \leftarrow \chi_e^m$, and $\mathbf{u} \leftarrow \mathbb{R}_q^m$.*

We rely on LWE where $\mathbf{e} \leftarrow \chi_e$ and $\mathbf{u} \leftarrow \mathbb{R}_q^m$ are *real* random variables (modulo q) to simplify our analysis. We omit the inclusion of m in the notation $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}$, as it has minimal impact on the hardness of the problem. The primary justification for the hardness of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}$ is that it admits reductions from worst-case hard lattice problems, initially due to Regev [32].

Theorem 2. *For any $m = n^{O(1)}$, any modulus $q \leq 2^{n^{O(1)}}$, let χ_e be any (discretized) Gaussian distribution χ of parameter $\sigma \geq 2\sqrt{n}$, and χ_{sk} be the uniform distribution on $\mathbb{Z}_q$. Then solving the decision $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}$ problem is at least as hard as quantumly solving GapSVP_γ and SIVP_γ on arbitrary n-dimensional lattices, where $\gamma = \tilde{O}(nq/\sigma)$.*

This work will not need definitions of of GapSVP_γ and SIVP_γ. We call attention to the approximation factor $\gamma = \tilde{O}(nq/\sigma)$, which controls the hardness of the problem, and depends on the "modulus to noise" ratio q/σ. The Gaussian parameter can often be set to a fixed polynomial $\sigma = 2\sqrt{n}$, so that larger values of q result in constructions that are both less efficient and less secure. Of particular interest will be the cases of *polynomial* $q/\sigma = n^{O(1)}$, and *superpolynomial* $q/\sigma = n^{\omega(1)}$ modulus to noise ratio.

2.7 Cryptographic Primitives

We will use the standard notion of IND-CPA security, as well as a less standard notion (that is better suited to lattice-based primitives) known as RND-CPA.

Definition 8 (IND-CPA). *An encryption scheme* $(\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ *is said to be* indistinguishable under chosen plaintext attack *if any efficient (probabilistic polynomial-time) adversary $\mathcal{A}$ can only achieve at most negligible advantage in the following game, parameterized by a bit $b \in \{0, 1\}$:*

1. $k \leftarrow \mathsf{KGen}(1^n)$,
2. $b' \leftarrow \mathcal{A}^{O_b(\cdot,\cdot)}$, *where* $O_b(m_0, m_1) = \mathsf{Enc}_k(m_b)$.

The adversary's advantage is defined to be $\mathsf{Adv}(\mathcal{A}) = |\Pr[b' = 1 \mid b = 0] - \Pr[b' = 1 \mid b = 1]|$.

Definition 9 (RND-CPA). *An encryption scheme* $(\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ *is said to be* pseudorandom under chosen plaintext attack *if any efficient (probabilistic polynomial-time) adversary $\mathcal{A}$ can only achieve at most negligible advantage in the following game, parameterized by a bit $b \in \{0, 1\}$:*

1. $k \leftarrow \mathsf{KGen}(1^n)$,
2. $b' \leftarrow \mathcal{A}^{O_b(\cdot)}$, *where* $O_b(m)$ *returns either*
   - $b = 0$*: an encryption* $\mathsf{Enc}_k(m)$ *of the message* m *under the key* k*, or*
   - $b = 1$*: a sample from a distribution that has support* $\{\mathsf{Enc}_k(m) \mid k \in \mathsf{supp}(\mathsf{KGen}(1^n)), m \in \mathcal{M}\}$.

The adversary's advantage is defined to be $\mathsf{Adv}(\mathcal{A}) = |\Pr[b' = 1 \mid b = 0] - \Pr[b' = 1 \mid b = 1]|$.

Note that the distribution in the $b = 1$ case is not dependent on k, m. A straightforward hybrid argument shows that RND-CPA-security implies IND-CPA-security, although the reverse implication does not hold[10]. We use the (standard) correctness notion of [20], specialized to the setting of private-key encryption.

Definition 10 (δ-Correctness). *A private-key encryption scheme* $(\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ *is said to be δ-correct if* $\mathbb{E}_{\mathsf{sk} \leftarrow \mathsf{KGen}(1^n)}[\max_{m \in \mathcal{M}}[\Pr[\mathsf{Dec}_{\mathsf{sk}}(c) \neq m \mid c \leftarrow \mathsf{Enc}_{\mathsf{sk}}(m)]]] \leq \delta$.

[10] Take an IND-CPA-secure cryptosystem, and modify encryption to output $\mathsf{Enc}_k(m) || H(k)$ for a hash function $H(\cdot)$, modeled as a random oracle. As k is not consistent between queries to $O_1(\cdot)$, there is a simple RND-CPA distinguisher, but the construction is still IND-CPA-secure.

3 The Encryption Framework

We next present and analyze a secret-key encryption framework. This is done for simplicity of presentation, as the main complication of the public-key setting is a more complex (but, by our results of Section 2.5, still log-concave) noise distribution.

To prove bounds in some framework, one must first

- define a sensible *rate* for the framework, and
- define a *ciphertext error distribution* for the framework.

We do this for our secret-key framework in this section. We additionally show cryptographic security of constructions in our framework, although this is relatively straightforward.

```
$\mathsf{KGen}(1^n)$
    $\mathbf{s} \leftarrow \chi_{\mathsf{sk}}^n$
    return $\mathbf{s}$

$\mathsf{Enc}_{\mathbf{s}}(\mathbf{m})$
    $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$
    $\mathbf{e} \leftarrow \chi_e^m$
    $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} + \mathsf{encode}_E(\mathbf{m})$
    return $(\mathbf{A}, \lfloor \mathbf{b} \rceil_Q)$

$\mathsf{Dec}_{\mathbf{s}}(\mathbf{A}, \mathbf{c})$
    return $\mathsf{decode}_E(\mathbf{c} - \mathbf{A}\mathbf{s})$
```

Fig. 2. Quantized Encryption $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$, defined relative to lattice codes $(E, \lfloor\cdot\rceil_E), (Q, \lfloor\cdot\rceil_Q)$.

Definition 11 (Quantized LWE Encryption). *Let $(E, \lfloor\cdot\rceil_E), (Q, \lfloor\cdot\rceil_Q)$ be lattice codes in $\mathbb{R}_q^m$. Let χ_{sk} be a distribution on $\mathbb{Z}_q$, and let χ_e be a distribution on $\mathbb{R}_q$. The Quantized LWE Encryption Scheme $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ is given by $(\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$, as defined in Figure 2.*

Definition 12. *Let $(E, \lfloor\cdot\rceil_E), (Q, \lfloor\cdot\rceil_Q)$ be lattice codes in $\mathbb{R}_q^m$. Let χ_{sk} be a distribution on $\mathbb{Z}_q$, and χ_e be a distribution on $\mathbb{R}_q$. We say the* asymptotic rate *of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ is the quantity*

$$\frac{\log_2 |E/q\mathbb{Z}^m|}{\log_2 |Q/q\mathbb{Z}^m|} = 1 - \frac{\log_2 \frac{\det E}{\det Q}}{\log_2 \frac{q^m}{\det Q}}.$$

This expression for rate does not include the cost of transmitting $\mathbf{A}$, as there are many ways to reduce (or amortize) this cost, such as appealing to algebraically structured forms of $\mathsf{LWE}_{\chi_{\mathsf{sk}},\chi_e}$, amortizing the cost of $\mathbf{A}$ across many (independent) communication sessions, or transmitting a short seed $s \in \{0,1\}^n$, which one deterministically expands with an extendable output function. In settings where these optimizations are not available (say if one wants to incorporate the cost of transmission of an LWE public key that will be used a single time), one should of course modify the rate to match the particular setting of interest.

We next define the *ciphertext error distribution* of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$. This is the distribution that E must correct for decryption to succeed.

