

Multi-ciphertext security degradation for lattices

Daniel J. Bernstein^{1,2}

¹ Department of Computer Science, University of Illinois at Chicago, USA

² Horst Görtz Institute for IT Security, Ruhr University Bochum, Germany

djb@cr.yp.to

Abstract. Typical lattice-based cryptosystems are commonly believed to resist multi-target attacks. For example, the New Hope proposal stated that it avoids “all-for-the-price-of-one attacks”. An ACM CCS 2021 paper from Duman–Hövelmanns–Kiltz–Lyubashevsky–Seiler stated that “we can show that $\text{Adv}_{\text{PKE}}^{\text{IND-CPA}} \approx \text{Adv}_{\text{PKE}}^{(n,q_C)\text{-IND-CPA}}$ ” for “lattice-based schemes” such as Kyber, i.e. that one-out-of-many-target IND-CPA is as difficult to break as single-target IND-CPA, assuming “the hardness of MLWE as originally defined for the purpose of worst-case to average-case reductions”. Meanwhile NIST expressed concern regarding multi-target attacks against non-lattice cryptosystems.

This paper quantifies the asymptotic impact of multiple ciphertexts per public key upon existing heuristic analyses of known lattice attacks. The qualitative conclusions are that typical lattice PKEs asymptotically degrade in heuristic multi-ciphertext IND-CPA security as the number of ciphertexts increases. These PKE attacks also imply multi-ciphertext IND-CCA2 attacks against typical constructions of lattice KEMs. This shows a contradiction between (1) the existing heuristics and (2) the idea that multi-target security matches single-target security.

The asymptotic heuristic security degradation is exponential in $\Theta(n)$ for decrypting many ciphertexts, cutting a constant fraction out of the total number of bits of security, and exponential in $\Theta(n/\log n)$ for decrypting one out of many ciphertexts, for conservative cryptosystem parameters. Furthermore, whether or not the existing heuristics are correct, (1) there are flaws in the claim of provable multi-target security based on MLWE, and (2) there is a 2^{88} -guess attack breaking one out of 2^{40} ciphertexts for a FrodoKEM-640 public key.

Keywords: algorithm analysis, multi-target attacks, lattices

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of the Excellence Strategy of the German Federal and State Governments—EXC 2092 CASA—390781972 “Cyber Security in the Age of Large-Scale Adversaries”; by the U.S. National Science Foundation under grant 1913167; by the Taiwan’s Executive Yuan Data Safety and Talent Cultivation Project (AS-KPQ-109-DSTCP); and by the Cisco University Research Program. “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” (or other funding agencies). Permanent ID of this document: cff870df5aeaaa8a4fd37778bac69658864897ef. Date: 2022.11.14.

1 Introduction

Given $(\text{AES}_k(0), \text{AES}_k(1))$ for an AES-128 key k , brute-force search finds k with at most 2^{128} AES-128 computations, about 2^{143} bit operations. Given $(\text{AES}_k(0), \text{AES}_k(1))$ for each k in a list of 2^{40} AES-128 keys, a batch version of brute-force search finds the entire list with at most 2^{128} AES-128 computations. The obvious batch attack is bottlenecked by memory access rather than computation, but the Hellman–Rivest–van Oorschot–Wiener parallel batch attack uses essentially the same amount of computation and much less communication; see generally [24].

These batch algorithms will find *one of the keys in the list* with only about 2^{103} bit operations. To put this in perspective, Bitcoin currently carries out roughly 2^{111} bit operations per year. The original 2^{143} bit operations sounded comfortably out of reach, but 2^{103} is certainly feasible for a large-scale attacker.

After these 2^{103} bit operations, the attacker has acquired a secret key that the attacker was not supposed to have, perhaps exposing other critical information—e.g., a password that can be used to launch further attacks. Any particular user is likely to be safe, but the full multi-user cryptographic system has been broken.

There are more examples in the literature of cryptographic systems for which known multi-target attacks are much more efficient than state-of-the-art single-target attacks. For example, state-of-the-art non-quantum algorithms to compute discrete logarithms in the group $(\mathbb{Z}/p)^*$, where p is a prime chosen in the usual way, can compute many discrete logarithms for only slightly higher total cost; see, e.g., [2]. More subtly, work can be shared across many primes p ; see [20, Chapter 7]. Finding *one of many discrete logarithms* for a single p is as expensive as attacking a single prespecified target, but the batch attacks have lower per-target cost and undermine a common argument for taking small p ; see [25, “difficulties for the ‘attacker economist’ philosophy”].

Given the tremendous diversity of ways that multi-target security has been shown to fail, some cryptosystem proposals presume that there will be $\lg T$ bits of security loss from T -target attacks—e.g., 40 bits of security loss from 2^{40} -target attacks—unless proven otherwise. See, e.g., [5, Section 8.3]. However, one can easily find many other cryptosystem proposals that choose low single-target security levels, implicitly or explicitly presuming that there is *no* multi-target security loss. For those systems, multi-target attacks can easily make the difference between being safe and being broken in the real world.

1.1. Bleeding-edge lattice systems. The current version of Kyber-512 was proposed in [17] in October 2020. The same document presented a new security analysis [17, Section 5] and concluded that “this preliminary analysis gives a cost of 2^{151} gates” to break Kyber-512. This might sound safely beyond 2^{143} , but there are at least three reasons for concern:

- The same document [17, page 26] said that “this number could be affected by a factor of up to 2^{16} in either direction”—i.e., possibly as low as 2^{135} bit operations—because of “known unknowns”.

- Subsequent work has added knowledge regarding the “known unknowns” but has also found exploitable gaps in the analysis in [17]. See, e.g., the attacks in [53] and [77]. It is not clear what security level is claimed for Kyber-512 after these attacks.
- These are single-target security analyses, leaving open the possibility of much more efficient multi-target attacks. Even if there are no further single-target improvements, why should users believe that multi-target attacks against Kyber-512 are not already feasible with current technology?

Regarding the third point, the literature quoted in Section 5 gives the impression that typical lattice systems such as Kyber are safe from all-for-the-price-of-one attacks, one-out-of-many attacks, and other types of multi-target attacks. For example, an ACM CCS 2021 paper [47, page 3] stated that “we can show that $\text{Adv}_{\text{PKE}}^{\text{IND-CPA}} \approx \text{Adv}_{\text{PKE}}^{(n,q_C)\text{-IND-CPA}}$,” for “lattice-based schemes”, meaning that single-target and one-out-of-many IND-CPA security are approximately the same for these PKEs. This statement relied on a hypothesis, but [47] argued plausibility of the hypothesis.

1.2. Contributions of this paper. This paper has five main contributions:

- Asymptotic analysis and optimization of parameters for well-known single-target lattice attacks, *assuming* accuracy of the existing heuristics regarding those attacks. See Section 2.
- Asymptotic analysis of an attack breaking many ciphertexts for one public key, *assuming* the existing heuristics. This attack has total heuristic cost asymptotically similar to the single-target message-recovery attacks, i.e., much lower heuristic cost per ciphertext. This does not rely on key recovery: it applies even if the cryptosystem is modified to put extra defenses around the secret key. See Section 3.
- Asymptotic analysis of an attack breaking one out of many ciphertexts. *assuming* the existing heuristics. For a wide range of cryptosystem parameters, this attack has heuristic cost beating the usual single-target message-recovery attack by a factor $2^{\Theta(n/\lg n)}$, contrary to the statement “ $\text{Adv}_{\text{PKE}}^{\text{IND-CPA}} \approx \text{Adv}_{\text{PKE}}^{(n,q_C)\text{-IND-CPA}}$ ”. See Section 4.
- Identification of logical gaps in how the previous literature arrived at various statements contrary to this paper’s calculations. See Section 5.
- A particularly easy attack specifically against FrodoKEM, disproving the claim in [12] that “the FrodoKEM parameter sets comfortably match their target security levels with a large margin”. See Section 6.

The gap identification and the FrodoKEM attack apply even if the existing heuristics are so inaccurate as to undermine the other analyses.

1.3. Caveats. In the literature on algorithms, papers introducing asymptotic speedups are often followed by papers doing more work to show corresponding concrete speedups. Sometimes, however, concrete input sizes of interest turn out to be too small for the asymptotic speedups to apply, even after further optimization effort. See, e.g., [52].

Establishing Kyber-512’s *single-target* security level remains a challenging research problem today; see, e.g., the open problems listed in [77]. Establishing Kyber-512’s *T-target* security level is even more challenging. Perhaps Kyber-512 is too small for the heuristic multi-target speedups to be applicable.

Furthermore, relying on heuristics is error-prone, as illustrated by previous lattice-attack literature endorsing application of the same heuristics to *S*-unit lattices, where the heuristics were later shown to be highly inaccurate; see [32]. Typical techniques to detect errors in analyses, such as checking proofs and carrying out experiments, are inherently unreliable for heuristic asymptotic analyses: saying that the analyses are heuristic implies that the results are not completely proven, and saying that the analyses are asymptotic implies that discrepancies from experiments are to be expected.

Sections 2, 3, and 4 of this paper are saying that two ideas visible in the literature—the idea that the existing heuristics accurately predict lattice security via the standard analysis, and the idea that multi-target lattice security matches single-target lattice security—cannot both be correct. This paper is *not* saying that one idea is accurate and the other idea is inaccurate; it seems more likely that there are serious inaccuracies in *both* ideas.

Yet another caveat is that, like most of the literature on lattice attacks, this paper works in an unrealistic model of computation where arbitrarily large arrays can be accessed for free. The multi-ciphertext attack analyzed here is not exceptionally memory-intensive—it handles one ciphertext at a time, and, for each ciphertext, uses memory in similar ways to the usual single-target attacks—but the issue is that all of these attacks could be outperformed by lower-memory approaches in more realistic models of computation.

Finally, the importance of multi-target attacks rests on the cryptographic community’s habit of taking narrow security margins. See generally [16] and [28]. This habit is not a rule; in particular, at the time of this writing, NIST has not committed to standardizing Kyber-512.³ However, NIST’s latest report claims that Kyber-512 is harder to break than single-target AES-128,⁴ and there are ongoing large-scale deployment experiments with Kyber-512, such as [39].

³ For example, NIST’s announcement [83] of the Kyber selection includes a statement that “NIST will seek input on specific parameter sets to include” in the resulting standard. Furthermore, NIST’s latest report [4, page 18], referring to patent agreements, says “If the agreements are not executed by the end of 2022, NIST may consider selecting NTRU instead of Kyber”.

⁴ [4, page 8] says “Figure 1 shows [performance] for Kyber, NTRU, and Saber for security categories 1 and 3”. The Kyber parameter sets in [4, Figure 1] are Kyber-512 and Kyber-768, evidently assigning Kyber-512 to “category 1” and Kyber-768 to “category 3”. [4, page 6] says “category 1” means “the best attack violating the security definition of a parameter set should cost more than a brute-force key search attack on a single instance of AES-128, according to any plausible assumption regarding the relative cost of the various computational resources involved in a real-world attack”.

2 Asymptotic heuristic cost of single-target attacks

The starting point for this paper is the standard heuristic analysis of the usual single-target message-recovery and key-recovery attacks against lattice PKEs. This analysis was introduced in the original New Hope paper [13, Section 6.3] in 2015, building on various earlier components, and was then copied into many newer cryptosystem proposals. Recent literature continues to report results of the same analysis; see, e.g., [12, Section 5.2.2]. There have been subsequent speedups and corrections, but the literature generally portrays these adjustments as minor; see, e.g., the recent paper [7], starting with the “history of refinements” title.

For definiteness, consider the Lyubashevsky–Peikert–Regev PKE [75, eprint version, page 4], using the ring $(\mathbb{Z}/q)[x]/(x^n+1)$ and sampling each error position from distribution χ . The standard heuristic analysis cares only about (n, q, s) , where s is the standard deviation of χ . As illustrated by the tables in [6], the same heuristic analysis also applies to a wide range of further cryptosystems; all the analysis needs to know about each cryptosystem is (n, q, s) , without regard to complications such as the error-correcting codes in NewHope or the matrices in Kyber.

A critical parameter in the usual attacks is a “block size” β . The standard analysis explains how to recognize when β is large enough for the attacks to succeed. The point of this section is to calculate how the heuristically required β scales with the PKE parameters (n, q, s) .

Specifically, fix real numbers Q_0, Q_1, S_0, S_1 . To simplify the analysis, assume $0 \leq S_0 \leq 1/2 < Q_0 - S_0$. Consider an infinite sequence of (n, q, s) where

- n runs through a subsequence of $2, 3, 4, \dots$;
- $\lg q \in Q_0 \lg n + Q_1 + o(1)$, or equivalently $q \in n^{Q_0}(2^{Q_1} + o(1))$, where $o(1)$ means the set of functions of n that converge to 0 as $n \rightarrow \infty$; and
- $\lg s \in S_0 \lg n + S_1 + o(1)$.

The conclusion of this section, in short, is that the heuristically required block size β has $\beta/n \in Z_0 + (Z_1 + o(1))/\lg n$ where $Z_0 = 2Q_0/(Q_0 - S_0 + 1/2)^2$ and

$$Z_1 = \left(2S_1 + \lg Z_0 - \left(S_0 - Q_0 + \frac{3}{2} \right) \lg \frac{Z_0}{2\pi \exp 1} - \frac{Q_1(Q_0 + S_0 - \frac{1}{2})}{Q_0} \right) \frac{2Q_0}{(Q_0 - S_0 + \frac{1}{2})^3}.$$

Unsurprisingly, this conclusion applies equally to (1) the usual key-recovery attack and (2) the usual single-target message-recovery attack. This makes the key-recovery attack more attractive, since key recovery trivially implies a multi-ciphertext attack that breaks *all* of the ciphertexts, with just minor cost per ciphertext to run the same decryption algorithm as the legitimate user.

Presumably the use of key-recovery attacks as multi-ciphertext attacks is the motivation for, e.g., [73, Section 6] saying “arguably, the secret key ought to be better-protected than any individual ciphertext”. One can easily modify the LPR cryptosystem to use one error distribution for encryption and a larger error distribution for key generation. This section’s asymptotics for β/n quantify the impact of moving from one pair (S_0, S_1) for encryption to a (lexicographically) larger pair (S_0, S_1) for key generation.

2.1. Comparison to the structure of concrete analyses. What this section does is conceptually similar to what the literature already does in evaluating attack costs for, e.g., New Hope or Kyber: the same heuristic analyses of the same attacks are used to predict attack costs for any given attack parameters; these predictions are then used to optimize attack parameters.

The big difference is that those predictions and optimizations are concrete while this section is asymptotic. Consider, e.g., [13, Table 1] saying the optimal β/n is $967/1024$ for the original version of NewHope, while this section says the optimal β/n is $270/289 + o(1)$ if the parameters (Q_0, S_0) are $(1.35, 0.15)$. These statements about β might seem similar, especially since in this example the numbers $967/1024 = 0.944\dots$ and $270/289 = 0.934\dots$ are close; but $967/1024$ is a statement about one parameter set, while $270/289 + o(1)$ is a statement about an infinite family of parameter sets.

This change forces calculations to be carried out symbolically, rather than through simple enumeration of all possible parameter choices. The compensating advantage is that a reader can use asymptotics to see “big” costs and “big” optimizations, such as changes in algorithm exponents, without having to worry about “small” costs and “small” optimizations, such as polynomial-factor speedups to an exponential-time algorithm. The literature contains many “small” improvements in lattice attacks, and it is useful to be able to skip those in seeing the “big” improvements—such as multi-target improvements.

2.2. Block size versus cost. The main bottleneck in the usual attacks is a BKZ- β computation. BKZ- β is a family of algorithms, not a single algorithm; its cost has dropped over the years, in part because underlying subroutines have been improved, notably for SVP- β , and in part because the use of those subroutines inside BKZ- β has been improved.

For example, Becker–Ducas–Gama–Laarhoven [21] reported heuristic costs $(3/2 + o(1))^{\beta/2} = 2^{(\log_4(3/2) + o(1))\beta}$ for a non-quantum⁵ algorithm to solve SVP- β . The exponent $\log_4(3/2) = 0.292\dots$ was smaller than in previous papers.

To first order, the heuristic BKZ- β cost is the same as the heuristic SVP- β cost: e.g., $2^{(\log_4(3/2) + o(1))\beta}$ using [21]. Combining this with the standard heuristic analysis of the required β , and with the first-order asymptotics $\beta/n \in Z_0 + o(1)$ from this paper where Z_0 is defined above, gives an overall heuristic attack cost of $2^{(Z_0 \log_4(3/2) + o(1))n}$, showing that n must be at least $\lambda / (Z_0 \log_4(3/2) + o(1))$ to reach λ bits of heuristic security.