Lemma 4. *Let $(E, \lfloor\cdot\rceil_E), (Q, \lfloor\cdot\rceil_Q)$ be lattice codes in $\mathbb{R}_q^m$. Let χ_{sk} be a distribution on $\mathbb{Z}_q$, and χ_e be a distribution on $\mathbb{R}_q$. Let $\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}$, $\mathbf{e} \leftarrow \chi_e^m$, $\mathbf{s} \leftarrow \chi_{\mathsf{sk}}^n$, and $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathsf{encode}_E(\mathbf{m}) + \mathbf{e}$. Then*

$$\mathsf{Dec}_{\mathbf{s}}(\mathsf{Enc}_{\mathbf{s}}(\mathbf{m})) = \mathbf{m} \iff \mathbf{e} - [\mathbf{b}]_Q \in \mathcal{V}_E.$$

Proof. We have that

$$\begin{aligned}\mathsf{Dec}_{\mathbf{s}}(\mathsf{Enc}_{\mathbf{s}})(\mathbf{m}) &= \mathsf{decode}_E(\lfloor \mathbf{b} \rceil_Q - \mathbf{A}\mathbf{s})\\ &= \mathsf{decode}_E(\mathbf{b} - [\mathbf{b}]_Q - \mathbf{A}\mathbf{s})\\ &= \mathbf{m} + \mathsf{decode}_E(\mathbf{e} - [\mathbf{b}]_Q).\end{aligned}$$

□

In principle the ciphertext error distribution may depend on $\mathbf{m}$. This and other annoyances (namely that $[\mathbf{b}]_Q$ and $\mathbf{e}$ may be dependent) lead us to introduce the following heuristic description of the ciphertext error distribution.

Heuristic 1. Let $(Q, \lfloor \cdot \rceil)$ be a lattice code, $\mathbf{m}$ be any message, $\mathbf{c} \in \mathcal{V}_Q$, $\mathbf{s} \leftarrow \chi_{\mathsf{sk}}^n$, $\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}$, and $\mathbf{e} \leftarrow \chi_e^m$. Then the error $\mathbf{e} - [\mathbf{A}\mathbf{s} + \mathbf{e} + \mathsf{encode}_E(\mathbf{m})]$ is distributed as $\mathbf{e} - \mathbf{u}$, where $\mathbf{u} \leftarrow \mathcal{V}_Q$ is independent from $\mathbf{e}$.

We present a modification of our cryptosystem in Section 3.2 that has the same rate as $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$, which (provably) has the above ciphertext error distribution.

We next derive a bound on δ in terms of the scheme parameters. Curiously, we get a better bound if we first separate-off the (bounded) quantization error and apply a *worst-case* bound over this quantity, rather than naively applying Corollary 2.

Lemma 5. *Let* $(E, \lfloor \cdot \rceil_E), (Q, \lfloor \cdot \rceil_Q)$ *be lattice codes in* $\mathbb{R}_q^m$. *Let* χ_{sk} *be a distribution on* $\mathbb{Z}_q$, *and* χ_e *be a distribution on* $\mathbb{R}_q$. *If* $\Sigma_{\mathbf{e}}$ *is the covariance matrix of* $\mathbf{e} \leftarrow \chi_e^m$, $\mathbf{u} \leftarrow \mathcal{V}_Q$, *then if for some* $p \in \{2, \infty\}$, $r_E^{(p)} > \sqrt{\mathsf{Tr}(\Sigma_{\mathbf{e}})}\left(\tilde{O}(\ln(1/\delta)) + \sqrt{m}\right) + R_Q^{(p)}$, *it follows that* $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ *is* δ*-correct.*

Proof. We have that $\delta = \Pr[\|\mathbf{e}^t - \mathbf{u}^t\|_E > 1] \leq \Pr[\|\mathbf{e}^t\|_E > 1 - \|\mathbf{u}^t\|_E]$. By definition we have that $r_E^{(p)} \cdot \mathcal{B}_m^{(p)} \subseteq \mathcal{V}_E$, and therefore for any $\mathbf{x}$, $\|\mathbf{x}\|_E \leq \frac{1}{r_E^{(p)}}\|\mathbf{x}\|_p$. It follows that $\delta \leq \Pr[\|\mathbf{e}^t\|_p > r_E^{(p)} - R_Q^{(p)}]$. Under the assumed bound on $r_E^{(p)}$, our claim follows by Corollary 2.

□

3.1 Cryptographic Properties of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$

We next establish $\mathsf{RND\text{-}CPA}$ security under the $\mathsf{LWE}_{\chi_{\mathsf{sk}},\chi_e}$ assumption. Note that we require no assumptions[11] on E, Q.

Theorem 3. $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ *is* $\mathsf{RND\text{-}CPA}$*-secure under the* $\mathsf{LWE}_{\chi_{\mathsf{sk}},\chi_e}$ *assumption.*

Proof. Given an adversary that breaks $\mathsf{RND\text{-}CPA}$-security of $\mathsf{LWE}_{\chi_{\mathsf{sk}},\chi_e}$, we describe how to break the decisional $\mathsf{LWE}_{\chi_{\mathsf{sk}},\chi_e}$ assumption. Let $O_b(\cdot)$ be an oracle that either returns samples from (when $b = 0$) $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$, or (when $b = 1$) $(\mathbf{A}, \mathbf{u}) \leftarrow \mathbb{Z}_q^{n\times m} \times \mathbb{R}_q^m$. Construct an encryption oracle that encrypts $\mathbf{m}$ by

[11] Part of this claim is an artifact of us using LWE samples with pseudorandom component $\mathbf{b} \in \mathbb{R}_q^m$. If we replace this with $\mathbb{Z}_q^m$, one can establish security if either $E_q \subseteq \mathbb{Z}_q^m$ or $Q_q \subseteq \mathbb{Z}_q^m$. This is still a relatively minor assumption, as it still implies security for E, Q sharing no common structure.

- sampling $(\mathbf{A}, \mathbf{b}) \leftarrow O_b(\cdot)$, and
- returning $(\mathbf{A}, \lfloor \mathbf{b} + \mathsf{encode}_E(\mathbf{m}) \rceil_Q)$.

When $b = 0$, this is exactly the oracle $O_0(\mathbf{m})$ of the RND-CPA game. When $b = 1$, we will show that it is a random ciphertext. Note that $\mathbf{v} := \mathbf{u} + \mathsf{encode}_E(\mathbf{m})$ is the sum of a uniformly random element $\mathbf{u}$ of a group $\mathbb{R}_q^m$ along with an independent element of that group. By a standard argument analogous to the security of the one-time pad, $\mathbf{v}$ is itself uniform over $\mathbb{R}_q^m$, and independent of $\mathsf{encode}_E(\mathbf{m})$. Finally, for uniform $\mathbf{v}$, it is straightforward to see (as $q\mathbb{Z}^m \subseteq Q$) that $\lfloor \mathbf{v} \rceil_Q$ is uniform, finishing the proof. □

We briefly remark that one could also achieve security of our cryptosystem using a "LWR-type" assumption, namely that $(\mathbf{A}, \lfloor \mathbf{As} \rceil_Q)$ is pseudorandom. This recovers the LWR assumption when Q is a scaling of $\mathbb{Z}^m$.