This paper computes more complicated second-order asymptotics $\beta/n \in Z_0 + (Z_1 + o(1)) / \lg n$ so that attack improvements that change block sizes by $\Theta(n / \lg n)$ become visible in the asymptotics. Examples of such improvements are “dimensions for free” from [45], heuristically reducing block size β to

⁵ The literature reports somewhat lower exponents for quantum algorithms, including improvements in 2021; see [37] and [55]. This directly produces better heuristic exponents for quantum lattice attacks. However, quantum algorithms in this context are typically dismissed as not producing enough speedup to overcome quantum overhead for realistic sizes (see, e.g., [8]) or as not being competitive with quantum speedups against AES. This paper focuses on non-quantum algorithms.

$\beta - (\lg(4/3) + o(1))\beta / \lg \beta = \beta - (0.415\dots + o(1))\beta / \lg \beta$; the improved techniques from [44], heuristically improving $\lg(4/3)$ to $\lg(13/9) = 0.530\dots$; and Section 4 of this paper.

Beware that the first-order asymptotics stated in the literature, such as $(3/2 + o(1))^{\beta/2}$ from [21], are logically insufficient to conclude that a second-order improvement in β such as [45] produces a corresponding improvement in attack cost. To see the issue, imagine a version of BKZ- β in which the cost depends only on $x = \lceil 0.01(\log \beta)^2 \rceil$. This is compatible with the cost having the form $2^{(c+o(1))\beta}$: for example, $10\sqrt{x} - \log \beta \in o(1)$, so $\exp(10\sqrt{x}) \in (1 + o(1))\beta$, so $2^{0.1 \exp(10\sqrt{x})} \in 2^{(0.1+o(1))\beta}$. However, decreasing β to $\beta - \beta / \lg \beta$ changes $0.01(\log \beta)^2$ by only about $0.02 \log 2 = 0.0138\dots$, and thus changes costs for under 1.4% of the values of β , perhaps missing the occasional values of β that appear, asymptotically, in a particular cryptosystem of interest.

It would be interesting to see second-order asymptotics for various algorithms for SVP- β and BKZ- β . For purposes of this paper, it suffices to observe that if costs are $2^{(c+o(1))\beta}$ then the total impact of $\Theta(\lg \beta)$ reductions in β , where each reduction replaces β with $\beta - (d + o(1))\beta / \lg \beta$, is an improvement in c . Even if this improvement is not divided evenly across the reductions in β , it is useful on average across any sufficiently large interval of β , and is thus useful on average across the broad spectrum of parameter choices considered in this paper.

2.3. Review of the LPR cryptosystem. This PKE has three parameters: an integer $n \geq 2$; an integer $q \geq 2$; and a probability distribution χ supported on a finite set of integers. Assume for simplicity that the average of χ is 0. Write R for the ring $\mathbb{Z}[x]/(x^n + 1)$.

Key generation works as follows. Generate uniform random $G \in R/q$. Generate $a, e \in R$ with coefficients drawn independently at random from χ . Compute $A = aG + e \in R/q$. The secret key is (a, e) . The public key is $(G, A) \in (R/q)^2$.

The set of messages is the set of elements of R with coefficients in $\{0, \lceil q/2 \rceil\}$. Encryption of a message M to a public key (G, A) works as follows. Generate $b, c, d \in R$ with coefficients drawn independently at random from χ . Compute $B = Gb + d \in R/q$ and $C = M + Ab + c \in R/q$. The ciphertext is $(B, C) \in (R/q)^2$.

Decryption of a ciphertext (B, C) works as follows. Compute $X = C - aB \in R/q$. Round each coefficient of X to the closest element of $\{0, \lceil q/2 \rceil\}$ in \mathbb{Z}/q , specifically 0 if both elements are equally close.⁶

The above PKE definition skips two requirements from the LPR paper, namely that n is a power of 2 and that q is a prime congruent to 1 modulo $2n$; see [75, Section 1.1]. Cryptosystems after [75] loosened the restrictions on q ; for example, Kyber's current prime 3329 is $1 + 13 \cdot 256$. As for n , readers concerned about attacks enabled by factors of $x^n + 1$ in $\mathbb{Z}[x]$ should feel free to substitute the marginally larger polynomial $x^n - x - 1$, as in [31]; this is orthogonal to the topic of this paper.

⁶ This rounding detail is not specified in [75]; also, [75] says $\lfloor q/2 \rfloor$ without specifying whether the rounding rounds 0.5 up or rounds to even. These details do not affect the standard analysis; they are specified here so as to have a complete PKE definition.

Correct decryption requires $X = C - aB = M + eb + c - ad$ to round to M , i.e., requires $eb + c - ad$ to be smaller than about $q/4$. If χ is, e.g., the uniform distribution on $\{-1, 0, 1\}$ then each coefficient of eb is a sum of n products where one expects about $4/9$ to be nonzero, evenly balanced between 1 and -1 , so typically the coefficient will be on the scale of \sqrt{n} , with considerable variation in the exact size. There are various proposals to reduce q close to this scale, and to avoid frequent decryption failures by applying an error-correcting code to M . In the opposite direction, sometimes cryptosystems take larger χ and correspondingly larger q ; sometimes cryptosystems pack more message bits into each coefficient, again taking larger q . To cover many different cases, this paper considers a spectrum of possibilities for the asymptotic sizes of q and s .

2.4. Review of the usual key-recovery attack. Consider the problem of recovering the private key $(a, e) \in R^2$ from the public key $(G, A) \in (R/q)^2$. Recovering a suffices, since $e = A - aG$ by definition. The usual “primal” attack works as follows.

There is an attack parameter $\kappa \leq n$. Define a function $\text{First}_\kappa : R \rightarrow \mathbb{Z}^\kappa$ that extracts the first κ coefficients from its input. This induces a function, also written First_κ , from R/q to $(\mathbb{Z}/q)^\kappa$.

Define L as the set of all $(\alpha, \epsilon, \beta) \in R \times \mathbb{Z}^\kappa \times \mathbb{Z}$ such that $\text{First}_\kappa(\beta A - \alpha G)$ is the same as ϵ modulo q . This is a lattice of full rank $d = n + \kappa + 1$ and determinant q^κ . Note that $\pm(a, \text{First}_\kappa(e), 1)$ are elements of this lattice.

There is another attack parameter β . The attack writes down a basis for L , applies BKZ- β to reduce this basis, and hopes that BKZ- β outputs one of the short nonzero vectors $\pm(a, \text{First}_\kappa(e), 1)$, in particular revealing a .

The problem being attacked here, the problem of finding a, e given a random G and $aG + e$, is typically called “Ring-LWE”, specifically “normal-form 1-sample search Ring-LWE”, where “normal form” refers to the secret a being small. The Ring-LWE problem is typically credited to [95] and [75]. However, this problem was already attacked in the 1998 Hoffstein–Pipher–Silverman NTRU paper, both for the homogeneous case $A = 0$ (see [56, Section 3.4.1]) and for general A (see [56, Section 3.4.2]). The problem statements in [95] and [75] merely generalize to more “samples”: e.g., finding a, e_1, e_2 given random $G_1, G_2, aG_1 + e_1, aG_2 + e_2$, or equivalently replacing $G \in R/q$ and $e \in R$ with row vectors $(G_1 \ G_2) \in (R/q)^2$ and $(e_1 \ e_2) \in R^2$ respectively.

The attack in [56] has $\kappa = n$. May–Silverman [79] generalized the attack to any $\kappa \leq n$. In the original 1996 NTRU handout [57], various concrete examples chose different sizes for a and e , motivating another generalization from Coppersmith and Shamir [41] to set up a lattice with, e.g., short vector $(3.14a, e, 1)$ rather than $(a, e, 1)$; this paper focuses on cryptosystems that take a and e of the same size, such as the LPR system. There can still be a tiny improvement from setting up a lattice with, e.g., short vector $(a, e, 3.14)$; this paper ignores this improvement for simplicity.

2.5. Review of the standard analysis. Beyond asking about the cost of BKZ- β (see Section 2.2), the standard analysis of Section 2.4’s attack asks

whether BKZ- β succeeds at finding the target vector. The standard heuristic conclusion is that BKZ- β succeeds if and only if the 2-norm of the target vector is below $(d/\beta)^{1/2}\delta^{2\beta-d-1}q^{\kappa/d}$, where $\delta = (\beta(\pi\beta)^{1/\beta}/(2\pi \exp 1))^{1/2(\beta-1)}$. The rationale for this inequality is as follows:

- One heuristic says that, for a “random” lattice L of rank d , BKZ- β finds a nonzero vector of length close to $\delta^{d-1}(\det L)^{1/d}$, with δ defined as above.
- Another heuristic says that the Gram–Schmidt lengths of the BKZ- β output are close to a geometric series. Combining this with the shortest length being close to $\delta^{d-1}(\det L)^{1/d}$ and the product of the lengths being $\det L$ says approximately how large each length is. In particular, the length at position $d - \beta + 1$ is close to $\delta^{2\beta-d-1}(\det L)^{1/d}$, which for this lattice is $\delta^{2\beta-d-1}q^{\kappa/d}$. The rationale treats this approximation as an equation.
- Another heuristic says that if the target vector has length t then its projection onto the space spanned by the last β Gram–Schmidt vectors has length approximately $t\sqrt{\beta/d}$. The rationale also treats this approximation as an equation.
- If the latter length $t\sqrt{\beta/d}$ is below the previous length $\delta^{2\beta-d-1}q^{\kappa/d}$ then the above heuristics seem to contradict each other, since the last SVP- β call in each “tour” of BKZ- β guarantees that the projection of the vector at position $d - \beta + 1$ is a minimum-length nonzero vector in the projection of L . Note, however, that the first heuristic was only for a “random” lattice. Another heuristic says that this seeming contradiction occurs if and only if BKZ- β detects the non-“randomness” of the lattice by finding the projection of v .
- A further heuristic says that BKZ- β finds the projection of v if and only if BKZ- β finds v . A slightly different statement appears in [9], which says that if BKZ- β finds the projection of v then BKZ- β finds v with “high probability” for large β .

The following paragraphs use the inequality as a black box without regard to the rationale, but the rationale matters for Section 3.

For the LPR PKE, the first $n + \kappa$ entries in the target vector $(a, \text{First}_\kappa(e), 1)$ are drawn independently and uniformly at random from χ . Each entry has square $\sum_i \chi_i i^2 = s^2$ on average (since χ has average 0 and standard deviation s), so the squared 2-norm of $(a, \text{First}_\kappa(e), 1)$ is $(n + \kappa)s^2 + 1$ on average. The standard heuristic analysis treats the squared 2-norm as being exactly its average, concluding for this PKE that BKZ- β works if and only if $((n + \kappa)s^2 + 1)^{1/2} < (d/\beta)^{1/2}\delta^{2\beta-d-1}q^{\kappa/d}$.

Other PKEs do not necessarily choose (a, e) this way: consider, e.g., a PKE that chooses a as a fixed-weight ternary vector. To apply the standard analysis to such cases, the literature calculates s so that $(n + \kappa)s^2 + 1$ is a reasonable estimate of the squared 2-norm of $(a, \text{First}_\kappa(e), 1)$, and concludes heuristically that BKZ- β works if and only if $((n + \kappa)s^2 + 1)^{1/2} < (d/\beta)^{1/2}\delta^{2\beta-d-1}q^{\kappa/d}$.

2.6. Asymptotic heuristic attack costs. The following theorem considers attack parameters $\kappa \in (X_0 + (X_1 + o(1)) / \lg n)n$ and $\beta \in (Y_0 + (Y_1 + o(1)) / \lg n)n$,

and asks whether the inequality $((n + \kappa)s^2 + 1)^{1/2} < (d/\beta)^{1/2}\delta^{2\beta-d-1}q^{\kappa/d}$ holds asymptotically.

In short, the cutoff for (Y_0, Y_1) is the (Z_0, Z_1) defined in the theorem as a function of $(Q_0, Q_1, S_0, S_1, X_0, X_1)$. Section 2.7 then chooses (X_0, X_1) to minimize (Z_0, Z_1) , obtaining the above formulas for (Z_0, Z_1) as a function of (Q_0, Q_1, S_0, S_1) .

Theorem 2.6.1 (asymptotics of the standard block-size inequality). *Let $Q_0, Q_1, S_0, S_1, X_0, X_1, Y_0, Y_1$ be real numbers such that $0 < Q_0$; $-1/2 < S_0$; $(0, 0) \leq (X_0, X_1) \leq (1, 0)$; $0 < Y_0$; $(Y_0, Y_1) \leq (1 + X_0, X_1)$; and $0 < 1 - 2S_0 + 2Q_0X_0/(1 + X_0)$. Define*

$$Z_0 = \frac{1 + X_0}{1 - 2S_0 + 2Q_0X_0/(1 + X_0)},$$

$$Z_1 = \left(2S_1 + \lg Z_0 + \frac{X_1}{Z_0} - \left(2 - \frac{1 + X_0}{Z_0}\right) \lg \frac{Z_0}{2\pi \exp 1} - \frac{2Q_1X_0}{1 + X_0} - \frac{2Q_0X_1}{(1 + X_0)^2}\right) \frac{Z_0^2}{1 + X_0}.$$

Let N be an infinite subset of $\{2, 3, 4, 5, \dots\}$. Let $n \mapsto q$, $n \mapsto s$, $n \mapsto \kappa$, $n \mapsto \beta$ be functions from N to \mathbb{R} such that

$$\begin{aligned} 2 &\leq q, & \lg q &\in Q_0 \lg n + Q_1 + o(1), \\ 0 &< s, & \lg s &\in S_0 \lg n + S_1 + o(1), \\ 1 &\leq \kappa \leq n, & \kappa/n &\in X_0 + (X_1 + o(1))/\lg n, \\ 2 &\leq \beta \leq n + \kappa + 1, & \beta/n &\in Y_0 + (Y_1 + o(1))/\lg n. \end{aligned}$$

Define $d = n + \kappa + 1$; $\delta = (\beta(\pi\beta)^{1/\beta}/(2\pi \exp 1))^{1/2(\beta-1)}$; and

$$\rho = \frac{((n + \kappa)s^2 + 1)^{1/2}}{(d/\beta)^{1/2}\delta^{2\beta-d-1}q^{\kappa/d}}.$$

(1) If $(Y_0, Y_1) > (Z_0, Z_1)$ then $\rho < 1$ for all sufficiently large n . (2) If $(Y_0, Y_1) < (Z_0, Z_1)$ then $\rho > 1$ for all sufficiently large n .

Proof. Note first that arithmetic on quantities of the form $A_0 + (A_1 + o(1))/\lg n$, where $A_0, A_1 \in \mathbb{R}$, matches arithmetic on elements $A_0 + A_1\epsilon$ of the ring $\mathbb{R}[\epsilon]/\epsilon^2$. For example,

$$\left(A_0 + \frac{A_1 + o(1)}{\lg n}\right) \left(B_0 + \frac{B_1 + o(1)}{\lg n}\right) \subseteq A_0B_0 + \frac{A_0B_1 + A_1B_0 + o(1)}{\lg n}$$

since $A_0o(1)/\lg n$, $A_1o(1)/\lg n$, and $(A_1 + o(1))(B_1 + o(1))/(\lg n)^2$ are all contained in $o(1)/\lg n$. As another example, $1/(A_0 + (A_1 + o(1))/\lg n)$ is $1/A_0 + (-A_1/A_0^2 + o(1))/\lg n$ if $A_0 \neq 0$.

In particular, $d/n = 1 + \kappa/n + 1/n \in 1 + X_0 + (X_1 + o(1))/\lg n$, so

$$\frac{n}{d} \in \frac{1}{1 + X_0} + \frac{-X_1/(1 + X_0)^2 + o(1)}{\lg n},$$

so

$$\begin{aligned} \frac{\kappa}{d} &\in \frac{X_0}{1+X_0} + \frac{X_1/(1+X_0) - X_0X_1/(1+X_0)^2 + o(1)}{\lg n} \\ &= \frac{X_0}{1+X_0} + \frac{X_1/(1+X_0)^2 + o(1)}{\lg n}. \end{aligned}$$

Meanwhile $(\lg q)/\lg n \in Q_0 + (Q_1 + o(1))/\lg n$. Multiply to see that

$$\frac{\kappa \lg q}{d \lg n} \in \frac{Q_0 X_0}{1+X_0} + \frac{Q_1 X_0/(1+X_0) + Q_0 X_1/(1+X_0)^2 + o(1)}{\lg n}.$$

The proof does not need as much precision for d/β : one has $n/\beta \in 1/Y_0 + o(1)$ and $d/n \in 1 + X_0 + o(1)$, so $d/\beta \in (1 + X_0)/Y_0 + o(1)$, so

$$\frac{\lg(d/\beta)}{\lg n} \in \frac{\lg(1 + X_0) - \lg Y_0 + o(1)}{\lg n}.$$

Next $s^2 \in n^{2S_0} 2^{2S_1 + o(1)} = n^{2S_0} (2^{2S_1} + o(1))$ and $n + \kappa \in n(1 + X_0 + o(1))$ so $(n + \kappa)s^2 \in n^{1+2S_0} ((1 + X_0)2^{2S_1} + o(1))$. The exponent $1 + 2S_0$ is positive, so also $(n + \kappa)s^2 + 1 \in n^{1+2S_0} ((1 + X_0)2^{2S_1} + o(1))$. Hence

$$\frac{\lg((n + \kappa)s^2 + 1)}{\lg n} \in 1 + 2S_0 + \frac{\lg(1 + X_0) + 2S_1 + o(1)}{\lg n}.$$

By hypothesis $\beta \in (Y_0 + o(1))n$ with $Y_0 > 0$, so $\beta \in \Theta(n)$, so $\lg \pi\beta \in O(\lg n)$, so $(\lg \pi\beta)/\beta \in O(\lg n)/n$. What matters here is merely that $(\lg \pi\beta)/\beta \in o(1)$, so

$$2(\beta - 1) \lg \delta = \lg \beta + \frac{\lg \pi\beta}{\beta} - \lg(2\pi \exp 1) \in \lg n + \lg Y_0 - \lg(2\pi \exp 1) + o(1).$$

Divide $(d-1)/n \in 1 + X_0 + (X_1 + o(1))/\lg n$ by $(\beta-1)/n \in Y_0 + (Y_1 + o(1))/\lg n$ to see that

$$\frac{d-1}{\beta-1} \in \frac{1+X_0}{Y_0} + \frac{X_1/Y_0 - (1+X_0)Y_1/Y_0^2 + o(1)}{\lg n}.$$

Consequently

$$\frac{2\beta - d - 1}{2(\beta - 1)} = 1 - \frac{d-1}{2(\beta-1)} \in 1 - \frac{1+X_0}{2Y_0} - \frac{X_1/2Y_0 - (1+X_0)Y_1/2Y_0^2 + o(1)}{\lg n}.$$

Multiply by $2(\beta-1)(\lg \delta)/\lg n \in 1 + (\lg Y_0 - \lg(2\pi \exp 1) + o(1))/\lg n$ to see that

$$\begin{aligned} \frac{(2\beta - d - 1) \lg \delta}{\lg n} &\in 1 - \frac{1+X_0}{2Y_0} \\ &+ \frac{-X_1/2Y_0 + (1+X_0)Y_1/2Y_0^2 + (1-(1+X_0)/2Y_0)(\lg Y_0 - \lg(2\pi \exp 1)) + o(1)}{\lg n}. \end{aligned}$$

Now add to see that

$$\begin{aligned} \frac{2 \lg \rho}{\lg n} &= \frac{\lg((n + \kappa)s^2 + 1)}{\lg n} - \frac{\lg(d/\beta)}{\lg n} - \frac{2 \lg \delta^{2\beta-d-1}}{\lg n} - \frac{2(\kappa/d) \lg q}{\lg n} \\ &\in 1 + 2S_0 + \frac{\lg(1+X_0) + 2S_1}{\lg n} - \frac{\lg(1+X_0) - \lg Y_0}{\lg n} - 2 + \frac{1+X_0}{Y_0} \\ &\quad + \frac{X_1/Y_0 - (1+X_0)Y_1/Y_0^2 - (2-(1+X_0)/Y_0)(\lg Y_0 - \lg(2\pi \exp 1))}{\lg n} \\ &\quad - \frac{2Q_0X_0}{1+X_0} - \frac{2Q_1X_0/(1+X_0) + 2Q_0X_1/(1+X_0)^2}{\lg n} + \frac{o(1)}{\lg n}. \end{aligned}$$

If $Y_0 > Z_0$, i.e., $Y_0 > (1 + X_0)/(1 - 2S_0 + 2Q_0X_0/(1 + X_0))$, then the limit $1 + 2S_0 - 2 + (1 + X_0)/Y_0 - 2Q_0X_0/(1 + X_0)$ of $(2 \lg \rho)/\lg n$ is negative, so $(2 \lg \rho)/\lg n < 0$ for all sufficiently large n , so $\rho < 1$ for all sufficiently large n as claimed.

If $Y_0 < Z_0$ then the limit is positive, so $\rho > 1$ for all sufficiently large n as claimed.

Assume from now on that $Y_0 = Z_0$. Then the limit is 0, and $(2 \lg \rho)/\lg n \in (\Delta + o(1))/\lg n$ where

$$\begin{aligned} \Delta &= 2S_1 + \lg Z_0 + \frac{X_1}{Z_0} - \frac{(1 + X_0)Y_1}{Z_0^2} \\ &\quad - \left(2 - \frac{1 + X_0}{Z_0}\right) \lg \frac{Z_0}{2\pi \exp 1} - \frac{2Q_1X_0}{1 + X_0} - \frac{2Q_0X_1}{(1 + X_0)^2} \\ &= \frac{(1 + X_0)(Z_1 - Y_1)}{Z_0^2}; \end{aligned}$$

note that the $\lg(1 + X_0)$ terms cancel.

If $Y_1 > Z_1$ then $\Delta < 0$ so $\rho < 1$ for all sufficiently large n as claimed. If $Y_1 < Z_1$ then $\Delta > 0$ so $\rho > 1$ for all sufficiently large n as claimed. The only remaining possibility is $(Y_0, Y_1) = (Z_0, Z_1)$, and the theorem statement makes no claims regarding this case. \square

2.7. Asymptotic heuristic optimization of attack parameters. Fix Q_0 and S_0 . Assume for simplicity that $0 \leq S_0 \leq 1/2 < Q_0 - S_0$; then also $1/2 < Q_0 + S_0$. Consider the problem of choosing (X_0, X_1) to minimize (Z_0, Z_1) in Theorem 2.6.1; this is the asymptotic version of choosing κ to allow the smallest possible β , subject to the standard inequality.

Define $W_0 = (Q_0 + S_0 - 1/2)/(Q_0 - S_0 + 1/2)$. Then $0 < W_0$ since $1 < 2Q_0 + 2S_0$; $W_0 \leq 1$ since $1 - W_0 = 2(1/2 - S_0)/(Q_0 - S_0 + 1/2)$; and $1 - 2S_0 + 2Q_0W_0/(1 + W_0) = Q_0 - S_0 + 1/2 > 1$. Consequently, taking $X_0 = W_0$ in Theorem 2.6.1 is compatible with the hypotheses that $0 \leq X_0 \leq 1$ and $0 < 1 - 2S_0 + 2Q_0X_0/(1 + X_0)$.

The theorem defines $Z_0 = (1 + X_0)/(1 - 2S_0 + 2Q_0X_0/(1 + X_0))$. If $X_0 = W_0$ then $Z_0 = (1 + W_0)/(Q_0 - S_0 + 1/2) = 2Q_0/(Q_0 - S_0 + 1/2)^2$. Note that $Z_0 < 1 + W_0$ since $Q_0 - S_0 + 1/2 > 1$. Taking $Y_0 = Z_0$, and taking any real number Y_1 ,

is thus compatible with the hypotheses that $0 < Y_0$ and $(Y_0, Y_1) \leq (1 + X_0, X_1)$. Note for future reference that taking $Y_0 = Z_0$ gives $d/\beta \in Q_0 - S_0 + 1/2 + o(1)$.

This is the minimum possible value of Z_0 among nonnegative real numbers X_0 : the derivative of Z_0 with respect to X_0 is

$$\frac{(1 + X_0)((Q_0 - S_0 + 1/2)X_0 - (Q_0 + S_0 - 1/2))}{2((1/2 - S_0)(1 + X_0) + Q_0 X_0)^2},$$

which, within nonnegative real numbers X_0 , is 0 for $X_0 = W_0$, negative for $X_0 < W_0$, and positive for $X_0 > W_0$.

Furthermore, taking $X_0 = W_0$ cancels the X_1 terms in the definition of Z_1 in the theorem: one has $1/Z_0 = 2Q_0/(1 + W_0)^2$, so

$$\begin{aligned} Z_1 &= \left(2S_1 + \lg Z_0 - \left(2 - \frac{1 + W_0}{Z_0} \right) \lg \frac{Z_0}{2\pi \exp 1} - \frac{2Q_1 W_0}{1 + W_0} \right) \frac{Z_0^2}{1 + W_0} \\ &= \left(2S_1 + \lg Z_0 - \left(S_0 - Q_0 + \frac{3}{2} \right) \lg \frac{Z_0}{2\pi \exp 1} - \frac{Q_1(Q_0 + S_0 - \frac{1}{2})}{Q_0} \right) \frac{2Q_0}{(Q_0 - S_0 + \frac{1}{2})^3}. \end{aligned}$$

The theorem now says that taking any $Y_1 > Z_1$ and any X_1 , along with $Y_0 = Z_0$ and $X_0 = W_0$, meets the standard inequality for all sufficiently large n .

In particular, pick a function $n \mapsto \kappa$ with $\kappa \in n(W_0 + o(1))/\lg n$, and, for each $j \geq 1$, pick a function $n \mapsto \beta_j$ with $\beta_j \in n(Z_0 + (Z_1 + 2^{-j} + o(1)))/\lg n$. The $o(1)$ for β_j has absolute value at most 2^{-j} for all sufficiently large n , say $n \geq n_j$; and (κ, β_j) meets the inequality for all sufficiently large n , say $n \geq m_j$. Assume without loss of generality that $m_j = n_j$; if not, replace both m_j, n_j with $\max\{m_j, n_j\}$. Assume without loss of generality that $n_1 < n_2$; if not, replace n_2 (and m_2) with $n_1 + 1$. Similarly assume $n_2 < n_3 < \dots$. Now glue β_1, β_2, \dots together into a function β that is β_1 for $n < n_2$, then β_2 for $n_2 \leq n < n_3$, then β_3 for $n_3 \leq n < n_4$, etc. Then (κ, β) meet the inequality for all $n \geq n_1$, and $\beta \in n(Z_0 + (Z_1 + o(1)))/\lg n$.

In short, this minimum value of (Z_0, Z_1) is achievable for β meeting the standard inequality. This does not mean that *every* $\beta \in n(Z_0 + (Z_1 + o(1)))/\lg n$ meets the inequality: on the contrary, (Z_0, Z_1) is also achievable for β *not* meeting the inequality.

This calculation also shows that β cannot do better than this minimum (Z_0, Z_1) : any functions $n \mapsto \kappa$ and $n \mapsto \beta$ of the form assumed in the theorem, with $(Y_0, Y_1) < (Z_0, Z_1)$, fail the inequality for all sufficiently large n .

What about *arbitrary* functions $n \mapsto \kappa$ and $n \mapsto \beta$ that satisfy the inequality for all sufficiently large n ? Suppose that, for some $\epsilon > 0$, an infinite set of n has $\beta/n \leq Z_0 - \epsilon$, with Z_0 as above. Then, by compactness, there is an infinite subset with $(\kappa/n, \beta/n)$ converging. Within this subset, κ/n has the form $X_0 + o(1)$, and β/n has the form $Y_0 + o(1)$, where $0 \leq X_0 \leq 1$, $0 \leq Y_0 \leq 1 + X_0$, and $Y_0 \leq Z_0 - \epsilon$. A simplified version of Theorem 2.6.1, suppressing the second-order terms, then obtains a contradiction. Consequently $\beta/n \geq Z_0 + o(1)$.

More effort is required for the second-order terms. Say there is an infinite set of n with $\beta/n \leq Z_0 + (Z_1 - \epsilon)/\lg n$. One wants to extract an infinite subset where

β/n has the form $Y_0 + (Y_1 + o(1))/\lg n$, but the compactness argument requires upper and lower bounds: for example, if β has the form $Z_0 + (Z_1 - \Theta(\lg \lg n))/\lg n$, then one cannot specialize n to put β in the form $Z_0 + (Y_1 + o(1))/\lg n$. A workaround is to increase β up to $Z_0 + (Z_1 - \epsilon + o(1))/\lg n$: the point is that, for sufficiently large n , if β satisfies the inequality then $\beta+1$ does too. One then needs yet another variant of Theorem 2.6.1 to prove that $\beta/n \in Z_0 + (Z_1 - \epsilon + o(1))/\lg n$ fails the inequality for all sufficiently large n whenever $\kappa/n \in W_0 + o(1)$, because of the X_1 cancellation noted above.

2.8. Asymptotics for the standard message-recovery attack. Consider now the problem of recovering the encryption secrets (b, d) from (G, B) where $B = Gb + d$. The standard analysis handles this exactly the same way as the key-recovery attack, starting with the distribution of (b, d) rather than the distribution of (a, e) .

In the case of the LPR PKE, the distributions are the same, so the conclusions are the same. The attacker will then prefer to carry out the key-recovery attack since it breaks many ciphertexts; this can be understood as motivation to modify the PKE to take a larger distribution for (a, e) than for (b, d) , as noted above.

The second-order asymptotics make it easy to see the effect of making errors 1 bit larger, i.e., increasing S_1 by 1: this increases Z_1 by $4Q_0/(Q_0 - S_0 + 1/2)^3$, increasing the heuristically required β by $(4Q_0/(Q_0 - S_0 + 1/2)^3 + o(1))n/\lg n$.

2.9. Further attacks. The standard analysis also considers the problem of recovering (b, c, d) given $G, A, B = Gb + d$, and $C = Ab + c$, i.e., from a public key and an encryption of 0. A successful recovery attack immediately gives an IND-CPA attack.

Structurally, this problem provides more “samples” to the attacker, allowing κ to be chosen as large as $2n$. However, the optimization presented in Section 2.7 ends up with $X_0 = W_0 \leq 1$ even if X_0 is initially allowed to be as large as 2.

Note that W_0 would be above 1 if S_0 were allowed to be above $1/2$. For essentially the same reason, there are some cryptosystem proposals for which [6] lists slightly lower numbers⁷ for “LWE $2n$ samples” than for “LWE n samples”.

The literature often reports that a “dual” attack allows marginally smaller β than the primal attack. However, [17, page 26, “Primal attack only”] argues that the dual attack is actually much slower than the primal attack. On the other hand, recent attack papers cited in Section 1 introduce new dual attacks that are claimed to be noticeably faster than primal attacks. Since single-target dual attacks are such an unstable topic at the moment, trying to optimize multi-target dual attacks would be premature. More to the point, this paper

⁷ This should already have raised questions regarding a common notion that releasing more samples preserves security levels. This notion is made explicit in, e.g., [47, page 3], which says “we believe that in practice the MLWE problem with k samples is *no easier* than with 1 sample” (emphasis added). This notion is also implicitly assumed by an argument in the literature that GAM/LPR cryptosystems are “at least as hard” to break as NTRU; see generally [85, Sections 3.14, 5.1, and 5.2] for analysis of that argument and further references.

is investigating how multiple ciphertexts asymptotically affect the standard analysis; the standard analysis does not include the new dual attacks.

IND-CCA2 attacks against lattice KEMs can be much easier than the usual lattice attacks against the underlying PKEs. See, e.g., the Round2 break in [22], or the FrodoKEM attack in Section 6. However, lattice attacks against PKEs seem to be the top threat for most lattice proposals.

3 Asymptotic heuristic cost of breaking many ciphertexts

This section analyzes what the existing heuristics say about another attack that breaks many ciphertexts. There is nothing new in the attack algorithm—in short, it plugs [46] into the algorithm briefly outlined in [48, Section 5, third paragraph], and applies this to multiple ciphertexts—but the impact on PKE security and KEM security was not clear from the literature.

Quantitatively, say cryptosystem parameters are chosen so that the usual single-target message-recovery attack reviewed in Section 2.8 has heuristic cost $2^{(1+o(1))\lambda}$. This section shows that total heuristic cost $2^{(1+o(1))\lambda}$ suffices to decrypt T ciphertexts, for any $T \leq 2^{(0.19\dots+o(1))\lambda}$. The heuristic per-ciphertext cost is just $2^{(1+o(1))\lambda}/T$, as low as $2^{(0.80\dots+o(1))\lambda}$. Because this is a first-order change in the heuristic per-ciphertext cost exponent as soon as $T \in 2^{\Theta(\lambda)}$, there is no need for this section to consider second-order terms in the asymptotics.

This is, at this level of asymptotic detail, an “all-for-the-price-of-one” attack heuristically breaking many ciphertexts. This is already achieved by the usual key-recovery attack from Section 2.4 *if* key generation and encryption use the same error distribution, as in the LPR PKE; but it is important to realize that modifying key generation to use a larger error distribution, so as to increase the heuristic cost of the key-recovery attack, fails to heuristically protect the ciphertexts. This section is also a stepping-stone to the one-out-of-many analysis in Section 4.