3.2 Quantized LWE Encryption with a Dither

```
KGen(1^n)
    $\mathbf{s} \leftarrow_\$ \chi_{\mathsf{sk}}^n$
    return $\mathbf{s}$

Enc_s(m)
    $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$
    $\mathbf{e} \leftarrow_\$ \chi_e^m$
    $\mathbf{b} = \mathbf{As} + \mathbf{e} + \mathsf{encode}_E(\mathbf{m})$
    $\mathbf{v} \leftarrow \mathcal{V}_Q$
    return $(\mathbf{A}, \lfloor \mathbf{b} - \mathbf{v} \rceil_Q, \mathbf{v})$

Dec_s(A, c, v)
    return $\mathsf{decode}_E(\mathbf{c} + \mathbf{v} - \mathbf{As})$
```

Fig. 3. Dithered Quantized Encryption $\mathsf{DithLWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$, defined relative to lattice codes $(E, \lfloor \cdot \rceil_E), (Q, \lfloor \cdot \rceil_Q)$. Sampling from $\mathcal{V}_Q$ can be done efficiently via sampling $\mathbf{v} \leftarrow [0, q)^m$, and then computing $[\mathbf{v}]_Q$.

We next describe a variant of quantized LWE for which Heuristic 1 holds. This utilizes what is known as the *subtractive dither* in coding theory, see Chapter 4 of [35] for more details. Security of our construction easily follows under the same conditions (and proof) of Theorem 3. We omit reproducing this proof for brevity, and instead show that the analogue of Heuristic 1 holds for $\mathsf{DithLWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$.

Lemma 6. *Let $(E, \lfloor \cdot \rceil_E), (Q, \lfloor \cdot \rceil_Q)$ be lattice codes in $\mathbb{R}_q^m$. Then the ciphertext error distribution of $\mathsf{DithLWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ satisfies Heuristic 1.*

Proof. For any message $\mathbf{m}$, we can compute that

$$\begin{aligned} \mathsf{Dec}_{\mathbf{s}}(\mathsf{Enc}_{\mathbf{s}}(\mathbf{m})) &= \mathsf{decode}_E(\lfloor \mathbf{b} - \mathbf{v} \rceil_Q + \mathbf{v} - \mathbf{As}) \\ &= \mathsf{decode}_E(\mathbf{b} - \mathbf{v} - [\mathbf{b} - \mathbf{v}]_Q + \mathbf{v} - \mathbf{As}) \\ &= \mathbf{m} + \mathsf{decode}_E(\mathbf{e} - [\mathbf{b} - \mathbf{v}]_Q). \end{aligned}$$

Now, as $\mathbf{v}$ is uniform over $\mathcal{V}_Q$, we have that $[\mathbf{b} - \mathbf{v}]_Q$ is uniform over $\mathcal{V}_Q$ as well, and independent of $\mathbf{b}$ (and therefore $\mathbf{e}$). It follows that $\mathsf{Dec}_{\mathbf{s}}(\mathsf{Enc}_{\mathbf{s}}(\mathbf{m})) = \mathbf{m}$, unless $\mathbf{e} - \mathbf{u} \notin \mathcal{V}_E$, for an independent uniform random variable $\mathbf{u} = [\mathbf{b} - \mathbf{v}]_Q$. □

We next argue that in practice, $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,Q]$ and $\mathsf{DithLWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,Q]$ have the same rate. Recall that we do not explicitly include the random matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ in our computations of rate. One justification for this was that typically, $\mathbf{A}$ itself is not transmitted, and instead a short seed $s \in \{0,1\}^n$ is transmitted, which is then expanded into $\mathbf{A} := H(s)$ using an extendable output function $H(\cdot)$. If this (common) optimization is used, one can simply generate $\mathbf{v}$ in this same manner, so $\mathbf{v}$ does not need to be explicitly included in ciphertexts.

4 Constructions of Quantized LWE Encryption

We next describe the rate achievable by several instantiations (parameterized by lattice codes E, Q) of our framework. The following choice of parameters will be used to enable uniform rate comparisons.

Definition 13. *We say the* standard choice of parameters *are the choice of* $\delta = \exp(-n)$, $\sigma = 2\sqrt{n}$, *and* $m = O(n)$.

4.1 Quantizing Regev's Encryption

We first analyze a quantized variant Regev's initial cryptosystem [32] in our framework, namely $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[(q/p)\mathbb{Z}^m, k\mathbb{Z}^m]$ for $k \in \mathbb{N}$. Regev's initial scheme corresponds to the cryptosystem with no quantization ($k = 1$). We will later optimize over the choice of k to attain a rate $1 - o(1)$ cryptosystem from polynomial modulus.

Definition 14 (Regev Encryption). *Let* $p, q, k \in \mathbb{N}$. *Regev Encryption is the Quantized LWE encryption scheme* $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[(q/p)\mathbb{Z}^m, k\mathbb{Z}^m]$.

Corollary 3. *Let* $p < q$, *and* $k \in \mathbb{N}$. *Then for any* $\delta > 0$, *provided* $\frac{q}{2p} > \tilde{\Omega}(n^{3/2}\sqrt{n + k^2})$, *one can parameterize Regev encryption to be* δ*-correct under the standard choice of parameters and of asymptotic rate at least*

$$1 - O\left(\frac{\log_2(n^2/k)}{\log_2(q/k)}\right).$$

We highlight three main takeaways from this example, namely that

1. for trivial quantization ($k = 1$), it is asymptotic rate $1-\Theta(1)$, i.e. asymptotic rate $1-\Omega(1)$ from polynomial modulus,
2. for non-trivial quantization ($k = \Omega(n^2)$), it is asymptotic rate $1 - o(1)$ from polynomial modulus, and
3. no parameterization (with polynomially-large q) can achieve asymptotic rate better than $1 - o\left(\frac{1}{\log_2 n}\right)$.

Proof. We get by Lemma 5 that this cryptosystem is δ-correct under the standard choice of parameters provided

$$\frac{q}{2p} > \tilde{\Omega}\left(n^{3/2}\sigma\right) + k. \tag{3}$$

Choosing q/p at most a constant-factor larger than this, we get a scheme of asymptotic rate

$$1 - O\left(\frac{\log_2(n^2/k)}{\log_2(q/k)}\right)$$

□

We briefly comment on the tightness of our bounds. Prior analysis of ours (not included in this work) that appealed to Gaussian-specific bounds[12] to optimize Eq. (3) yielded a different bound on q/p, namely the bound

$$\frac{q}{2p} > \sqrt{2}\sigma(\sqrt{\log_2 m} + \sqrt{\ln n}) + k,$$

i.e. with no implicit constants[13], and a bound of $q/2p > \Omega(n)$ rather than $q/2p > \tilde{\Omega}(n^2)$. This yields a scheme of asymptotic rate $1 - O\left(\frac{\log_2(n/k)}{\log_2(q/k)}\right)$. We say this to highlight that the more general log-concave analysis (compared to the Gaussian analysis, only relevant for private-key encryption) does result in *some* loss, but only impacts the three points we highlighted above via requiring a larger parameter $k = \Omega(n^2)$.

4.2 Quantizing the Cryptosystem of [18]

To demonstrate the breadth of our framework, we next show that it contains the high-rate cryptosystems of [18]. This work proposed two high-rate cryptosystems, namely

- **Section 4.1**: an (unquantized) form of what we call Regev encryption, and
- **Section 4.2**: an (unquantized) form of encryption that uses a lattice generated by a“nearly square gadget matrix” H for error-correction.

As we have already analyzed the first construction, we focus on the second construction in this sub-section. [18] constructs the matrix H as the kernel modulo q of an explicit matrix[14] $F' \otimes I_k$, where (for $q = p^\ell - 1$)

$$F' = \begin{pmatrix} p^{\ell-1} & 1 & \dots p^{\ell-2} \\ p^{\ell-2} & p^{\ell-1} & p^{\ell-3} \\ \vdots & \vdots & \ddots \ \vdots \\ 1 & p & \dots p^{\ell-1} \end{pmatrix}.$$

One can verify that F' is precisely what one gets when reducing the collection of $\ell+1$ vectors given by $[\mathbf{g}_p, q\mathbf{e}_1, \dots, q\mathbf{e}_\ell]$ to a basis, i.e. is a basis of the lattice $\Lambda_q(\mathbf{g}_p^t)$. It then follows that the desired matrix H is a basis for the lattice $\Lambda_q^\perp(\mathbf{g}_p^t) \otimes \mathbb{Z}^k$ for some k, as we claimed in Table 1.