“Multi-ciphertext security degradation” in the title of this paper refers to the fact that the existing lattice heuristics claim less security against these multi-ciphertext attacks—both in the all-for-price-of-one scenario of this section, and the one-out-of-many scenario in Section 4—than they claim against the usual single-ciphertext attacks. Extrapolating from recent history (see, e.g., [85, Section 1]) suggests that the known attacks are not the end of the story; further attacks could easily change the single-target/multi-target cost ratio. Certainly there is no proof to the contrary: for example, no proof rules out the extreme possibility of an attack breaking typical lattice-based cryptosystems in time $n^{O(1)}$, in which case attacking $2^{\Theta(n)}$ ciphertexts would certainly not produce a $2^{\Theta(n)}$ speedup.

3.1. High-level structure of multi-target lattice attacks. In 2007, Howgrave-Graham [60] introduced a “hybrid” attack algorithm against lattice problems. “Hybrid” here refers to a mixture of combinatorial searches and lattice-basis reduction. The algorithm was stated as an attack against NTRU but also applies to LWE, Ring-LWE, etc.

Howgrave-Graham’s algorithm starts with a single attack target. It carries out a brute-force search for some error positions. Each guess of those error positions produces a lower-dimensional target; note that there are many of these targets. Each of these lower-dimensional targets is handled by an algorithm typically called the “nearest-plane algorithm”, after a preliminary *target-independent* reduction of the relevant lattice.

The nearest-plane algorithm, also known as “weak reduction”, finds lattice vectors close to a given target. It takes a lattice basis as input, and the resulting closeness depends on how well the lattice basis has been reduced. An analysis of the nearest-plane algorithm appears implicitly as part of the analysis in the LLL paper [72], and explicitly in a paper by Babai [19].

The previous two paragraphs oversimplify Howgrave-Graham’s algorithm in two ways:

- [60, Section 4] presented a more sophisticated meet-in-the-middle search, after presenting a brute-force search (“an algorithm that enumerates all possible v ”) as a warmup.⁸
- [60, Section 7] briefly suggested “using a better CVP algorithm than Babai’s closest plane algorithm (e.g. mixing Babai’s CVP (which is essentially blocksize 1) with searching in higher blocksizes 2, 3, ...)”, while keeping the total time under control.

This paper attacks errors of any size, and follows the standard analysis of lattice security in ignoring the possibility of combinatorially searching for small errors.⁹ Consequently, the distinction between a brute-force search and a meet-in-the-middle search is not relevant here. What matters is simply the portion of Howgrave-Graham’s algorithm that searches for lattice vectors close to a batch of targets, after precomputation of short lattice vectors.

The question of what block size β to use—and how to handle β -dimensional close-vector computations after precomputation—remains important in this context. The nearest-plane algorithm, the case $\beta = 1$, runs in polynomial time, but will it find vectors at the necessary distance, after a given amount of effort precomputing short lattice vectors? The necessary distance could be too large for the nearest-plane algorithm to find; perhaps a larger β succeeds within the same total time budget for precomputations and main computations.

In 2015, the original New Hope proposal [13, page 5] briefly considered the following attack against cryptosystems that share a lattice across many

⁸ A paper years later from Lindner–Peikert [73] claimed a “new” lattice attack on LWE, a “decoding” attack that “combines basis reduction with an enumeration algorithm”; [73] does not explain any differences between this decoding attack and the warmup in [60].

⁹ If the error width s is $\Theta(1)$ then, heuristically, hybrid attacks provide an exponential speedup, reducing the asymptotic attack exponent. This raises questions regarding the wisdom of choosing parameters based on attack analyses that ignore hybrid attacks, such as the Kyber security analysis. Even for somewhat larger s , the hybrid speedup should be visible in second-order asymptotics. However, quantifying the hybrid improvement is outside the scope of this paper.

ciphertexts: “finding *once* a good enough basis of the lattice” and then compromising “*all* communications, using for example Babai’s decoding algorithm”. This is the case $\beta = 1$ of the same multi-ciphertext attack. This again begs the question of whether this attack is competitive with the usual single-target attacks. Each ciphertext is handled efficiently, but one also has to quantify the time spent on precomputation and analyze whether the ciphertexts are decoded successfully.

In 2020, Espitau and Kirchner [48], introduced the name “nearest-colattice algorithm” for the block-size- β generalization of the nearest-plane algorithm, wrote “we believe that this algorithm has been in the folklore for some time”, and analyzed the asymptotic performance of this generalization. Specifically, [48, Theorem 3.3] states, assuming heuristics and assuming $d > 2\beta$, that the nearest-colattice algorithm finds a vector in a d -dimensional lattice L at distance $\Theta(\beta)^{d/2\beta}(\det L)^{1/d}$ from any given target. The algorithm uses a series of closest-vector computations in β -dimensional lattices derived from L ; the lattices are independent of the target, so one can “use CVP algorithms after precomputations”, as noted in [48, Section 3].

Espitau and Kirchner also briefly outlined a variant of this algorithm [48, Section 5, third paragraph] that, heuristically, finds an element of L when it is known that the element has distance below $(d/\beta)^{1/2}\Theta(\beta)^{1-d/2\beta}(\det L)^{1/d}$ from a given target vector, asymptotically the same inequality as in Section 2.5:

- Start by finding a “highly reduced basis” of the lattice L . Presumably this means applying BKZ- β .
- Compute a “CVP on the tail of the basis”. This is just one closest-vector computation (rather than a series of closest-vector computations) for the given target, finding a vector in the β -dimensional projection of L closest to the projection of the target.
- “Finish with Babai’s algorithm”, hopefully recovering the closest vector in L to the target.

The rationale for claiming distance $(d/\beta)^{1/2}\Theta(\beta)^{1-d/2\beta}(\det L)^{1/d}$ is not spelled out in [48], but presumably is intended to be as follows, by analogy to the rationale reviewed in Section 2.5:

- The δ heuristic and the geometric-series heuristic imply that the shortest nonzero vector in the projection has length approximately $\delta^{2\beta-d-1}(\det L)^{1/d}$.
- Another existing heuristic says that, for “random” lattices and “random” targets, the distance to the closest vector is approximately the length of the shortest nonzero vector, hence approximately $\delta^{2\beta-d-1}(\det L)^{1/d}$.
- If the target is at distance t from L then its projection is, heuristically, at distance approximately $t\sqrt{\beta/d}$ from the β -dimensional projection of L .
- The approximations are again treated as equations, giving a contradiction if $t\sqrt{\beta/d} < \delta^{2\beta-d-1}(\det L)^{1/d}$. Heuristically, this contradiction occurs if and only if the computation finds the target.
- Finally, δ is $\Theta(\beta)^{1/2\beta}$ as $\beta \rightarrow \infty$, so $\delta^{2\beta-d-1}$ is $\Theta(\beta)^{1-d/2\beta}$.

The $\sqrt{\beta/d}$ factor is claimed in [48] to not be “significant” under the assumption $d > 2\beta$, and is thus suppressed in [48]. However, if $d = 2\beta + 1$, then $\Theta(\beta)^{1-d/2\beta}$ is $\Theta(\beta)^{-1/2\beta}$, which is in $1 + o(1)$, while $\sqrt{\beta/d} = \sqrt{1/2} + o(1)$. To derive the asymptotics stated in [48] from the existing heuristics, one needs to make a slightly stronger assumption, such as $d > 2.01\beta$, or modify the asymptotics to include the $\sqrt{\beta/d}$ factor. When the $\sqrt{\beta/d}$ factor is included, there is nothing in the existing heuristic analysis of this algorithm that relies on the assumption $d > 2\beta$, so this paper does not make that assumption.¹⁰

If the closest-vs.-shortest ratio is heuristically assumed to be $1 + o(1)$, as in [48, Section 3.3, last sentence], then this rationale says that the algorithm works when the distance is below $(1 + o(1))(d/\beta)^{1/2}\delta^{2\beta-d-1}(\det L)^{1/d}$. This is within a factor $1 + o(1)$ of the inequality in Theorem 2.6.1. The second-order asymptotics there generalize immediately to this situation, since $(\lg(1 + o(1)))/\lg n$ is $o(1)/\lg n$. The weaker $\Theta(\beta)$ statement in [48] suffices for first-order asymptotics.

3.2. CVP after precomputation. The next step in the analysis is to see how precomputation changes the exponent of CVP algorithms, algorithms to find *closest* vectors, since these algorithms are used as subroutines above.

The basic idea used in the literature is to precompute a database of all short nonzero vectors in L , and then apply the following greedy reduction algorithm:¹¹ start from a target vector t ; repeatedly replace t with the shortest vector having the form $t - ju$ where u is a database entry and j is an integer; stop when no $t - ju$ is shorter than t . The algorithm output is the sum of vectors ju that were used, i.e., the original t minus the final t . This is a lattice vector close to t , hopefully the closest vector.

Sommer, Feder, and Shalvi [94] showed that a database of $2^{d+1} - 2$ vectors suffices for perfect reduction. Laarhoven [71] argued heuristically that $2^{(1/2+o(1))d}$ vectors suffice. The literature presents two important ways to further reduce the heuristic cost exponent:

- Use a smaller database, i.e., precompute fewer short vectors in L . This reduces the heuristic probability of finding the closest vector, but hopefully saves more time. The literature then amplifies the heuristic probability close to 1 by repeating the algorithm many times, each time adding a random lattice vector to t .
- Hash the database into buckets using a “locality-sensitive hash function”, and search only buckets close to the hash of t . This again trades heuristic probability for time. The literature then amplifies the heuristic probability close to 1 by trying many hash functions.

¹⁰ Recall from Section 2.7 that the optimal d/β there is $Q_0 - S_0 + 1/2 + o(1)$. Readers who wish to restrict attention to d/β being asymptotically above 2, so as to be able to more directly apply [48], can restrict attention to parameter choices with $Q_0 - S_0 + 1/2 > 2$.

¹¹ Often this greedy algorithm, the “iterative slicer”, is credited to [94]. However, the algorithm had already appeared in a 2000 textbook by Cohen [40, pages 375–376]. See [32, Appendix D.3] for the relevant quotes.

The latter “near-neighbor” techniques are also used in, e.g., the aforementioned paper [21] to heuristically save time in building the database of short vectors in the first place.

Ducas–Laarhoven–van Woerden [46], heuristically improving on results from Laarhoven [71] and Doulgerakis–Laarhoven–de Weger [43], concluded¹² that one can “heuristically solve CVPP instances in $2^{0.234d+o(d)}$ amortized time, for batches of size at least $2^{0.058d+o(d)}$ ”. Here $(0.234, 0.058)$ actually means an approximation to a pair of real numbers with sum $\log_4(3/2)$. Formulas for the real numbers can in principle be calculated from [46]; this paper takes 0.058 as axiomatic and does not review the calculations.

The minimum batch size comes from the fact that there is an initial heuristic cost of $(3/2 + o(1))^{d/2} = 2^{(0.292\dots+o(1))d}$ for computing the database of all short vectors. The algorithm also works for smaller batches, but then the total heuristic cost is dominated by this initial heuristic cost, so the per-target heuristic cost is this initial heuristic cost divided by the number of targets.

3.3. Breaking many ciphertexts. The results cited above are used essentially as black boxes in the following analysis of the heuristic cost of a multi-ciphertext attack that breaks *all* of the ciphertexts.

By definition, the j th ciphertext is $(B_j, C_j) = (Gb_j + d_j, M_j + Ab_j + c_j)$; all of b_j, c_j, d_j are small, M_j is the j th message, and the public key reveals G and A . The attack will use G and B_j to find b_j , then subtract Ab_j from C_j to find $M_j + c_j$, then round to find M_j as desired. The critical step is finding b_j , which works as follows.

As in Section 2.4, there is an attack parameter κ , and a map $\text{First}_\kappa : R \rightarrow \mathbb{Z}^\kappa$ that selects the first κ coefficients of its input, inducing a map $R/q \rightarrow (\mathbb{Z}/q)^\kappa$.

Lift $\text{First}_\kappa(B_j) \in (\mathbb{Z}/q)^\kappa$ to \mathbb{Z}^κ : i.e., choose a vector $v_j \in \mathbb{Z}^\kappa$ such that v_j is the same modulo q as $\text{First}_\kappa(B_j)$, i.e., as $\text{First}_\kappa(Gb_j) + \text{First}_\kappa(d_j)$.

Define L as the set of all $(\alpha, \epsilon) \in R \times \mathbb{Z}^\kappa$ such that $\text{First}_\kappa(G\alpha)$ is the same as ϵ modulo q . This is a lattice of full rank $d = n + \kappa$ and determinant q^κ .

This lattice contains a vector close to $(0, v_j)$, namely $(b_j, v_j - \text{First}_\kappa(d_j))$, since $\text{First}_\kappa(Gb_j)$ is the same as $v_j - \text{First}_\kappa(d_j)$ modulo q . To quantify “close”: the difference is $(b_j, -\text{First}_\kappa(d_j))$, a vector consisting of coefficients that were generated independently at random from χ .

Now use the algorithm from [48, Section 5, third paragraph], with the CVP subroutine from [46], to try to find this lattice vector:

¹² The same paper [46, Section 7] also considered the special case of “bounded-distance decoding” (BDD), the case that the target is guaranteed to have a particular closeness to the lattice, as in the problem of decrypting a ciphertext. The paper found that the algorithm obtained only limited benefit from this guarantee as the distance bound decreases, and said “An open problem would be to adapt the iterative slicer to make better use of this guarantee”. One solution to this problem is as follows: use the slicer inside the CVP algorithm from [46], plug the CVP algorithm into the algorithm from [48, Section 5, third paragraph], and then allow the block size in that algorithm to drop in the usual way with the distance bound.

- As precomputation for [48], apply BKZ- β to a basis for L . This costs, heuristically, $(3/2 + o(1))^{\beta/2}$. This work is shared across all j .
- Apply the precomputation step from [46] to the β -dimensional projection of L , computing a database of all short vectors in that projection. (Some versions of BKZ- β produce this database as a side effect.) This also has heuristic cost $(3/2 + o(1))^{\beta/2}$, again shared across all j .
- For each j , apply the main step from [46] to find a vector in the projection of L close—hopefully closest—to the projection of $(0, v_j)$. The heuristic analysis of [46] says that this costs just $2^{(0.234+o(1))\beta}$; this is not a bottleneck if there are fewer than $2^{(0.058+o(1))\beta}$ targets.
- “Finish with Babai’s algorithm” as in [48], hopefully finding the desired $(b_j, v_j - \text{First}_\kappa(d_j))$.

The existing heuristics say that this works if the difference $(b_j, -\text{First}_\kappa(d_j))$ has length below $(d/\beta)^{1/2}\Theta(\beta)^{1-d/2\beta}(\det L)^{1/d}$; again, this is already stated in [48], modulo the issues noted above regarding d/β .

At the level of detail of first-order asymptotics, this is the same inequality as in Section 2.8, so the smallest β that heuristically works here is again $(2Q_0/(Q_0 - S_0 + 1/2)^2 + o(1))n$. The heuristic cost is the same $(3/2 + o(1))^{\beta/2} = 2^{((\log_2(3/2))Q_0/(Q_0 - S_0 + 1/2)^2 + o(1))n}$, if there are at most $2^{(0.058+o(1))\beta}$ targets. The critical difference is that this attack heuristically breaks *all* of the targets for this total cost, whereas Section 2.8 was for just one target.

In particular, if n is chosen so that the usual single-target message-recovery attack has heuristic cost $2^{(1+o(1))\lambda}$, then this attack breaks T targets, for any $T \leq 2^{(0.058+o(1))\beta} = 2^{(0.19\dots+o(1))\lambda}$, with total heuristic cost $2^{(1+o(1))\lambda}$.

4 Asymptotic heuristic cost of breaking one out of many ciphertexts

This section analyzes what the existing heuristics say about an attack that decrypts one out of many ciphertexts.¹³ The conclusion, in short, is that one can heuristically reduce the block size β from Section 3.3 by $\Theta(n/\lg n)$ for breaking one out of T ciphertexts, provided that $\lg T \in \Theta(n)$.

The Θ constant in $\Theta(n/\lg n)$ depends on how many ciphertexts are being attacked and, more subtly, on details of how the error distribution χ is chosen. As a numerical example, consider the following conservative choices:

- Choose χ as the binomial distribution on $\{-16, \dots, 16\}$, exactly as in the original New Hope paper. Then $S_0 = 0$ and $S_1 = 3/2$.
- Since the literature says that smaller moduli are safer, and since any Q_0 above $1/2$ together with error correction suffices to asymptotically eliminate decryption failures (given that $S_0 = 0$), choose $Q_0 = 0.501$.

¹³ A decoding-one-out-of-many attack against code-based cryptography was dubbed a “DOOM” attack in [91], but using the title “Lattice DOOM” for this paper would be unnecessarily theatrical.

One can then heuristically reduce the block size β by $(0.60537\dots + o(1))\beta/\lg\beta$ for breaking one out of $2^{(0.0574+o(1))\lg\beta}$ ciphertexts. This number of ciphertexts fits inside the per-ciphertext budget from Section 3, since $0.0574 < 0.058$. The calculation of $0.60537\dots$ is explained in detail below.

This is a larger heuristic improvement than the heuristic “dimensions for free” improvements reported in [45] and [44] for SVP- β . Consequently, the standard message-recovery attack *plus* the improvements from [45] and [44] is, for this conservative (χ, Q_0) , heuristically outperformed by this one-out-of-many attack.

Qualitatively, compared to the security consequences of an attack decrypting many ciphertexts, there are different, often stronger, security consequences of a faster attack decrypting one out of many ciphertexts. For example, consider what the attacks say regarding security levels:

- Section 3 illustrates the danger of the “attacker economist” philosophy, choosing the security level λ so that 2^λ is slightly beyond the attacker’s benefit of breaking *one* ciphertext.
- This section illustrates the danger of choosing the security level λ so that 2^λ is slightly beyond a feasible computation.

As another example of the consequences, consider again the statement that “ $\text{Adv}_{\text{PKE}}^{\text{IND-CPA}} \approx \text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}$ ”. If the time limit on attacks is too low for a single-target attack but is high enough for a faster one-out-of- q_C attack, then the first Adv is essentially 0 and the second Adv is far above 0 (where exact numbers depend on probabilistic aspects of the attacks), contradicting the statement. For comparison, an attack breaking *all* targets is at least as expensive as a single-target attack, so it cannot similarly contradict the same statement.

4.1. Designing one-out-of-many algorithms. There is a large overlap from the algorithm designer’s perspective between solving many problems and solving one of many problems, despite the differences in security consequences.

In particular, starting from an algorithm to solve many problems—especially an algorithm that performs a shared precomputation and then handles each problem separately—one can usually save time by simply reducing algorithm parameters. There is then nothing new to design; what matters is the algorithm analysis. Often it turns out that the reduced algorithm continues to successfully break one or more of the problems, because of randomness in the problems, randomness in the algorithm, or both. In the simple AES-key-guessing example from Section 1, guessing only 2^{88} keys has a good chance (more precisely, probability close to $1 - \exp(-1) = 0.632\dots$) of colliding with one of the 2^{40} keys chosen at random by users.

In the T -target attack from Section 3.3, reducing the block size β correspondingly reduces the heuristic cost of the initial BKZ- β computation and the precomputation of a database of short vectors in a β -dimensional lattice; under a reasonable assumption on T , other steps are not heuristic bottlenecks. The critical question is then how far one can reduce β so that the resulting algorithm, when applied to one target, still has success probability at least $1/T$.

There are large gaps in the literature at this point. It is clear from small-scale experiments that there are probabilistic effects in lattice attacks;¹⁴ however, for many of the necessary subroutines, the heuristic analyses in the literature consider only the time taken for high success probability, rather than the whole time-probability tradeoff curve. This is a defensible simplification, but it poses an obstacle for the analysis of multi-target attacks.¹⁵ It seems necessary to revisit a large part of the lattice-attack literature, systematically filling in every missing probability analysis and building new heuristics accordingly.

Fortunately, the standard analysis has one step with two features that are important for this paper:

- A probabilistic analysis of that step can be rigorously carried out as a separate module. This is the main technical content of this section.
- Plugging the results for probability $1/2^{(c+o(1))n}$ into the rest of the standard analysis obtains a block size $\Theta(n/\lg n)$ smaller than before: more precisely, $(d + o(1))n/\lg n$ smaller, with d quantified below.

This makes no changes in the existing heuristics. The analysis here applies the existing heuristics to a modified PKE; each ciphertext in the original PKE has a noticeable chance, rigorously quantified below, of matching a ciphertext in the modified PKE.

4.2. Computing 2-norm distributions from error distributions. Take any probability distribution χ supported on a finite set of integers such that 0 and 1 each have nonzero probability and the average of χ is 0. Let n be a positive integer. Let v be an element of \mathbb{Z}^n with each entry drawn independently at random from χ .

The goal here is to analyze the chance that v has a particularly small 2-norm. Section 4.4 will replace n with $n + \kappa$ to see the chance of a ciphertext being particularly weak in the context of attacks against the LPR PKE.

Define g as the polynomial $\sum_{i \in \mathbb{Z}} \chi_i z^{i^2} \in \mathbb{R}[z]$. Then, for each j , the generating function of v_j^2 is g . This means that the chance of $v_j^2 = \ell$ is the coefficient of z^ℓ in g . See [98] for an introduction to generating functions.

One of the useful properties of generating functions is that adding independent random variables corresponds to multiplying generating functions. In particular, the generating function of $\sum_j v_j^2$ is g^n : the chance of $\sum_j v_j^2 = \ell$ is the coefficient of z^ℓ in g^n .

Fix a positive rational number y below the degree of g . If $yn \in \mathbb{Z}$ then the coefficient of z^{yn} in g^n is $(c + o(1))n^{-1/2}(g(\rho)/\rho^y)^n$; here ρ is the unique positive

¹⁴ Consider, e.g., the first batch of experiments in [9, Table 1], reporting that probabilities dropped from 93.3% to 52.8% to 4.8% as β dropped from the cutoff mandated by the standard analysis to 5 or 10 lower.

¹⁵ Even for single-target attacks, understanding the cost of lower-probability attacks is important. Presumably many users would be unhappy to hear that a feasible attack succeeds with probability 10%.

real number satisfying $\rho g'(\rho) = yg(\rho)$, and

$$c = \left(2\pi \left(\frac{g(\rho)g''(\rho) - g'(\rho)^2}{g(\rho)^2} + \frac{y}{\rho^2} \right) \right)^{-1/2}.$$

The previous sentence is an example of [49, Proposition VIII.8], which in turn is an application of the saddle-point method in analytic combinatorics. For the asymptotics in this paper, the important factor is $(g(\rho)/\rho^y)^n$; but the extra precision of $(c+o(1))n^{-1/2}(g(\rho)/\rho^y)^n$ is used in the sanity check described below.

Take, for example, χ as the binomial distribution on $\{-16, \dots, 16\}$, and take $y = 8$. Then $g(1) = 1$, $g'(1) = 8$, and $g''(1) = 180$; $\rho = 1$ is a root of $\rho g'(\rho) = yg(\rho)$, the unique positive root; $g(\rho)/\rho^y = 1$; and $c = (248\pi)^{-1/2} = 0.035826\dots$. The coefficient of z^{8n} in g^n , the chance of $\sum_j v_j^2$ being exactly its average value, is $((248\pi)^{-1/2} + o(1))n^{-1/2}$.

The script in Figure 4.2.1 uses the Sage [90] mathematics system to carry out the same computation for the same χ but with $y = 21/4 = 5.25$. The script uses Sage's built-in interval arithmetic to track rounding errors, printing question marks for unknown digits. The output of the script indicates that $\rho = 0.96713\dots$, $g(\rho)/\rho^y = 2^{-0.057241\dots}$, and $c = 0.054093\dots$. The coefficient of z^{yn} in g^n is thus $2^{(-0.057241\dots+o(1))n}$ if yn is an integer.

As a sanity check, the script also runs through various small n where yn is an integer, printing out tuples $(n, yn, (g^n)_{8n}, (g^n)_{8n}/cn^{-1/2}(g(\rho)/\rho^y)^n$. The ratio here, the last tuple element, is $1 + o(1)$ according to the proposition cited above, meaning that *if* the script is performing the right computations then the printed ratios are also $1 + o(1)$. Concrete examples of the printed ratios include 0.87716... for $n = 8$; 0.99683... for $n = 16$; 0.99442... for $n = 32$; 0.99718... for $n = 64$; 0.99859... for $n = 128$; 0.99929... for $n = 256$; and 0.99964... for $n = 512$. A proper test, rather than just a sanity check, would require explicit bounds on the $o(1)$ in $c + o(1)$.

For comparison, naively modeling the distribution of the squared 2-norm of v as a normal distribution with average $8n$ and variance $124n$ would say that squared 2-norm $\leq (21/4)n$ appears with probability $(1 - \operatorname{erf}((11/4)\sqrt{n/248}))/2 \in \exp(-((11/4)^2/248 + o(1))n) = 2^{(-(11/4)^2/(248 \log 2) + o(1))n} = 2^{(-0.043993\dots+o(1))n}$, overestimating the actual $2^{(-0.057241\dots+o(1))n}$ chance by an exponential factor.

Changing y to $5511515/1048576 = 5.2561\dots$ in the script produces $\rho = 0.96724\dots$, $g(\rho)/\rho^y = 2^{-0.056943\dots}$, and $c = 0.054030\dots$; unsurprisingly, these are close to the results for $y = 5.25$. For this value of y , the coefficient of z^{yn} in g^n is $2^{(-0.056943\dots+o(1))n}$ if yn is an integer. The script skips the sanity check in this case given the size of 1048576.

What follows is a simpler example that can be checked by hand for variable y —assuming Stirling's well-known asymptotic formula for factorials. Stirling's formula can be proven as an example of the saddle-point method, but can also be proven more directly as in [89].

Take χ as the uniform distribution on $\{-1, 0, 1\}$, and take any rational number y with $0 < y < 1$. Then $g = (1 + 2z)/3$. The equation $\rho g'(\rho) = yg(\rho)$ is $\rho(2/3) = y(1 + 2\rho)/3$, i.e., $2(1 - y)\rho = y$, which has a unique positive root $\rho = y/2(1 - y)$.

```

QQx.<x> = QQ[]

s2 = 16
B = (1+x)^(2*s2)/2^(2*s2)
g = sum(B[i+s2]*x^(i^2) for i in range(-s2,s2+1))

y = 21/4

assert all(gi >= 0 for gi in list(g))
assert g[0] > 0
assert g[1] > 0
assert y in QQ
assert y > 0
assert y < g.degree()

g1 = g.diff()
g2 = g1.diff()
foundrho = False
for rho in (x*g1-y*g).roots(RIF,multiplicities=False):
    if rho > 0:
        foundrho = True
        break
assert foundrho
base = g(rho)/rho^y
xi = (g(rho)*g2(rho)-g1(rho)^2)/g(rho)^2+y/rho^2
c = 1/(rho*sqrt(2*RIF(pi)*xi))

print('y',y)
print('rho',rho)
print('base',base)
print('lgbase',log(base,2.0))
print('lgnaive',-RIF((g1(1)-y)^2/(2*(g2(1)+g1(1)-g1(1)^2)*log(2))))
print('c',c)
sys.stdout.flush()

D = QQ(y).denominator()
if D < 1024:
    g = g.change_ring(RIF)
    gD = g^D
    n,gn = D,gD
    while n < 1024:
        est = c*base^n/sqrt(RIF(n))
        print('n',n,'yn',y*n,'coeff',gn[y*n],'coeff/est',gn[y*n]/est)
        sys.stdout.flush()
        n,gn = n+D,gn*gD

```

Fig. 4.2.1. Computing the asymptotic probability of a specific 2-norm of an error vector; and computing the exact probability for various n as a sanity check. See text for details.

Here $g(\rho)/\rho^y = ((1 + y/(1 - y))/3)/(y/2(1 - y))^y = 2^y/3y^y(1 - y)^{1-y}$, so the coefficient of z^{yn} in g^n is $(2^y/3y^y(1 - y)^{1-y})^n/n^{1/2+o(1)}$ when yn is an integer.

The check, given that this g is linear, is to use the binomial formula for g^n to see that the coefficient of z^i in g^n is $\binom{n}{i}2^i/3^n$. Equivalently, observe that there are exactly $\binom{n}{i}2^i$ ways for a vector $v \in \{-1, 0, 1\}^n$ to have i nonzero positions, i.e., squared 2-norm i . By Stirling's formula, the chance $\binom{n}{i}2^i/3^n$ is $(2^y/3y^y(1 - y)^{1-y})^n/n^{1/2+o(1)}$ where $y = i/n$.

4.3. A modified cryptosystem. Fix a probability distribution χ . Assume as above that χ is supported on a finite set of integers, that 0 and 1 each have nonzero probability, and that the average of χ is 0. Fix a rational number y with $0 < y < s^2$, where s is the standard deviation of χ . Also fix a rational number $Q_0 > 1/2$, and define $W_0 = (Q_0 - 1/2)/(Q_0 + 1/2)$; then $0 < W_0 < 1$.

Consider a cryptosystem SqueezedLPR that is identical to the LPR PKE with distribution χ except for two restrictions (to simplify the analysis) and one critical change. The two restrictions are that W_0n and $y(1 + W_0)n$ are integers. The change is that encryption rejects (b, d) and starts over¹⁶ unless the sum of the squares of the first $(1 + W_0)n$ components of the $2n$ -component vector (b, d) is exactly $y(1 + W_0)n$.

Now apply the standard heuristic analysis to SqueezedLPR, specifically the analysis of the usual single-target message-recovery attack, specifically with the attack parameter κ chosen as W_0n .

For LPR, the squared 2-norm of the target vector $(b, \text{First}_\kappa(d), 1)$ was estimated as $(n + \kappa)s^2 + 1$. For SqueezedLPR, the squared 2-norm of the target vector $(b, \text{First}_\kappa(d), 1)$ is exactly $y(1 + W_0)n + 1 = (n + \kappa)y + 1$, in effect replacing s^2 with y . The standard analysis proceeds identically aside from this difference.

Now fix Q_1 , and fix a function $n \mapsto q$ with $q \in \mathbb{Z}$, $q \geq 2$, and $\lg q \in Q_0 \lg n + Q_1 + o(1)$. Asymptotically, SqueezedLPR is just like LPR with cryptosystem parameters (n, q, χ) in that the asymptotics of the standard β inequality are covered by Theorem 2.6.1, specifically with $S_0 = 0$ since the error distribution is independent of n . The only difference is that S_1 drops from $\lg s$ to $(\lg y)/2$.

The optimized parameters in Section 2.7 have $Z_0 = 2Q_0/(Q_0 + 1/2)^2$ and $Z_1 = (2S_1 + \dots)2Q_0/(Q_0 + 1/2)^3$. The reduction of S_1 by $\lg s - (\lg y)/2$ reduces Z_1 by $(2 \lg s - \lg y)2Q_0/(Q_0 + 1/2)^3$, and thus reduces β by

$$\frac{((2 \lg s - \lg y)2Q_0/(Q_0 + 1/2)^3 + o(1))n}{\lg n} = \frac{((2 \lg s - \lg y)/(Q_0 + 1/2) + o(1))\beta}{\lg \beta}$$

where the equation comes from $\beta \in (Z_0 + o(1))n$ and $2Q_0/(Q_0 + 1/2)^3 = Z_0/(Q_0 + 1/2)$.

For the multi-ciphertext attack in Section 3.3, the existing heuristics say that the requirement on β is the same inequality—aside from a $1 + o(1)$ factor, which has no effect on the first-order or second-order asymptotics. Again this change of cryptosystem reduces the block size β by $((2 \lg s - \lg y)/(Q_0 + 1/2) + o(1))\beta/\lg \beta$.

¹⁶ There are much more efficient ways to sample the resulting distribution, but the speed of encryption has no relevance to the standard attack analysis.

4.4. Attacks against the modified cryptosystem as attacks against the original cryptosystem. Consider again LPR with distribution χ , and assume that W_0n and $y(1 + W_0)n$ are integers.

An LPR ciphertext has a chance of being a SqueezedLPR ciphertext; this chance is quantified in Section 4.2. In a large enough collection of LPR ciphertexts, one expects to find some SqueezedLPR ciphertexts. The standard analysis of SqueezedLPR says, under the existing heuristics, that a smaller block size β for the multi-ciphertext attack breaks those SqueezedLPR ciphertexts; this is quantified in Section 4.3. Putting this together gives a one-out-of-many-ciphertexts attack against LPR, heuristically faster than the original multi-ciphertext attack.

Here is a numerical example of the asymptotic impact:

- Take $Q_0 = 0.501$. Then $W_0 = 1/1001$ and $Z_0 = 1002000/1002001$.
- Take χ as the binomial distribution on $\{-16, \dots, 16\}$. This has standard deviation $s = 2^{3/2}$.
- Take $y = 5511515/1048576$. Restrict attention to n such that $(1 + W_0)n$ and $y(1 + W_0)n$ are integers.
- The chance of the first $(1 + W_0)n$ entries of (b, d) having squared 2-norm exactly $y(1 + W_0)n$ is $2^{(-0.056943\dots + o(1))(1+W_0)n} = 2^{(-0.056999\dots + o(1))n} = 2^{(-0.056999\dots + o(1))\beta}$. This is simply replacing n with $(1 + W_0)n$ in one of the examples from Section 4.2, and then rewriting n as $(1/Z_0 + o(1))\beta$.
- In a collection of, say, $2^{(0.0574 + o(1))\beta}$ LPR ciphertexts, there will be, with probability $1 - o(1)$, at least one ciphertext meeting this condition—in other words, a SqueezedLPR ciphertext.¹⁷ This is also a small enough number of ciphertexts to not be a heuristic bottleneck in the attack analyzed in Section 3, since 0.0574 is safely below 0.058.
- Finally, $(2 \lg s - \lg y)/(Q_0 + 1/2)$ is $0.60537\dots$. The analysis from Section 4.3 says that reducing β by $(0.60537\dots + o(1))\beta/\lg \beta$ suffices to break the SqueezedLPR ciphertexts, according to the existing heuristics.