[12] To handle the (bounded) uniform component $\mathbf{u}$, we appealed to worst-case bounds on its size.

[13] For this reason, we use this tighter (yet standard) analysis to compute the curves in Figure 1.

[14] The matrix we copy down is actually the transpose of the matrix of [18], as we have different conventions for whether lattices are generated by rows/columns of their basis.

Definition 15 (Gentry-Halevi Encryption, [18]). *For $p, q \in \mathbb{N}$, the Gentry-Halevi Encryption scheme is the Quantized LWE encryption scheme $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[\Lambda_q^\perp(\mathbf{g}_p^t)\otimes\mathbb{Z}^{m/\lceil\log_p q\rceil}, k\mathbb{Z}^m]$.*

Corollary 4. *Let $p < q$, and let $\ell = \lceil \log_p q\rceil$. Assume that $q/p^\ell = O(1)$ with respect to p. Then provided $p > \tilde{\Omega}(n^2)+k$, one can parameterize Gentry-Halevi Encryption to be δ-correct under the standard choice of parameters, and of asymptotic rate at least*

$$1 - O\left(\frac{\log_2(n^2/k)}{\log_2(q/k)}\right).$$

Proof. By Lemma 5, this is δ-correct provided $\frac{q}{p^\ell}\frac{(p-1)}{2} > \sqrt{n}\sigma(\tilde{\Omega}(n)+\sqrt{m})+k$. Under the standard choice of parameters (and assuming $\frac{q}{p^\ell} = O(1)$, independently of p), we get that it suffices to take $p > \tilde{\Omega}(n^2)+k$. This yields a cryptosystem of rate

$$1 - O\left(\frac{\log_2(p/k)}{\log_2(q/k)}\right) = 1 - O\left(\frac{\log_2(n^2/k)}{\log_2(q/k)}\right).$$

□

Note that for large-enough $k = \Omega(n^2)$ this is asymptotic rate $1 - o(1)$ from polynomial modulus, while [18] required super-polynomial modulus to attain rate $1 - o(1)$.

4.3 Optimizing the Quantized Cryptosystem of [6]

We next consider the only cryptosystem in the literature that uses a quantizer that is not of the form $\mathbb{Z}^{m/k}\otimes Q'$, namely the cryptosystem of [6], which the authors of that work refer to as "linearly homomorphic encryption with ciphertext shrinking". We claim this defines exactly the cryptosystem $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[(q/2)\mathbb{Z}^m, \Lambda_{q/2}(\mathbf{u}_m^t)]$. As this equivalence is not obvious, we briefly recall their construction.

The construction starts with an (unquantized) Regev ciphertext $(\mathbf{A}, \mathbf{A}\mathbf{s}+\mathbf{e}+(q/2)m)$. It then shows (existentially) that one can find a scalar $r \in \mathbb{Z}_q$ such that the pair $(\mathbf{w} := \mathsf{decode}_{(q/2)\mathbb{Z}^m}(\mathbf{c}_2 + r\cdot\mathbf{u}_m^t), r) \in \mathbb{Z}_2^m \times \mathbb{Z}_q$ suffice for decryption. We view this pair $(\mathbf{w}, r)$ as defining an element of the lattice $\Lambda_{q/2}(\mathbf{u}_m^t) = (q/2)\mathbb{Z}^m + \mathbb{Z}\cdot\mathbf{u}_m$ via the obvious mapping $(\mathbf{w}, r)\mapsto (q/2)\mathbf{w} + r\cdot\mathbf{u}_m$. Note that this mapping is almost a bijection[15]. Under this identification, the pair $(\mathbf{w}, r)$ is simply equal to $\mathsf{decode}_{\Lambda_{q/2}(\mathbf{u}_m^t)}(\mathbf{c}_2)$ (for a decoding algorithm which need not solve CVP on $\Lambda_{q/2}(\mathbf{u}_m^t)$). If one then attempts to decrypt this ciphertext (using the decryption formula of our work), we have that

$$\begin{aligned}\mathsf{decode}_{(q/2)\mathbb{Z}^m}(\mathsf{encode}_{\Lambda_{q/2}(\mathbf{u}_m^t)}((\mathbf{w}, r)) - \mathbf{A}\mathbf{s}) &= \mathsf{decode}_{(q/2)\mathbb{Z}^m}((q/2)\mathbf{w} + r\cdot\mathbf{u}_m - \mathbf{A}\mathbf{s})\\ &= \mathbf{w} + \mathsf{decode}_{(q/2)\mathbb{Z}^m}(r\cdot\mathbf{u}_m - \mathbf{A}\mathbf{s})\\ &= \mathsf{decode}_{(q/2)\mathbb{Z}^m}(\mathbf{c}_2 + r\cdot\mathbf{u}_m)\\ &- \mathsf{decode}_{(q/2)\mathbb{Z}^m}(\mathbf{A}\mathbf{s} - r\cdot\mathbf{u}_m).\end{aligned}$$

[15] When working modulo q, it is instead a bijection between $\mathbb{Z}_2^m\times\mathbb{Z}_{q/2}$ and our lattice, rather than $\mathbb{Z}_2^m\times\mathbb{Z}_q$ and our lattice. This extra bit in the r component can be removed from [6], i.e. it is not a difference between our schemes. While saving 1 bit does not matter much, for $p \neq 2$ one will save $\log_2 p$ bits, which can start to matter for $p = \omega(1)$.

This is precisely the decryption formula that [6] proposed for their cryptosystem, and therefore their "linearly homomorphic encryption with ciphertext shrinking" is precisely our cryptosystem $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[(q/2)\mathbb{Z}^m, \Lambda_{q/2}(\mathbf{u}^t_m)]$.

We next analyze this construction in our framework, again for a parameterized (by k) family of quantizers that reduces to the cryptosystem of [6] when $k = 1$. The family we choose is given by $k\Lambda_{q/(kp)}(\mathbf{u}^t_m) = (q/p)\mathbb{Z}^m + k\mathbf{u}^t_m \cdot \mathbb{Z}$, i.e. we only sparsify the quantizer in a single dimension (parallel to $\mathbf{u}^t_m$). This yields a *much* smaller (non-asymptotic) improvement. We include this more general analysis so we can refer to it during the conclusion.

Our analysis is done where one decodes with respect to the CVP algorithm (Corollary 1) we have previously derived for this lattice.

Definition 16 (Modified BDGM Encryption). *Let $p, q, k \in \mathbb{N}$. The Modified BDGM Cryptosystem is $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[(q/p)\mathbb{Z}^m, k\Lambda_{q/(kp)}(\mathbf{u}^t_m)]$.*

Corollary 5. *For any $\delta > 0$, let k be such that $kp \mid q$ and $m \mid q/(kp)$. Then one can parameterize the Modified BDGM Cryptosystem under the standard parameters to be δ-correct, and of asymptotic rate at least*

$$1 - O\left(\frac{\log_2(\frac{n^{5/2}}{k})}{m\log_2 p}\right).$$

Proof. Note that by Proposition 3 we have that $R^{(\infty)}_{k\Lambda_{q/(kp)}(\mathbf{u}^t_m)} \leq \frac{q}{2p}\left(1 - \frac{1}{m}\right) + \frac{k}{2}$. By Lemma 5, we have that this cryptosystem is δ-correct provided

$$\frac{q}{2p} > \sqrt{m}\sigma(\tilde{O}(\ln(1/\delta)) + \sqrt{m}) + \frac{q}{2p}\left(1 - \frac{1}{m}\right) + \frac{k}{2}, \tag{4}$$

Under standard parameters, this follows provided $q/p \geq \tilde{\Omega}(n^{5/2}) + kn$. Choosing q/p that is at most a constant factor larger than this, we get (as $\det k\Lambda_{q/(kp)}(\mathbf{u}^t_m) = k(q/p)^{m-1}$) that the asymptotic rate is at least

$$1 - \frac{\log_2 q/kp}{\log_2(q/kp)p^m} \geq 1 - \frac{1}{1 + m\frac{\log_2 p}{\log_2 q/kp}} \geq 1 - O\left(\frac{\log_2(\frac{n^{5/2}}{k})}{m\log_2 p}\right).$$

□

We comment the loss in Eq. (4) (compared to a Gaussian analysis) is smaller for this scheme — we require $q/p = \tilde{\Omega}(n^{5/2}) + kn$ rather than $q/p > \Omega(n^2) + kn$.