The gap between $0.056999\dots$ and the 0.058 from [46] means that $0.60537\dots$ can be increased, but finding the exact limit would require a tedious recalculation of the 0.058; scripts for [46] do not seem to be publicly available. The choice of 0.0574 here is designed to protect this paper against the possibility that 0.058 actually means something as small as 0.0575. As a separate matter, readers who prefer n to be a power of 2 can instead take $Q_0 = 1025/2046$ with $W_0 = 1/1024$; this marginally changes the other numbers listed above.

The possibility of distance variations being important for security was already pointed out in the NTRU Prime proposal [31, page 41] in 2019: “The analysis does not take into account the variations in the 2-norm of g ; this simplification could overestimate or underestimate security”. However, that was not the same as pointing out the possibility of exploiting these variations for an efficient multi-ciphertext attack, let alone analyzing the resulting impact.

¹⁷ In fact, one expects an exponential number of SqueezedLPR ciphertexts, so this is a some-out-of-many attack. More precisely, one expects to find more than $2^{(0.0004 + o(1))\beta}$ SqueezedLPR ciphertexts.

4.5. A sanity check on block sizes. The conclusion that the heuristically required β drops by $((2 \lg s - \lg y)/(Q_0 + 1/2) + o(1))\beta/\lg \beta$, when the variance drops from s^2 to y , relies on a chain of asymptotic calculations, including Theorem 2.6.1 and the optimization in Section 2.7. As a sanity check on these calculations (analogous to the sanity check in Section 4.2 on the calculations of how often the squared 2-norm is yn), the following paragraphs ask what the same drop in variance does to concrete values of β , according to an estimator already in the literature that calculates β for concrete (n, q, s) .

Specifically, consider [6], which reported security estimates for a wide range of NIST submissions, specifically by plugging a variety of estimates for the cost of BKZ- β into an estimator using the standard heuristic analysis. It is straightforward to run the same estimator with different choices of variance to see how the variance affects the choice of β for concrete examples of (n, q) .

Take, in particular, $n \in \{2^5, 2^{10}, 2^{15}, 2^{20}\}$ and $q \in \{\lfloor 64n^{0.501} \rfloor, 64n\}$. These are realistic sizes of q for $n = 1024$:

- FrodoKEM-976 uses the modulus $64 \cdot 1024 = 65536$.
- Back-of-the-envelope calculations suggest that using BCH codes as in LAC [74] should make a modulus as small as $64\sqrt{1024} = 2048$ work with the New Hope error distribution.

Extrapolating from these sizes of q for $n = 1024$ to formulas for q that differ only in the exponent of n , namely $64n^{0.501}$ and $64n$, corresponds to purely changing Q_0 . Further motivations for considering $Q_0 = 0.501$ and $Q_0 = 1$ are mentioned elsewhere in this paper. More to the point, these two values of Q_0 are separated by enough to easily see the impact on $1/(Q_0 + 1/2)$.

Table 4.5.1 reports, for each of these pairs (n, q) , what happens to β when variance drops from 8 to $5511515/1048576$, according to the estimator used in [6], with κ allowed to be as large as $2n$. The “reldrop” column in the table is the drop in β divided by $\beta/\lg \beta$. For comparison, the asymptotics say that “reldrop” is asymptotically $0.605\dots + o(1)$ for $Q_0 = 0.501$, and about $2/3$ as large, asymptotically $0.403\dots + o(1)$, for $Q_0 = 1$. The concrete numbers in the table are not far from this.

The numbers in the table are printed out by the Sage script in Figure 4.5.2. As a check that it is using the estimator correctly, the script also prints out (β, κ) for the original New Hope parameters $(n, q, s) = (1024, 12289, \sqrt{8})$ and the revised new Hope parameters $(n, q, s) = (1024, 12289, 2)$, namely $(968, 1071)$ and $(886, 1019)$ respectively. This matches the $(886, 1019)$ produced by the the “LWE $2n$ samples” NewHope-1024 script provided by [6], and the $(968, 1071)$ produced if the script is modified to say `sd = sqrt(8)` instead of `sd = 2`.

A slightly smaller $\beta = 967$ is claimed for the original New Hope in [13, Table 1] and in [11, Table 12], but those attack parameters do not satisfy the inequality stated in [13]. A closer look at the literature shows that there are actually variations in the standard analysis, and presumably these variations are the source of the discrepancy between 967 and 968. Consider, e.g., the $(a, e, 3.14)$ variant mentioned in Section 2.4; as another example, the literature sometimes replaces δ with $\delta^{d/(d-1)}$.

n	q	β_0/n	κ_0/n	β_1/n	κ_1/n	reldrop
32	363	1.25000	0.90625	1.25000	0.75000	0.00000
1024	2062	1.22461	1.03809	1.14941	1.02930	0.63198
32768	11706	1.25000	0.72910	1.20096	0.69708	0.60113
1048576	66450	1.22636	0.55642	1.19031	0.53552	0.59660
32	2048	1.25000	0.53125	1.25000	0.46875	0.00000
1024	65536	0.76465	1.08887	0.72949	1.03320	0.44197
32768	2097152	0.86414	0.83578	0.83987	0.81613	0.41523
1048576	67108864	0.89010	0.71349	0.87176	0.69796	0.40855

Table 4.5.1. The impact of changes in variance upon block size for various pairs (n, q) , according to the estimator used in [6]. The β_0/n and κ_0/n columns are for variance 8. The β_1/n and κ_1/n columns are for variance 5511515/1048576. The “reldrop” column is $(\beta_0 - \beta_1)/(\beta_0/\lg \beta_0)$. The last five columns are rounded to 5 digits after the decimal point.

The table also reports the values of β/n and κ/n from the same estimator. For comparison, the asymptotics say that β/n and κ/n converge to the aforementioned 1002000/1002001 and 1/1001 respectively for $Q_0 = 0.501$, and to 8/9 and 1/3 respectively for $Q_0 = 1$. The concrete numbers in the table for both β/n and κ/n seem to be considerably above these limits (except that β/n seems close for $Q_0 = 1$), motivating the following investigation of larger n .

The estimators in the literature become extremely slow as n increases. However, for any particular (n, q, s) , one can instead minimize β on the curve of *real* pairs $(\kappa, \beta) \in \mathbb{R} \times \mathbb{R}$ for which $\varphi(\kappa, \beta) = 0$, where φ is chosen so that the standard inequality is $\varphi(\kappa, \beta) < 0$. Lagrange multipliers convert this minimization problem into the problem of solving the two equations $\varphi(\kappa, \beta) = 0$ and $\varphi_x(\kappa, \beta) = 0$ where φ_x is the derivative of φ with respect to its first variable. Figure 4.5.3 is a Sage script that finds, to reasonably high precision, a solution at reasonable speed by well-known equation-solving techniques.

For, e.g., $(n, q, s) = (1024, 65536, \sqrt{8})$, the output of this script says that $\beta/n \approx 0.76508$ and $\kappa/n \approx 1.10203$. This is a slightly larger value of β/n than the 0.76465 (i.e., $\beta = 783$) reported in Table 4.5.1 from the estimator used in [6], whereas a minimum over $\mathbb{R} \times \mathbb{R}$ cannot be larger than a minimum of the same function over $\mathbb{Z} \times \mathbb{Z}$. One possible explanation would be that the local minimum found by Figure 4.5.3 is not a global minimum; a simpler explanation is the tiny variations mentioned above in the inequalities considered in the literature.

For $n = 2^{2048}$, again with $q = 64n$ and $s = \sqrt{8}$, the same script outputs $(\beta/n, \kappa/n) = (0.88945\dots, 0.33709\dots)$, much closer to the (8/9, 1/3) limit. For $n = 2^{2048}$, $q = \lfloor 64n^{0.501} \rfloor$, and $s = \sqrt{8}$, the script outputs $(\beta/n, \kappa/n) = (1.00343\dots, 0.00665\dots)$, almost as close to the (0.99999\dots, 0.00099\dots) limit. Finally, the relative drops in β for $n = 2^{2048}$ are 0.40394\dots for $Q_0 = 1$ and 0.60513\dots for $Q_0 = 0.501$, very close to the limits.

4.6. Countermeasures. Switching from the binomial distribution on $\{-16, \dots, 16\}$ to the binomial distribution on $\{-8, \dots, 8\}$, as in newer versions

```

load('estimator.py')

def bkzmodel(beta,d,B):
    return (3/2)^RR(beta/2)

def betakappa(n,q,V):
    alpha = sqrt(2*pi*V)/RR(q)
    samplesavailable = 2*n # limiting kappa
    est = primal_usvp(n,alpha,q,m=samplesavailable,
                    success_probability=0.99,
                    secret_distribution='normal',
                    reduction_cost_model=bkzmodel)
    return est['beta'],est['m']

for n,q,V in (1024,12289,8),(1024,12289,4):
    beta,kappa = betakappa(n,q,V)
    print('n',n,'q',q,'V',V,'kappa',kappa,'beta',beta)
    sys.stdout.flush()

for j in 5,10,15,20:
    n = 2^j
    for q in floor(64*n^0.501),64*n:
        betas = []
        for V in 8,5511515/1048576:
            beta,kappa = betakappa(n,q,V)
            print('n',n,'q',q,'V',V,'kappa/n',RR(kappa/n),'beta/n',RR(beta/n))
            sys.stdout.flush()
            betas += [beta]
        reldrop = RR((betas[0]-betas[1])/(betas[0]/log(betas[0],2.0)))
        print('n',n,'q',q,'reldrop',reldrop)
        sys.stdout.flush()

```

Fig. 4.5.2. A Sage script that calculated the numbers in Table 4.5.1, using the same estimator used in [6].

of New Hope, and adjusting y appropriately, reduces the above 0.60537... to 0.59739.... Switching from 8 to 3, one of the distributions used in the latest version of Kyber, produces 0.56984.... Switching from 3 to 2, another distribution used in the latest version of Kyber,¹⁸ produces 0.54668.... These numbers were calculated by the Sage script in Figure 4.6.1.

¹⁸ A full analysis of Kyber would also have to account for the rounding in Kyber ciphertexts. This rounding is mostly in the C component but also a little in the B component, slightly increasing the effective size of d . The latest Kyber documentation claims that the rounding in the latest version of Kyber-512 gains several bits of security for attacks against the ciphertexts. The concrete question—not addressed by this paper’s asymptotics; see Section 1.3—is then the extent to which security is damaged by variations in the effective size of (b, d) .

```

x,y = var('x y')
for Q0 in 1,0.501:
  for V in 8,5511515/1048576:
    for j in 10,15,20,25,32,64,128,256,512,1024,2048:
      RR = RealField(2*j+50)
      n = 2^j
      q = floor(64*n^Q0)
      d = n+x+1
      delta = (y*(pi*y)^(1/y)/(2*pi*exp(1)))^(1/(2*(y-1)))
      rho = ((n+x)*V+1)^(1/2)
      rho /= (d/y)^(1/2)*delta^(2*y-d-1)*q^(x/d)
      g = log(rho)
      gx = g.derivative(x)
      gy = g.derivative(y)
      gxx = gx.derivative(x)
      gxy = gx.derivative(y)

      ok = False
      while not ok:
        a = RR.random_element(1/10,19/10)*n
        b = RR.random_element(n/10,9*n/10+a)

        for loop in range(30):
          gab = RR(g(x=a,y=b))
          gxab = RR(gx(x=a,y=b))
          if abs(gab) < 1/2^30 and abs(gxab) < 1/2^30:
            ok = True
            break
          gyab = RR(gy(x=a,y=b))
          gxxab = RR(gxx(x=a,y=b))
          gxyab = RR(gxy(x=a,y=b))
          J = matrix([[gxab,gyab],[gxxab,gxyab]])
          d,e = J.solve_right(vector([gab,gxab]))
          step = 1
          while True:
            a1,b1 = a-step*d,b-step*e
            if a1 > 0 and b1 > 0 and a1 < 2*n and b1 < a1+n:
              a,b = a1,b1
              break
          step *= RR(1/2)

      print('n 2^%d Q0 %.8f V %s beta/n %.8f kappa/n %.8f'
            % (j,Q0,V,b/n,a/n))

```

Fig. 4.5.3. A script that, for various (n, q, s) , finds real numbers (β, κ) to approximately locally minimize β subject to $((n + \kappa)s^2 + 1)^{1/2} = (d/\beta)^{1/2}\delta^{2\beta-d-1}q^{\kappa/d}$.

```

batchtarget = 0.057 # safely below 0.058
QQz.<z> = QQ[]

def reldrop(g,Q0):
    W0 = (Q0-1/2)/(Q0+1/2)
    Z0 = 2*Q0/(Q0+1/2)^2
    g = QQz(g)
    assert all(gi >= 0 for gi in list(g))
    assert g[0] > 0 and g[1] > 0 and g(1) == 1
    g1 = g.diff()
    g2 = g1.diff()
    V = g1(1)

def rhoy(y):
    return max((z*g1-y*g).roots(ring=RR,multiplicities=False))

def grhorhoy(y):
    rho = rhoy(y)
    assert rho > 0
    return g(rho)/rho^y

def yevaluation(y):
    # require y(n+kappa) = sum of squares of n+kappa positions
    lgbasenk = -log(grhorhoy(y))/log(2.0)
    # probability is 2^((-lgbasenk+o(1))*(n+kappa))
    lgbasen = (1+W0)*lgbasenk # n+kappa is (1+W0+o(1))n
    # probability is 2^((-lgbasen+o(1))*n)
    lgbasebeta = lgbasen/Z0 # beta is (Z0+o(1))n
    # probability is 2^((-lgbasebeta+o(1))*beta)
    return lgbasebeta-batchtarget

y = find_root(yevaluation,0.001,V)
y = RR(ceil(y*1048576)/1048576)
return log(V/y,2.0)*Z0^2/(1+W0)

for Q0 in 0.501,1:
    for c in 16,8,3,2:
        g = sum(binomial(2*c,j)*z^((j-c)^2)
                for j in range(2*c+1))/4^(c)
        print('binomial -%d,...,%d Q0 %.8f reldrop %.8f'
              %(c,c,Q0,reldrop(g,Q0)))

    g = (1+2*z)/3
    print('uniform -1,0,1 Q0 %.8f reldrop %.8f'%(Q0,reldrop(g,Q0)))

```

Fig. 4.6.1. A script to calculate the asymptotic drops of the heuristically required β , relative to $\beta/\lg \beta$, for various χ and $Q_0 \in \{0.501, 1\}$, with $S_0 = 0$, with batch size fitting into at most $2^{(0.057+o(1))\beta}$ ciphertexts.

More interestingly, switching to the uniform distribution on $\{-1, 0, 1\}$ produces just $0.32860\dots$, smaller than the constants from [45] and [44]. Compared to New Hope, this distribution has much smaller squared 2-norms on average, but considerably less variation *relative to the average*: i.e., the variation in heuristic security from one ciphertext to another is lower.

There is even less 2-norm variation in cryptosystem proposals that take, e.g., fixed-weight ternary vectors b and d . More generally, starting from any χ , one can efficiently require a specific 2-norm (or a narrow range of 2-norms), limiting the size variations exploited in this section. There is still *some* variation: $\text{First}_\kappa(d)$ selects some positions from d , and the attacker can vary which positions are taken.

Ternary distributions with a prespecified number of 1 entries and a prespecified number of -1 entries were already highlighted in the 1998 NTRU proposal [56] as an attractive choice. Subsequent analysis consistently indicated that proposed wider-than-ternary distributions damaged security against known attacks for any given key size, as noted in, e.g., [31, Section 4.6].

New Hope’s use of a wide close-to-normal distribution on $\{-16, \dots, 16\}$ can be traced to the advertising of “worst-case-to-average-case reductions” for Ring-LWE in [75]. On the other hand, New Hope is still not large enough for the worst-case-to-average-case reductions to apply; see generally [69]. Descendants of New Hope typically returned to smaller distributions, putting more emphasis on security against known attacks—but often avoiding ternary distributions, for example because of a common misimpression that “hybrid” attacks apply only to ternary distributions.

Changing error distributions is not the only way to reduce the frequency of particularly weak ciphertexts. Increasing Q_0 , while leaving the vector distribution unchanged, generally reduces the variations that one can expect to see within $2^{(0.058+o(1))\beta}$ ciphertexts. Some cryptosystem proposals already choose larger q (and correspondingly larger n to reach the same heuristic security level), asymptotically with $Q_0 = 1$, for another reason: namely, to provide proofs that there are no decryption failures.