4.4 Novel Quantized "Gadget" Encryption

We next describe $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[\Lambda_q(\mathbf{g}^t_p) \otimes \mathbb{Z}^{m/\ell}, \Lambda_{q/p}(\mathbf{u}^t_m)]$, which combines the quantizer of [6] with the (standard) gadget $\Lambda_q(\mathbf{g}^t_p) \otimes \mathbb{Z}^{m/\ell}$. We find this combination has the exact same rate as [6], while still encoding under an error-correcting code that is a gadget, i.e. we combine the relative strengths of both known constructions of high-rate encryption [18,6].

Definition 17. *Let $p, q, k \in \mathbb{N}$. The Quantized Gadget Cryptosystem is* $\mathsf{LWE}_{\chi_{\mathsf{sk}},\chi_e}[\Lambda_q(\mathbf{g}_p^t) \otimes \mathbb{Z}^{m/\ell}, k\Lambda_{q/(kp)}(\mathbf{u}_m^t)]$.

Corollary 6. *For any $\delta > 0$, let k be such that $kp \mid q$ and $m \mid q/(kp)$. Let $q = p^\ell$ for some $\ell > 0$. Then one can parameterize the Quantized Gadget cryptosystem under the standard parameters to be δ-correct, and of asymptotic rate at least*

$$1 - O\left(\frac{\log_2(\frac{n^{5/2}}{k})}{m \log_2 p}\right).$$

Proof. Note that the proof of Corollary 5 only depends on $E = (q/p)\mathbb{Z}^m$ through $\mathcal{V}_E$ and $\det E$, and that by Proposition 1 these quantities are equal for $(q/p)\mathbb{Z}^m$ and $\Lambda_q(\mathbf{g}_p^t) \otimes \mathbb{Z}^{m/\ell}$.

5 Rate Impossibility Results

We next establish rate upper bounds (i.e. impossibility) results in two separate noise models, namely that of perfectly correct encryption (with respect to bounded noise), and that of δ-correct encryption (with respect to log-concave noise).

5.1 Bounded Noise Model

Recall that (with high probability), a Gaussian $\mathbf{e} \leftarrow \chi_e$ concentrates tightly within a ball of radius $\sigma\sqrt{m}$. We first assume that $\|\mathbf{e}\|_2 \leq \sigma\sqrt{m}$ (say by replacing χ_e^m with a Gaussian that is truncated to be contained in this set), and bound the rate of quantized encryption that has $\delta = 0$, i.e. no decryption failures. This setting is amenable to strong packing arguments.

Theorem 4. *Let $(E, \lfloor\cdot\rceil)$ be a lattice code in $\mathbb{R}^m$. Let χ_e be a distribution such that $\mathsf{supp}(\chi_e^m) = \sqrt{m}\sigma\cdot\mathcal{B}_m$. Then, if $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, \mathbb{Z}^m]$ is 0-correct, it has asymptotic rate at most $1 - \Omega\left(\frac{\log_2(\sqrt{m}\sigma)}{\log_2 q}\right)$, i.e. asymptotic rate $1 - o(1)$ encryption from polynomial modulus is impossible.*

Proof. For $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ to be perfectly correct, we need that $\delta = \Pr_{\mathbf{e},\mathbf{b}}[\mathbf{e} - [\mathbf{b}]_Q \notin \mathcal{V}_E] = 0$. As we have that $Q = \mathbb{Z}^m$, we have that $[\mathbf{b}]_Q \in [-1/2, 1/2)^m$, and our condition reduces to $\Pr_{\mathbf{e}}[\mathbf{e} + [\mathbf{b}]_Q \notin \mathcal{V}_E] = 0$, or equivalently $\Pr_{\mathbf{e}}[\mathbf{e} + [-1/2, 1/2)^m \subseteq \mathcal{V}_E] = 1$, i.e. $\mathsf{supp}(\chi_e^m) \subseteq \mathsf{supp}(\chi_e^m) + [-1/2, 1/2)^m = \sqrt{m}\sigma \cdot \mathcal{B}_m + [-1/2, 1/2)^m \subseteq \mathcal{V}_E$.

Now, as $E_q + \mathcal{V}_E = \mathbb{R}_q^m$ is a partition, we have that $E_q + \sqrt{m}\sigma \cdot \mathcal{B}_m \subseteq \mathbb{R}_q^m$ is a packing, meaning the sets $\{e + \sqrt{m}\sigma \cdot \mathcal{B}_m\}_{e \in E_q}$ are disjoint. Taking volumes of both sides, we have that

$$\mathsf{vol}(E_q + \mathsf{supp}(\chi_e^m)) \overset{1}{=} |E_q|\mathsf{vol}(\sqrt{m}\sigma \cdot \mathcal{B}_m) \leq q^m = \mathsf{vol}(\mathbb{R}_q^m),$$

where (1) easily follows from the aforementioned disjointness condition.

Now, we have that $|E_q| = \frac{q^m}{\det E}$. Rearranging, we get that $\det E \geq \mathsf{vol}(\sqrt{m}\sigma\cdot\mathcal{B}_m)$. Stirling's approximation gives that $\mathsf{vol}(\sqrt{m}\sigma \cdot \mathcal{B}_m) \approx \frac{1}{\sqrt{m\pi}}\left(\frac{2\pi e}{m}\right)^{m/2}(\sqrt{m}\sigma)^m$. Finally, we have that the asymptotic rate is

$$R = 1 - \frac{\log_2 \frac{\det E}{\det \mathbb{Z}^m}}{\log_2 \frac{q^m}{\det \mathbb{Z}^m}} = 1 - \frac{\log_2 \det E}{m \log_2 q} \leq 1 - \Omega\left(\frac{\log_2(\sqrt{m}\sigma)}{\log_2 q}\right).$$

□

For our next result, we need the Brunn-Minkowski inequality.

Proposition 7 (Brunn-Minkowski). *Let A, B be non-empty compact subsets of $\mathbb{R}^m$. Then $\sqrt[m]{\mathsf{vol}(A+B)} \geq \sqrt[m]{\mathsf{vol}(A)} + \sqrt[m]{\mathsf{vol}(B)}$.*

Theorem 5. *Let $(E, \lfloor\cdot\rceil_E)$, and $(Q, \lfloor\cdot\rceil_Q)$ be lattice codes in $\mathbb{R}^m$. Let χ_e be a distribution such that $\mathsf{supp}(\chi_e^m) = \sqrt{m}\sigma \cdot \mathcal{B}_m$. Assume that Heuristic 1 holds. Then, if $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ is 0-correct, it has asymptotic rate at most*

$$1 - \frac{\log_2(1 + \frac{\sqrt{2\pi e}\sigma}{\sqrt[m]{\det Q}})}{\log_2 \frac{q}{\sqrt[m]{\det Q}}}.$$

Proof. For $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ to be perfectly correct, we need that $\delta = \Pr_{\mathbf{e},\mathbf{b}}[\mathbf{e} - [\mathbf{b}]_Q \notin \mathcal{V}_E] = 0$. Equivalently, we need that $\Pr_{\mathbf{e},\mathbf{b}}[\mathbf{e} - [\mathbf{b}]_Q \in \mathcal{V}_E] = 1$. Under Heuristic 1, we have that the random variable $\mathbf{e} - [\mathbf{b}]_Q$ has support $\mathsf{supp}(\chi_e^m) + (-\mathcal{V}_Q)$. Note that $\mathcal{V}_Q$ is centrally symmetric, so $-\mathcal{V}_Q = \mathcal{V}_Q$. We therefore have that $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ is 0-correct if and only if $\sqrt{m}\sigma \cdot \mathcal{B}_m + \mathcal{V}_Q \subseteq \mathcal{V}_E$.

Now, as $E_q + \mathcal{V}_E = \mathbb{R}_q^m$ is a partition, we have that $E_q + (\sqrt{m}\sigma \cdot \mathcal{B}_m + \mathcal{V}_Q) \subseteq \mathbb{R}_q^m$ is a packing, i.e. the sets $\{e + (\sqrt{m}\sigma \cdot \mathcal{B}_m + \mathcal{V}_Q)\}_{e \in E_q}$ are disjoint. Taking volumes, we have that

$$\mathsf{vol}(E_q + (\sqrt{m}\sigma \cdot \mathcal{B}_m + \mathcal{V}_Q)) = |E_q| \,\mathsf{vol}(\sqrt{m}\sigma \cdot \mathcal{B}_m + \mathcal{V}_Q) \leq q^m = \mathsf{vol}(\mathbb{R}_q^m).$$

As $|E_q| = \frac{q^m}{\det E}$, this inequality is equivalent to $\sqrt[m]{\det E} \geq \sqrt[m]{\mathsf{vol}(\sqrt{m}\sigma \cdot \mathcal{B}_m + \mathcal{V}_Q)}$. Applying the Brunn-Minkowski inequality and Stirling's Approximation, we get that

$$\sqrt[m]{\det E} \geq \sqrt{2\pi e}\sigma + \sqrt[m]{\det Q}.$$

This immediately implies that the asymptotic rate is

$$R = 1 - \frac{\log_2 \frac{\det E}{\det Q}}{\log_2 \frac{q^m}{\det Q}} = 1 - \frac{\log_2(1 + \frac{\sqrt{2\pi e}\sigma}{\sqrt[m]{\det Q}})}{\log_2 \frac{q}{\sqrt[m]{\det Q}}}.$$

□

Note that the upper bound becomes $1 - o\left(\frac{1}{\log_2 \frac{q}{\sigma}}\right)$ if $\sqrt[m]{\det Q} \approx \sigma$, i.e. rate $1 - o(1)$ encryption is no longer impossible provided one quantizes even a relatively small amount.

5.2 Results for Unbounded Errors

We next return to the setting of χ_e an arbitrary log-concave distribution, and bounding δ-correct encryption for $\delta > 0$. Here, we rely on the anti-concentration inequality of Proposition 6, rather than the prior packing arguments. We first give a bound that is mostly useful in the case of trivial quantization, i.e. where $Q = \mathbb{Z}^m$.

Theorem 6. *Let $\epsilon > 0$. Let $(E, \lfloor\cdot\rceil_E)$, $(Q, \lfloor\cdot\rceil_Q)$ be any lattice codes in $\mathbb{R}^m$. Let the ciphertext error distribution has covariance matrix Σ. If $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ is δ-correct, then the asymptotic rate of $\mathsf{LWE}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ is at most*

$$1 - \frac{\log_2 \Omega\left(\frac{\sqrt{\mathsf{Tr}(\Sigma)}}{\overline{R}_E}\right)}{\log_2 q/\sqrt[m]{\det Q}} + o(1).$$

Proof. We have that

$$1 - \delta \leq \Pr_{\mathbf{e}}[\|\mathbf{e}\|_E \leq 1] \leq \Pr_{\mathbf{e}}[\|\mathbf{e}\|_2 \leq R_E] \leq O\left(\frac{R_E}{\sqrt{\mathsf{Tr}(\Sigma)}}\right).$$

The first inequality is from Lemma 1, and the second from Proposition 6. We then easily get the bound $\sqrt[m]{\det E} \geq \Omega\left(\frac{1-\delta}{\overline{R}_E}\sqrt{\mathsf{Tr}(\Sigma)}\right)$, and the asymptotic rate is

$$R = 1 - \frac{\log_2 \sqrt[m]{\det E}}{\log_2 q/\sqrt[m]{\det Q}} \leq 1 - \frac{\log_2 \Omega\left(\frac{\sqrt{\mathsf{Tr}(\Sigma)}(1-\delta)}{\overline{R}_E}\right)}{\log_2 q/\sqrt[m]{\det Q}},$$

Finally, we separate off the $1 - \delta$ term, and note that $-\log_2(1-\delta)/\log_2 q$ is easily $o(1)$ to get the claimed result. □

The presence of $\overline{R}_E$ in this bound is peculiar, and we cannot remove it by appealing to a universal upper bound on $\overline{R}_E$ (no such bound exists, even if we restrict $\mathcal{V}_E$ to be the Voronoi cell of a lattice). If we assume $\overline{R}_E$ is not too large (either absolutely, or in comparison to $\overline{r}_E$), we can prove impossibility of rate $1 - o(1)$ encryption.

Corollary 7. *Let $\epsilon > 0$, and let $(E, \lfloor\cdot\rceil_E)$ be a lattice code in $\mathbb{R}^m$. If either*

- *$\overline{R}_E \leq O(m^{1-\epsilon})$, or*
- *$\overline{R}_E/\overline{r}_E \leq O(m^{1/2-\epsilon})$,*

and q is polynomially large, then $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, \mathbb{Z}^m]$ is of rate $1 - \Omega(1)$, i.e. under these conditions rate $1 - o(1)$ encryption is impossible.

Proof. We show that the second condition implies the first. This is simple, as the bound $\overline{r}_E \leq O(m^{1/2})$ implies that $\overline{R}_E \leq O(\overline{r}_E m^{1/2-\epsilon}) \leq O(m^{1-\epsilon})$. Next, note that by Theorem 6, we have that the asymptotic rate is at most

$$1 - \frac{\log_2 \Omega\left(\frac{\sqrt{m}\sigma}{m^{1-\epsilon}}\right)}{\log_2 q} + o(1) = 1 - \epsilon\frac{\log_2 \Omega(m)}{\log_2 q} - \frac{\log_2 \frac{\sigma}{\sqrt{m}}}{\log_2 q} + o(1).$$

As q is polynomially large, this suffices for the claimed result. □

Corollary 8. *There exist lattice codes E with $\overline{r}_E \geq \Omega(\sqrt{m})$, i.e. within a constant factor of optimal, such that $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,\mathbb{Z}^m]$ is of rate $1-\Omega(1)$.*

Proof. Choose E with $\frac{\overline{R}_E}{\overline{r}_E} \leq 2+o(1)$, which are known to exist [9], and then apply Corollary 7. □

Therefore, any result establishing rate $1-o(1)$ encryption from $Q=\mathbb{Z}^m$ and $q=n^{O(1)}$ must do more than simply appeal to the packing radius $\overline{r}_E = \Theta(\sqrt{m})$ being nearly optimal.

We next extend our bound on $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,Q]$ for ciphertext error distribution the sum of a log-concave random variable and $\mathbf{u} \leftarrow \mathcal{V}_Q$ uniform, in a similar way to how we got sharper upper bounds on δ by considering this special case.