Instead of, or as a supplement to, trying to reduce variations, one can modify the standard analysis to account for those variations, increasing parameter sizes to protect the weakest ciphertexts. The obvious approach is as follows: decide a limit on the number of ciphertexts, decide a limit on the acceptable success probability of a multi-ciphertext attack, use generating-function techniques to calculate the smallest 2-norm that will be found with that probability among that number of ciphertexts, and use this 2-norm in place of an “expected” 2-norm.

5 Reconciliation with previous work

This section traces various quotes from the literature that seem likely to have led many readers to believe that multi-target attacks against lattice KEMs are no more efficient than single-target attacks. To the extent that the sources appear to be stating arguments for this belief, this section explains how this paper’s

analysis exploits gaps in those arguments. The deepest part of this section is identifying flaws in a provable-security argument in [47]; the other gaps are easier to see.

5.1. New Hope 2015. A subsection labeled “all-for-the-price-of-one attacks” in the New Hope paper [13] expressed concern about the possibility¹⁹ of attacks that break T targets (“compromise *all* communications”) for the cost of breaking 1 target. The paper continued by saying that “all those pitfalls can be avoided by having the communicating parties generate a fresh \mathbf{a} at each instance of the protocol (as we propose)”.

The argument that there are no all-for-the-price-of-one attacks was not stated explicitly, but appears to be the following:

- Multi-target attacks exploit the fact that multiple ciphertexts share one lattice.
- An instance of the 2015 New Hope protocol, by definition, involves just one ciphertext. Generating a fresh \mathbf{a} means generating a fresh public key—a fresh lattice—to be used for a single ciphertext. Ergo, multiple ciphertexts do not share one lattice.

This paper’s attack fits the first part of the argument, but straightforwardly dodges the second part by attacking KEMs that are designed for IND-CCA2 security, KEMs that allow many ciphertexts for a single public key; the second part of the argument fails for those KEMs.

Note that the first part of the argument is merely a prediction, not a proof, so it could fail too. For example, multiple New Hope public keys share a ring, and thus share an S -unit lattice. There are some other lattice problems that appear to be vulnerable to S -unit precomputations (see, e.g., [34] and [86]); perhaps there are similar vulnerabilities in New Hope. However, this possibility is not what this paper’s attack exploits. This paper is not attacking KEMs that support only one ciphertext per public key, such as the 2015 version of New Hope.

5.2. New Hope 2017. A revised version [10] of New Hope was submitted to the NIST Post-Quantum Cryptography Standardization Project in 2017. In particular, this revision included two cryptosystems: NewHope-CPA for one ciphertext per key, and NewHope-CCA claiming IND-CCA2 security.

The text “all those pitfalls can be avoided by having the communicating parties generate a fresh \mathbf{a} at each instance of the protocol (as we propose)” was copied from [13], where it was correct, into [10], where it was incorrect: NewHope-CCA was not limited to a single ciphertext per key.

The erroneous description was eliminated in the 2019 version of the New Hope submission [11, page 18], which instead made the following claim: “all those pitfalls can be avoided by having the communicating parties generate a

¹⁹ See Section 3.1 regarding questions left open by the specific multi-target attack strategy outlined in [13].

fresh \mathbf{a} for each public key (as we propose)”. It is not clear that the consequences of this change were ever considered.

“All those pitfalls” in [11] appears to be referring to “all-for-the-price-of-one attacks”, as in [13]. Saying that the pitfalls “can be avoided” appears to be saying that there are no all-for-the-price-of-one attacks. But the argument that there are no all-for-the-price-of-one attacks is again not stated explicitly.

The argument stated in Section 5.1 relies on each public key having only one ciphertext. This argument is broken by the subsequent change from “each instance of the protocol” (i.e., each ciphertext) to “each public key”. Each public key has its own lattice, but that lattice is shared by all of the ciphertexts for that public key.

NewHope has similar heuristic security levels for key-recovery attacks and message-recovery attacks. This opened up key recovery as an obvious “all-for-the-price-of-one attack” against multiple ciphertexts, given that New Hope was modified to allow multiple ciphertexts per key. Similar comments apply to many other lattice proposals that claim IND-CCA2 security and that do not try to make the usual key-recovery attacks more difficult than the usual message-recovery attacks.

One could try to argue that, quantitatively, attacks against many ciphertexts for one public key are a smaller problem than attacks against many ciphertexts for many public keys. This is abandoning the claim that the pitfalls are *avoided*, and is instead arguing that the pitfalls are quantitatively *reduced*. This begs further questions, such as the following:

- How many ciphertexts are sent to the most popular public keys?
- What security level is achieved by that number of ciphertexts?
- If Kyber-512 is not simply discarded, should users deploying Kyber-512 be told that Kyber-512 keys are not safe to use for more than, say, 2^{10} ciphertexts per key?

The New Hope documentation already stated in 2015 [13, page 5] that it is “rather uncomfortable to have the security of all connections rely on a single instance of a lattice problem”. Reusing a lattice for all connections for one public key is quantitatively different from reusing a lattice for all connections for all public keys, but still creates the same basic issue of multi-target attacks being more efficient than single-target attacks.

5.3. NIST 2021. A 2021 NIST posting “Requirements for security against multi-target attacks, for McEliece and other code-based cryptosystems?” [88] said “it seems reasonable to be worried about 2^{64} target ciphertexts” and criticized code-based cryptosystems for allowing “decoding one out of many” attacks.

As background, the Classic McEliece submission [5, Section 8.3] had cited examples of multi-ciphertext attacks from 2002 [62] and 2011 [91] giving roughly a $T^{1/2}$ speedup for breaking one out of T ciphertexts, and had handled those attacks as follows:

Rather than analyzing multi-message security in detail, we rely on the general fact that attacking T targets cannot gain more than a factor T . For example, with our recommended 6688/6960/8192 parameter sets, one ciphertext is expected to be secure against an attacker without the resources to find an AES-256 key, and 2^{64} ciphertexts are expected to all be secure against an attacker without the resources to find an AES-192 key.

Taking a very high single-target security level straightforwardly protects against multi-target attacks, whether all-for-the-price-of-one attacks or one-out-of-many attacks or something intermediate.

Instead of acknowledging the submission for this protection against multi-target attacks, NIST criticized the submission for the fact that multi-target attacks existed in the first place. Meanwhile NIST did not criticize lattice submissions that (1) did not recommend such high security levels and (2) did not cite any multi-target cryptanalysis. In essence, one area of cryptography was punished for having investigated and proactively eliminated a threat, while another area of cryptography was rewarded for not even being aware of the threat, never mind protecting against it.

Readers who know the context of NIST comparing submissions to a competition, and who see NIST arguing that one-out-of-many attacks are of concern for code-based cryptography, can easily leap to the conclusion that there are no one-out-of-many attacks against lattice-based cryptography. It is easy to conflate lack of *knowledge* of such lattice attacks with lack of *existence* of those attacks; however, until a problem has survived extensive cryptanalysis, there is little reason to believe that the problem is hard to break.

5.4. ACM CCS 2021. A typical statement that cryptosystem X is “provably secure” communicates two ideas to the informed reader. The first idea is that there is a proof saying that there cannot be an attack against X *unless* there is an attack against some underlying problem P . The second idea is that there are good reasons to think that there are no attacks against P .

If a cryptanalysis paper then claims attacks against X , obviously the inconsistency requires explanation. Here are three possibilities to consider:

- Perhaps the attack analysis is incorrect. This is a cryptanalysis failure.
- Perhaps the proof is incorrect. This is a provable-security failure.
- Perhaps the attack is correct and the proof is correct—so the idea that there are no attacks against P has to be wrong; the proof is actually an attack tool, converting the attack against X into an attack against P . This is a different type of provable-security failure, where a problem P without sufficient cryptanalysis was being used to portray X as safe.

The case considered below is that X is one-out-of-many-ciphertext lattice security and P is many-sample Module-LWE.

Here is the context. Duman–Hövelmanns–Kiltz–Lyubashevsky–Seiler [47] gave a proof, not the proof at issue here, relating multi-target ROM IND-CCA2

security of a KEM to multi-target IND-CPA security of an underlying PKE. This appears to quantitatively improve upon previous proofs *for PKEs that do not suffer a $T \times$ loss from one-out-of- T attacks*.

The same paper gave a provable-security argument that for “lattice-based schemes” there is *no* loss from one-out-of- T attacks. Specifically, the paper said for these systems that “we can show that $\text{Adv}_{\text{PKE}}^{\text{IND-CPA}} \approx \text{Adv}_{\text{PKE}}^{(n,qc)\text{-IND-CPA}}$.” This is not an unconditional statement: it assumes “the hardness of MLWE as originally defined for the purpose of worst-case to average-case reductions [28, 30, 32] where the number of samples (using the same secret) is unlimited”.

The quantity $\text{Adv}_{\text{PKE}}^{(n,q)\text{-IND-CPA}}$ is defined in [47, Figure 2] as the attacker’s advantage at guessing a bit r in the following scenario. There are n legitimately generated public keys. The attacker carries out at most q queries (j, M_0, M_1) . The response to query (j, M_0, M_1) is a legitimately generated ciphertext for M_r under the j th public key. The quantity $\text{Adv}_{\text{PKE}}^{\text{IND-CPA}}$ is defined the same way for the single-target case $n = q = 1$.

The class of “lattice-based schemes” is not defined in [47], but presumably includes LPR, New Hope, Kyber, and further noisy-ElGamal schemes surveyed in [27, Section 8]. The features of these schemes that matter for the following analysis are as follows: key generation generates G uniformly at random, generates small a, e at random, computes $A = aG + e$, and outputs public key (G, A) ; encryption of message M generates small b, c, d at random, computes $B = Gb + d$, computes $C = M + Ab + c$, and outputs ciphertext (B, C) .

The conventional single-target “security proof” for such a scheme is as follows. Assume that $(G, aG + e)$ is indistinguishable from a uniform random pair (G, A') ; this is an example of a 1-sample small-secret Module-LWE indistinguishability assumption.²⁰ Also assume that $(G, A', Gb + d, A'b + c)$ is indistinguishable from a uniform random tuple (G, A', B', X') ; this is an example of a 2-sample small-secret Module-LWE indistinguishability assumption. Then (G, A, B, C) is indistinguishable from $(G, A', B, M + A'b + c)$, which is indistinguishable from $(G, A', B', M + X')$. Finally, one cannot distinguish $M_0 + X'$ from $M_1 + X'$ when X' is chosen uniformly at random.

Sometimes s -sample Module-LWE is defined only for large secrets a . One might think that an attack against, e.g., 3-sample small-secret Module-LWE can be trivially converted into an attack against 4-sample large-secret Module-LWE as follows:

- We are given 4 Module-LWE samples for a uniform random secret r (or a secret r from any distribution; the distribution will not matter). In other words, we are given uniform random G_0, G_1, G_2, G_3 , and given A_0, A_1, A_2, A_3 where $A_j = rG_j + e_j$ for small secrets e_0, e_1, e_2, e_3 .
- Compute $G'_j = -G_0^{-1}G_j$ and $A'_j = A_0G'_j + A_j$. The point of this computation is that $A'_j = e_0G'_j + e_j$.

²⁰ It can also be an example of a small-secret LWE indistinguishability assumption, and an example of a small-secret Ring-LWE indistinguishability assumption. The naming depends on details of the algebraic structures containing G, A , etc.

- We now have uniform random G'_1, G'_2, G'_3 , along with A'_1, A'_2, A'_3 where $A'_j = e_0 G'_j + e_j$. These are 3 Module-LWE samples for the small secret e_0 . Apply the 3-sample small-secret Module-LWE attack.

Caution is required here: what happens if G_0 is not invertible? One can try to dodge this difficulty by searching for an invertible G_j , but in general this can use more samples, possibly many more if invertibility is infrequent. For the case of LWE, [15, Lemma 2] takes as many samples as necessary to build an invertible matrix.

Now consider what happens if one replaces the ratios $G_0^{-1}G_1, G_0^{-1}G_2, G_0^{-1}G_3$ with the ratios $G_0^{-1}G_1, G_2^{-1}G_3$, and adjusts the A_j calculations analogously, assuming invertibility of G_0 and G_2 . Instead of obtaining 3 samples with a small secret e_0 , one obtains a batch of 2 independent problems with small secrets e_0 and e_2 , each with 1 sample. More generally, aside from the invertibility issue, $2s$ samples reduce to a batch of s independent 1-sample problems; $3s$ samples reduce to a batch of s independent 2-sample problems; etc.

This generalized reduction appears to be the point of the brief proof outline in [47, page 3, bottom of second column]. However, the conclusions drawn in [47] go beyond this in three ways that are not justified in [47]:

- The many-sample Module-LWE problem is claimed to be difficult: “we believe that in practice the MLWE problem with k samples is no easier than with 1 sample”. However, the Arora–Ge attack cited in [47, footnote 2] is already known to break the Module-LWE problem for (e.g.) the Kyber error distribution with a feasible number of samples, showing that a proof based on that problem would be vacuous for attacks against a feasible number of Kyber targets.
- Single-target and multi-target IND-CPA advantages are claimed to be proven to be close. However, this does not logically follow from a one-way conversion of IND-CPA attacks into Module-LWE attacks. The obvious way to close this gap would be to prove a conversion the other way (presumably also breaking the many-ciphertext case via the Arora–Ge attack)—but the words “plausibly much harder problem” in [47, footnote 2] indicate that a conversion is not known the other way.
- The type of multi-target attack for which [47] claims provable security is not merely an attack against many independent keys, but an attack against many ciphertexts per key. However, if one wants to convert an attack against s ciphertexts for one key (with many ciphertexts multiplied by the same public G , making G a natural target for precomputation) into an attack against $\Theta(s)$ Module-LWE samples, it is not obviously sufficient to convert an attack against s independent keys into an attack against $\Theta(s)$ Module-LWE samples.

Perhaps the second and third items have been addressed by a full proof, but the proof outline in [47] does not seem to have been written as a summary of a full proof. For example, the stated number of samples, “ $\max(n, q_C)$ ”, does not account for losing at least n samples as denominators, and for the possibility of

needing more samples for the necessary invertibility. What is cited here in [47], namely the transformation from [15], does not count the number of samples, and does not handle Ring-LWE or Module-LWE.

The set of T for which this paper gives heuristic one-out-of- T -ciphertext speedups is not covered by the set of T for which there is a $\Theta(T)$ -sample Arora–Ge attack, so the fact that the proof outlined in [47] is vacuous for the number of samples covered by the Arora–Ge attack is not logically sufficient to reconcile [47] with this paper. However, if there *is* a conversion from a multi-ciphertext attack into a multi-sample Module-LWE attack, the simplest explanation is not “these attacks do not exist, except that the Arora–Ge attack suddenly breaks Module-LWE with enough samples”, but rather “Module-LWE security degrades with the number of samples”.

More fundamentally, it is not clear that such a conversion exists. There is a proof gap at this point in [47]. It is clear how the proof techniques mentioned in [47] apply to multiple keys, but it is not clear how those techniques justify claiming multi-ciphertext security for a single key.

6 The curious case of FrodoKEM

We reiterate the crucial point: if the reduction proving security is “loose,” like the one above, the efficiency of the scheme is impacted, because we must move to a larger security parameter. —1996 Bellare–Rogaway [23]

In terms of security, Frodo’s conservative design choices are laudable. —2022 NIST [4, Section 4.3.1]

This section focuses specifically on FrodoKEM. Unusual features of the FrodoKEM design might at first seem to make FrodoKEM quantitatively more resistant to one-out-of-many-ciphertext attacks than many lattice systems are (although perhaps not quite as resistant as fixed-weight ternary systems; see Section 4.6). However, it turns out that the same design approach opens up FrodoKEM to a trivial, well-known attack strategy.

The current version of FrodoKEM includes three proposed parameter sets:²¹ FrodoKEM-640, FrodoKEM-976, and FrodoKEM-1344, targeting the security levels of brute-force single-key search for AES-128, AES-192, and AES-256 respectively. For each parameter set, the attack highlighted in this section is *almost* below the target security level. This disproves the claim in [12, page 44] that “the FrodoKEM parameter sets comfortably match their target security levels with a large margin”.

When T ciphertexts are sent to one FrodoKEM public key, the same well-known attack strategy immediately produces an all-for-the-price-of-one attack, and a one-out-of-many attack that is approximately $T \times$ faster. These

²¹ Actually six, since each parameter set has an AES-based version and a SHAKE-based version. The difference has a marginal effect on the performance of the attack here.

attacks against FrodoKEM are valid even if the standard lattice heuristics are incorrect.

6.1. Multiple independent ciphertexts. Say E is a PKE with message space X . Write E^j for the PKE with message space X^j that encrypts (M_1, \dots, M_j) by using E to independently encrypt each of M_1, \dots, M_j .

An attack that completely decrypts jT ciphertexts for E is also an attack that completely decrypts T ciphertexts for E^j . But decrypting *some* of the jT ciphertexts for E is, as j increases, less and less effective as an attack against T ciphertexts for E^j .