Theorem 7. *Let $(E,\lfloor\cdot\rceil_E),(Q,\lfloor\cdot\rceil_Q)$ be lattice codes in $\mathbb{R}^m$, and assume that $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,Q]$ is δ-correct. Assume that Heuristic 1 holds, i.e. one can write the ciphertext error distribution as the independent sum of a log-concave random variable (with covariance matrix Σ) and $\mathbf{u} \leftarrow \mathcal{V}_Q$. Then $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,Q]$ is of asymptotic rate at most*

$$1-\frac{\log_2 \Omega\left(\frac{\sqrt{m}\sigma}{R_Q}\right)}{m\log_2 \frac{q}{\sqrt[m]{\det Q}}}+o(1).$$

Proof. Throughout, let $p(x)$ be the density of the log-concave random variable $\mathbf{e}$. By the law of total probability, we have that

$$\begin{aligned}
\Pr[\mathbf{e}-\mathbf{u}\in\mathcal{V}_E] &= \frac{1}{\det Q}\int_{\mathcal{V}_Q}\int_{\mathcal{V}_E} p(\mathbf{e}-\mathbf{x})d\mathbf{e}d\mathbf{x}\\
&\leq \frac{1}{\det Q}\int_{\mathcal{V}_E}\Pr[\|\mathbf{e}-\mathbf{x}\|_2 \leq R_Q]d\mathbf{e}\\
&\leq O\left(\frac{\det E}{\det Q}\frac{R_Q}{\sqrt{\mathsf{Tr}(\Sigma)}}\right),
\end{aligned}$$

where the first inequality is the containment $\mathcal{V}_Q \subseteq R_Q\cdot\mathcal{B}_n$ (as well as Fubini's theorem), and the second inequality is Proposition 6. It follows that the asymptotic rate is

$$1-\frac{\log_2\Omega\left(\frac{\mathsf{Tr}(\Sigma)}{R_Q}\right)}{\log_2|Q/q\mathbb{Z}^m|}-\frac{\log_2(1-\delta)}{\log_2|Q/q\mathbb{Z}^m|}.$$

We finish by applying the same bound to $1-\delta$ as we did in Theorem 6. □

Corollary 9. *Let $(E,\lfloor\cdot\rceil_E),(Q,\lfloor\cdot\rceil_Q)$ be lattice codes in $\mathbb{R}^m$, and let $\epsilon > 0$. Assume the validity of Heuristic 1. If $\overline{R}_Q \leq O(\sqrt{m})$ is within a constant factor of optimal, then the asymptotic rate of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E,Q]$ is at most*

$$1-\frac{\log_2\Omega\left(\frac{\sigma}{\sqrt[m]{\det Q}}\right)}{m\log_2\frac{q}{\sqrt[m]{\det Q}}}+o(1).$$

In particular, if $\sqrt[m]{\det Q} \leq O(\sigma)$, this quantity is at most $1-\Omega\left(\frac{1}{m\log_2\frac{q}{\sigma}}\right)+o(1)$.

Proof. This follows directly from plugging the bounds we assume into Theorem 7. □

Note that, as our modification of BDGM encryption (Corollary 5) and the Quantized Gadget cryptosystem (Corollary 6) have rate $1 - O\left(\frac{1}{m}\right)$, under the standard choice of parameters this bound is tight up to an $O(\log_2 m)$ factor for quantizers with $\sqrt[m]{\det Q} \leq O(\sigma)$.

5.3 Exponentially Stronger Bounds Against a Common Design Paradigm

We finish by showing that the bound Corollary 9 can be significantly strengthened when restricting to $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ where $E = \mathbb{Z}^{m/\dim E'} \otimes E'$, $Q = \mathbb{Z}^{m/\dim E'} \otimes Q'$ are the direct sum of m/k identical (smaller) codes for $k = \dim E' = \dim Q'$. In what follows we *solely* change the dimension, and keep the other parameters q, δ, σ, n fixed.

Lemma 7. *Let $E = \mathbb{Z}^{m/k} \otimes E'$, and $Q = \mathbb{Z}^{m/k} \otimes Q'$, where E', Q' are k-dimensional lattice codes. Then the asymptotic rate of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E, Q]$ is equal to the asymptotic rate of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E', Q']$.*

Proof. Note that $\det E = (\det E')^{m/k}$, and similarly for $\det Q$. We then have that the asymptotic rate is

$$\frac{\log_2 \frac{q^m}{\det E}}{\log_2 \frac{q^m}{\det Q}} = \frac{\log_2 \frac{q^m}{(\det E')^{m/k}}}{\log_2 \frac{q^m}{(\det Q')^{m/k}}} = \frac{\log_2 \frac{q^k}{\det E'}}{\log_2 \frac{q^k}{\det Q'}}.$$

□

Corollary 10. *Let $(E', \lfloor\cdot\rceil_{E'}), (Q', \lfloor\cdot\rceil_{Q'})$ be lattice codes in $\mathbb{R}^k$, let $k \mid m$, and let $\epsilon > 0$. Assume the validity of Heuristic 1. If $\overline{R}_{Q'} \leq O(\sqrt{k})$ is within a constant factor of optimal, then the asymptotic rate of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[\mathbb{Z}^{m/k} \otimes E', \mathbb{Z}^{m/k} \otimes Q']$ is at most*

$$1 - \frac{\log_2 \Omega\left(\frac{\sigma}{\sqrt[k]{\det Q}}\right)}{k \log_2 \frac{q}{\sqrt[k]{\det Q}}} + o(1).$$

In particular, if $\sqrt[k]{\det Q'} \leq O(\sigma)$, this quantity is at most $1 - \Omega\left(\frac{1}{k \log_2 \frac{q}{\sigma}}\right) + o(1)$.

Proof. Use Corollary 9 to bound the rate of $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[E', Q']$. By Lemma 7, this implies the same bound for $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[\mathbb{Z}^{m/k} \otimes E', \mathbb{Z}^{m/k} \otimes Q']$. □

Note that in the literature, k is typically at most $O(\log_2 m)$, so this bound is exponentially stronger than Corollary 9 in this common setting.

6 Conclusion and Open Problems

Conclusion We propose a framework that reduces the design of LWE-based encryption to a handful of coding-theoretic choices. We then prove bounds on any instantiation of this framework, and find that a preexisting cryptosystem in the literature [6] is within an $O(\log_2 m)$ factor of optimal rate. We additionally prove bounds against the common situation of building lattices for error-correction and quantization by setting $L = \bigoplus_{i=1}^{m/\log_2 m} L'$ for $\dim L' = \Theta(\log_2 m)$. We establish exponentially stronger bounds against this setting, which we validate via practical rate computations.

Open Problems We find an $O(\log_2 m)$ gap between the best-known construction and our bound for any construction. This gap is surprisingly significant — if there exists a construction meeting our bound, it implies constant (independent of the amount of data to transmit) overhead lattice-based encryption, i.e. a lattice-based cryptosystem that is similar to (standard) hybrid encryption. Does such a cryptosystem exist, or can one establish the impossibility of such a construction? Note that our cryptosystem $\mathsf{LWE}^{n,q}_{\chi_{\mathsf{sk}},\chi_e}[(q/p)\mathbb{Z}^m, k\Lambda_{q/(kp)}(\mathbf{u}^t_m)]$ gets quite close. If we did not have the divisibility requirement $m \mid q/(kp)$, it would suffice to close the gap itself. Can this requirement be removed? Finally, our work suggests the quantizer $\Lambda_{q/p}(\mathbf{u}^t_m)$ is much better than $k\mathbb{Z}^m$, which is implicitly used to define the LWR assumption. Can one obtain secure and practical LWR-type constructions using this quantizer?

References

1. M. Albrecht, M. Chase, H. Chen, J. Ding, S. Goldwasser, S. Gorbunov, S. Halevi, J. Hoffstein, K. Laine, K. Lauter, S. Lokam, D. Micciancio, D. Moody, T. Morrison, A. Sahai, and V. Vaikuntanathan. Homomorphic encryption security standard. Technical report, HomomorphicEncryption.org, Toronto, Canada, November 2018.
2. M. R. Albrecht, L. Grassi, C. Rechberger, A. Roy, and T. Tiessen. MiMC: Efficient encryption and cryptographic hashing with minimal multiplicative complexity. In J. H. Cheon and T. Takagi, editors, *Advances in Cryptology – ASIACRYPT 2016, Part I*, volume 10031 of *Lecture Notes in Computer Science*, pages 191–219, Hanoi, Vietnam, Dec. 4–8, 2016. Springer, Heidelberg, Germany.
3. T. Ashur, M. Mahzoun, and D. Toprakhisar. Chaghri - A FHE-friendly block cipher. In H. Yin, A. Stavrou, C. Cremers, and E. Shi, editors, *ACM CCS 2022: 29th Conference on Computer and Communications Security*, pages 139–150, Los Angeles, CA, USA, Nov. 7–11, 2022. ACM Press.
4. L. Babai. On lovász'lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.
5. A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In Pointcheval and Johansson [30], pages 719–737.
6. Z. Brakerski, N. Döttling, S. Garg, and G. Malavolta. Leveraging linear decryption: Rate-1 fully-homomorphic encryption and time-lock puzzles. In Hofheinz and Rosen [21], pages 407–437.
7. Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In S. Goldwasser, editor, *ITCS 2012: 3rd Innovations in Theoretical Computer Science*, pages 309–325, Cambridge, MA, USA, Jan. 8–10, 2012. Association for Computing Machinery.
8. Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In R. Ostrovsky, editor, *52nd Annual Symposium on Foundations of Computer Science*, pages 97–106, Palm Springs, CA, USA, Oct. 22–25, 2011. IEEE Computer Society Press.
9. G. Butler. Simultaneous packing and covering in euclidean space. *Proceedings of the London Mathematical Society*, 3(4):721–735, 1972.
10. A. Carbery and J. Wright. Distributional and l-q norm inequalities for polynomials over convex bodies in r-n. *Mathematical Research Letters*, 8:233–248, 2001.
11. J. H. Conway and N. J. A. Sloane. *Sphere Packings, Lattices and Groups*, volume 290 of *Grundlehren der mathematischen Wissenschaften.* Springer New York, New York, NY, 1999.
12. J.-P. D'Anvers, A. Karmakar, S. S. Roy, and F. Vercauteren. Saber: Module-LWR based key exchange, CPA-secure encryption and CCA-secure KEM. In A. Joux, A. Nitaj, and T. Rachidi, editors, *AFRICACRYPT 18: 10th International Conference on Cryptology in Africa*, volume 10831 of *Lecture Notes in Computer Science*, pages 282–305, Marrakesh, Morocco, May 7–9, 2018. Springer, Heidelberg, Germany.
13. H. Davenport. The covering of space by spheres. *Rendiconti del Circolo Matematico di Palermo*, 1(1):92–107, Jan. 1952.
14. L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé. CRYSTALS-Dilithium: A lattice-based digital signature scheme. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2018(1):238–268, 2018. `https://tches.iacr.org/index.php/TCHES/article/view/839`.
15. L. Ducas and W. P. van Woerden. The closest vector problem in tensored root lattices of type a and in their duals. *Designs, Codes and Cryptography*, 86:137–150, 2018.
16. R. E. Gaunt. The basic distributional theory for the product of zero mean correlated normal random variables. *Statistica Neerlandica*, 2022.

17. N. Genise, D. Micciancio, and Y. Polyakov. Building an efficient lattice gadget toolkit: Subgaussian sampling and more. In Y. Ishai and V. Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part II*, volume 11477 of *Lecture Notes in Computer Science*, pages 655–684, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany.
18. C. Gentry and S. Halevi. Compressible FHE with applications to PIR. In Hofheinz and Rosen [21], pages 438–464.
19. S. Guo, P. Kamath, A. Rosen, and K. Sotiraki. Limits on the efficiency of (ring) LWE-based non-interactive key exchange. *Journal of Cryptology*, 35(1):1, Jan. 2022.
20. D. Hofheinz, K. Hövelmanns, and E. Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In Y. Kalai and L. Reyzin, editors, *TCC 2017: 15th Theory of Cryptography Conference, Part I*, volume 10677 of *Lecture Notes in Computer Science*, pages 341–371, Baltimore, MD, USA, Nov. 12–15, 2017. Springer, Heidelberg, Germany.
21. D. Hofheinz and A. Rosen, editors. *TCC 2019: 17th Theory of Cryptography Conference, Part II*, volume 11892 of *Lecture Notes in Computer Science*, Nuremberg, Germany, Dec. 1–5, 2019. Springer, Heidelberg, Germany.
22. A. Jambulapati, Y. T. Lee, and S. S. Vempala. A slightly improved bound for the kls constant. *arXiv preprint arXiv:2208.11644*, 2022.
23. Z. Jin and Y. Zhao. Generic and practical key establishment from lattice. In R. H. Deng, V. Gauthier-Umaña, M. Ochoa, and M. Yung, editors, *ACNS 19: 17th International Conference on Applied Cryptography and Network Security*, volume 11464 of *Lecture Notes in Computer Science*, pages 302–322, Bogota, Colombia, June 5–7, 2019. Springer, Heidelberg, Germany.
24. Y. T. Lee and S. S. Vempala. The kannan–lovász–simonovits conjecture. *Current Developments in Mathematics*, 2017(1):1–36, 2017.
25. J. Martinet. *Perfect Lattices in Euclidean Spaces*, volume 327 of *Grundlehren der mathematischen Wissenschaften.* Springer Berlin Heidelberg, Berlin, Heidelberg, 2003.
26. R. G. McKilliam, W. D. Smith, and I. V. L. Clarkson. Linear-Time Nearest Point Algorithms for Coxeter Lattices. *IEEE Transactions on Information Theory*, 56(3):1015–1022, Mar. 2010.
27. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In Pointcheval and Johansson [30], pages 700–718.
28. D. Micciancio and Y. Polyakov. Bootstrapping in fhew-like cryptosystems. In *Proceedings of the 9th on Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, pages 17–28, 2021.
29. C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In M. Mitzenmacher, editor, *41st Annual ACM Symposium on Theory of Computing*, pages 333–342, Bethesda, MD, USA, May 31 – June 2, 2009. ACM Press.
30. D. Pointcheval and T. Johansson, editors. *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, Cambridge, UK, Apr. 15–19, 2012. Springer, Heidelberg, Germany.
31. A. v. Poppelen. Cryptographic Decoding of the Leech Lattice. Master's thesis, Utrecht University, 2016. https://studenttheses.uu.nl/handle/20.500.12932/24606.
32. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In H. N. Gabow and R. Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 84–93, Baltimore, MA, USA, May 22–24, 2005. ACM Press.
33. C. Saliba, L. Luzzi, and C. Ling. A reconciliation approach to key generation based on module-lwe. In *2021 IEEE International Symposium on Information Theory (ISIT)*, pages 1636–1641, 2021.
34. A. Saumard and J. A. Wellner. Log-concavity and strong log-concavity: a review. *Statistics surveys*, 8:45, 2014.
35. R. Zamir, B. Nazer, Y. Kochman, and I. Bistritz. *Lattice Coding for Signals and Networks: A Structured Coding Approach to Quantization, Modulation and Multiuser Information Theory.* Cambridge University Press, 2014.