Quantitatively, if each of the jT ciphertexts for E is decrypted with probability p , then each ciphertext for E^j is decrypted with probability just p^j . For essentially the same reason, as j increases, ROM IND-CPA for a tweak of E^j becomes provably almost as hard to break as OW-CPA for E ; see [58, Section 3.4].

It is hard to imagine how this could be the best way to defend against one-out-of-many attacks (or against IND-CPA attacks). The problem is that E^2 is twice as expensive as E . Normally, for the same cost as E^2 , one can instead take E with larger parameters, aiming at a much higher security target, presumably making *all* attacks more difficult. On the other hand, perhaps one is forced to use E^2 for other reasons, as in Section 6.2.

6.2. PKEs with small message spaces. Say one is trying to build an IND-CCA2 KEM out of a limited PKE E that sends only 32-bit messages.

The limited message space does not stop E from achieving IND-CPA security: the encryption procedure can generate many more bits of randomness. However, if one builds a KEM by applying any common variant of the Fujisaki–Okamoto transform to E , then the FO derandomization step eliminates the additional randomness. Ciphertexts and session keys for any particular public key have just 32 bits of entropy and are easily breakable.

A straightforward fix, at the cost of making everything $8\times$ less efficient, is to instead apply the FO transform to E^8 . Ciphertexts and session keys then have 256 bits of entropy.

This is essentially how FrodoKEM and its underlying PKE, FrodoPKE, are built. See, e.g., [12, page 17] saying “Several (\bar{m}) ciphertexts are generated at once”. The “ \bar{m} ” parameter *could* be chosen as 1, but is instead chosen as 8 for performance reasons derived from a design goal of minimizing “the sum of the bit lengths of FrodoPKE’s ciphertext and its public key”; see [12, page 23].

Concretely, FrodoKEM-1344 can be viewed as sending 8 independent ciphertexts. Each ciphertext, in turn, communicates just 4 message bits at each position in a C vector of length 8 (while B is much larger), for a total of 32 message bits in each ciphertext, or 256 bits across the 8 ciphertexts.

6.3. Multi-ciphertext attacks on FrodoKEM. One might think that this 8-ciphertext structure reduces the quantitative impact of attacks trying to decrypt one out of many FrodoKEM ciphertexts for the same public key.

Concretely, to exploit the shared lattice to break one of 2^{40} ciphertexts, one would want to break $1/2^5$ of the underlying length-8 ciphertexts. This factor 2^5 is much less concerning than the original 2^{40} , especially since FrodoKEM is a “conservative” lattice design claiming that its lattice problems have even more security than needed. However, the same performance pressure that drove FrodoKEM to use 8 ciphertexts turns out to create an even bigger security problem for FrodoKEM.

Consider FrodoKEM-640, for which [12, Table 11] estimates “LWE hardness” of $2^{175.1}$ “gates” and $2^{110.4}$ “memory in bits”. This appears to be the basis of the claim in [12, page 44] that “the FrodoKEM parameter sets comfortably match their target security levels with a large margin”—but FrodoKEM has an even larger attack surface than the underlying LWE problem does. It is an error to conflate LWE hardness with FrodoKEM security.

Like FrodoKEM-1344, FrodoKEM-640 sends 8 ciphertexts, each having 8 positions, but—again, because of performance pressure—FrodoKEM-640 has just 2 message bits per position, a total of 128 message bits.

This limited message space, combined with FO derandomization, means that there are only 2^{128} possible FrodoKEM-640 ciphertexts for any particular public key. Consequently, FrodoKEM-640 is vulnerable to a trivial brute-force search through messages.

A full search through all 2^{128} messages is much more efficient than the $2^{175.1}$ “gates” and $2^{110.4}$ “memory in bits” estimated in [12] for FrodoKEM’s “LWE hardness”.²² An attacker faced with 2^{40} ciphertexts for a FrodoKEM-640 public key will successfully decrypt one of the ciphertexts after trying about 2^{88} messages, a feasible computation.

There is nothing new about this attack: it is an example of one of the workhorses of symmetric cryptanalysis, namely “guess possibilities for secret x , checking each guess against published $F(x)$ values”. But the literature does not seem to have pointed out that FrodoKEM-640 opens up this attack with a 128-bit x ; that this contradicts the claim that “the FrodoKEM parameter sets comfortably match their target security levels with a large margin”; and that F depends only on the public key, allowing a $2^{128}/T$ attack against T FrodoKEM-640 ciphertexts for that key.

6.4. Broader impacts. It is interesting to compare this FrodoKEM-640 attack to Section 5.3. Recall that, in 2021, NIST criticized code-based proposals on the basis of known one-out-of- T -ciphertext attacks that save roughly \sqrt{T} from

²² FrodoKEM encryption is slow, but not *that* slow. Also, computing well under 1% of each ciphertext, using well under 1% of the multiplications in \mathbb{Z}/q , is sufficient for the attack. Readers interested in more detailed optimization should note that the IND-CCA2 definition provides a session key to the attacker, so one can skip all the multiplications and simply check the session-key hash. Furthermore, most of the hash computation for the session key can be precomputed, since FrodoKEM’s hash input puts the secret message *after* the public ciphertext. For comparison, one rule of thumb for the order of hash inputs—see, e.g., [30]—is that whatever is least likely to be predictable by the attacker should come first.

the starting cost 2^b of decoding—even when the code-based proposals cited the attacks, presumed that $2^b/T$ is possible, and recommended large b as a defense. The presumed $2^b/T$ is far beyond the $2^{128}/T$ cost of the one-out-of- T -ciphertext attack against FrodoKEM-640, because b has intentionally been chosen much larger than 128.

It is also interesting to compare this FrodoKEM-640 attack to NIST’s stated reason [87] for continuing to recommend AES-128 in the context of post-quantum KEMs:

Finally, you seem to be advocating that NIST respond to the possibility of multikey attacks by withdrawing AES128, rather than by advocating for modes of operation that have built-in multi-key security. Given that

- 1) AES-128 is the most widely used block cipher at present, and it has never come even close to being practically attacked based on an insufficiently large key size.
- 2) Most widely-used high-volume protocols, where multi-key security is actually a concern (e.g. TLS and IPsec) already have built-in protections against multi-key attacks.

It seems premature to pull AES128.

The protections mentioned here are randomizing the inputs to various protocols using AES so that any specific input is not encrypted under many AES keys. This fails to protect a broader system that uses FrodoKEM-640 to communicate AES-128 keys in the first place: the attack directly breaks FrodoKEM-640 ciphertexts, independently of how the resulting session keys are used.

Trying to review multi-target security of every aspect of a cryptographic system, and trying to patch every multi-target attack that is found, is much more complicated and much less robust than systematically requiring higher security levels. This particular attack is very easy, but it still appears to have been missed in the FrodoKEM design. Note that multi-target security was within scope for the FrodoKEM proposal: see, e.g., [12, page 33], claiming that a tweak elsewhere in FrodoKEM “has the potential to provide stronger multi-target security”.

Moving up to FrodoKEM-976 makes these multi-target attacks infeasible today. But this is not the end of the story: FrodoKEM-976 claims a correspondingly higher security level, and claims to have “a large margin” beyond that security level. The claim of a large margin is false: FrodoKEM-976’s performance-driven choice to use just 192-bit messages exposes FrodoKEM-976 to a straightforward 2^{192} -guess single-target brute-force attack, and to correspondingly more efficient multi-ciphertext attacks.

Most other noisy-DH lattice proposals (e.g., Kyber) have less performance pressure on the size of the message space and simply take 256-bit messages for all security levels. When the target security level is AES-256, the same easy attack of searching through messages means that none of these proposals have large security margins; and this attack becomes $T\times$ faster for decrypting one out of T ciphertexts to the same public key. Of course, $2^{256}/T$ is of much less real-world concern than $2^{128}/T$; pushing $2^{128}/T$ up to 2^{256} is an improvement if it can be achieved, but setting $2^{256}/T$ as a goal is simpler and safer.

6.5. The importance of accounting for looseness. FrodoKEM was not the only submission to the NIST Post-Quantum Cryptography Standardization Project putting heavy emphasis on proofs. See, e.g., [38, Section 14]: “MQDSS is the first multivariate signature scheme that is provably secure, and whose security relies solely on the MQ problem.”

The MQDSS proofs said that a ROM attack against MQDSS could be converted into an attack against the MQ problem, a simpler problem that had already attracted cryptanalysis. However, these proofs were *loose*. The proofs did not say that the cost of the resulting MQ attack was *close to* the cost of the MQDSS attack; the proofs merely put some sort of quantitative limit on the change in attack costs. The proofs thus allowed MQDSS to be much easier to break than the underlying MQ problem.

As Bellare and Rogaway noted in [23], applying a loose proof means that, logically, one “must move to a larger security parameter” to compensate for the looseness. But MQDSS chose cryptosystem parameters without regard to the looseness. MQDSS then had its security claims disproven by an attack [64] that exploited the looseness of the proofs. See also [68] for a survey of many more provable-security failures, including failures directly attributable to looseness.

Similarly, various loose theorems are falsely advertised in [12, Section 5] as “supporting the security of FrodoKEM”. To the extent that the security of the underlying problems has been studied, these theorems would support the security of a hypothetical CarefulFrodoKEM with parameters chosen to systematically account for looseness; but it is an error to conflate the security of FrodoKEM with the security of CarefulFrodoKEM. This error has been remarkably persistent:

- I had, in [26], already pointed out this error in an earlier version of [12], and yet the same error remained in [12].
- In [29], I recommended that NIST add “procedural protections against loose proofs”. In [14], a NIST team member replied “Stop propagandizing”.
- Secret NIST discussions included, e.g., slides [3] from another NIST team member repeating the false advertising of FrodoKEM’s proofs, categorically dismissing objections as “a lot of DJBFUD about the relevance of asymptotic security reductions to security”,²³ and not rationally addressing the objections. What makes this particularly bizarre is that the same slides also admit, regarding the advertising, that “not all of this applies for practically relevant parameters”.
- A public NIST report [4] also admitted that “these theorems do not hold for the concrete parameter choices used in Frodo” but tried to defend the false advertising as follows: (1) labeling this as “typical”; (2) claiming that the theorems “do indicate some fundamental soundness in the core idea underlying the Frodo approach”—which is not the question at hand.

NIST was supposed to be asking whether FrodoKEM is secure. Merely asking whether there is “some fundamental soundness in the core idea” is a poor

²³ One might wonder why there is no credit to, e.g., [23] for the BRFUD and [68] for the KMFUD.

substitute for this; the history of cryptography is littered with the corpses of cryptosystems that have “some fundamental soundness in the core idea”. NIST was also supposed to be evaluating “the quality of the analysis provided by the submitter”; having loose theorems incorrectly labeled as “supporting the security of FrodoKEM” should have been treated negatively, not positively.

To see the relevance of this error to the sudden collapse of FrodoKEM’s “large margin”, consider the fact that CarefulFrodoKEM would be forced to account for, among other things,

- $2q + 1$ divided by the size of the message space in [12, Lemma 5.3], where q is the number of hash calls in an attack; and
- the looseness factor T in the generic proof that a T -target attack implies a 1-target attack.

The trivial FrodoKEM attack highlighted in this section exploits exactly these two looseness issues. For a single ciphertext, the attack searches the message space (which is dangerously small for FrodoKEM-640, and does not have a “large margin” for any of the parameter sets), checking message guesses via hash calls; for T ciphertexts, the attack gains a factor T . The success of this attack in disproving one of FrodoKEM’s security claims is thus directly attributable to FrodoKEM’s failure to account for proof looseness, a failure endorsed by NIST. CarefulFrodoKEM would be bigger and slower than FrodoKEM, but it would not have this type of security failure.

According to [4, page 38], the only reason that NIST did not select FrodoKEM was performance. Other organizations are continuing to consider FrodoKEM on the basis of its “conservative” design. It is deeply concerning to see that the evaluations of FrodoKEM are based primarily upon the volume of allegedly related proofs, with no insistence upon investigating what—if anything—those proofs actually say about FrodoKEM. Much more work needs to be done to systematically study the FrodoKEM attack surface.

References

- [1] — (no editor), *IEEE international symposium on information theory, ISIT 2007, Nice, France, June 24–29, 2007*, IEEE, 2007. URL: <https://ieeexplore.ieee.org/xpl/conhome/4497218/proceeding>. See [94].
- [2] David Adrian, Karthikeyan Bhargavan, Zakir Durumeric, Pierrick Gaudry, Matthew Green, J. Alex Halderman, Nadia Heninger, Drew Springall, Emmanuel Thomé, Luke Valenta, Benjamin VanderSloot, Eric Wustrow, Santiago Zanella Béguelin, Paul Zimmermann, *Imperfect forward secrecy: how Diffie-Hellman fails in practice*, *Communications of the ACM* **62** (2019), 106–114. URL: <https://weakdh.org>. Citations in this document: §1.
- [3] Gorjan Alagic, *FrodoKEM in the third round* (2021). URL: <https://nist.pqcrypto.org/foia/20221107/Frodo-final.pptx>. Citations in this document: §6.5.

- [4] Gorjan Alagic, Daniel Apon, David Cooper, Quynh Dang, Thinh Dang, John Kelsey, Jacob Lichtinger, Carl Miller, Dustin Moody, Rene Peralta, Ray Perlner, Angela Robinson, Daniel Smith-Tone, Yi-Kai Liu, *Status report on the third round of the NIST Post-Quantum Cryptography Standardization Process* (2022). NISTIR 8413. URL: <https://csrc.nist.gov/publications/detail/nistir/8413/final>. Citations in this document: §3, §4, §4, §4, §6, §6.5, §6.5.
- [5] Martin R. Albrecht, Daniel J. Bernstein, Tung Chou, Carlos Cid, Jan Gilcher, Tanja Lange, Varun Maram, Ingo von Maurich, Rafael Misoczki, Ruben Niederhagen, Kenneth G. Paterson, Edoardo Persichetti, Christiane Peters, Peter Schwabe, Nicolas Sendrier, Jakub Szefer, Cen Jung Tjhai, Martin Tomlinson, Wen Wang, *Classic McEliece: conservative code-based cryptography* (2020). URL: <https://classic.mceliece.org/nist/mceliece-20201010.pdf>. Citations in this document: §1, §5.3.
- [6] Martin R. Albrecht, Benjamin R. Curtis, Amit Deo, Alex Davidson, Rachel Player, Eamonn W. Postlethwaite, Fernando Virdia, Thomas Wunderer, *Estimate all the {LWE, NTRU} Schemes!*, in SCN 2018 [36] (2018), 351–367. URL: <https://eprint.iacr.org/2018/331>. Citations in this document: §2, §2.9, §4.5, §4.5, §4.5.1, §4.5.1, §4.5, §4.5.2, §4.5.2, §4.5.
- [7] Martin R. Albrecht, Léo Ducas, *Lattice attacks on NTRU and LWE: a history of refinements* (2021). URL: <https://eprint.iacr.org/2021/799>. Citations in this document: §2.
- [8] Martin R. Albrecht, Vlad Gheorghiu, Eamonn W. Postlethwaite, John M. Schanck, *Estimating quantum speedups for lattice sieves*, in Asiacrypt 2020 [81] (2020), 583–613. URL: <https://eprint.iacr.org/2019/1161>. Citations in this document: §5.
- [9] Martin R. Albrecht, Florian Göpfert, Fernando Virdia, Thomas Wunderer, *Revisiting the expected cost of solving uSVP and applications to LWE*, in Asiacrypt 2017 [96] (2017). URL: <https://eprint.iacr.org/2017/815>. Citations in this document: §2.5, §14.
- [10] Erdem Alkim, Roberto Avanzi, Joppe Bos, Leo Ducas, Antonio de la Piedra, Thomas Poppelmann, Peter Schwabe, Douglas Stebila, *NewHope: algorithm specifications and supporting documentation* (2017). URL: https://web.archive.org/web/20190411045044/https://newhopecrypto.org/data/NewHope_2017_12_21.pdf. Citations in this document: §5.2, §5.2.
- [11] Erdem Alkim, Roberto Avanzi, Joppe Bos, Leo Ducas, Antonio de la Piedra, Thomas Poppelmann, Peter Schwabe, Douglas Stebila, Martin R. Albrecht, Emmanuela Orsini, Valery Osheter, Kenneth G. Paterson, Guy Peer, Nigel P. Smart, *NewHope: algorithm specifications and supporting documentation* (2019). URL: https://web.archive.org/web/20191015124615/https://newhopecrypto.org/data/NewHope_2019_04_10.pdf. Citations in this document: §4.5, §5.2, §5.2.
- [12] Erdem Alkim, Joppe W. Bos, Léo Ducas, Patrick Longa, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Chris Peikert, Ananth Raghunathan, Douglas Stebila, *FrodoKEM: Learning With Errors key encapsulation: algorithm specifications and supporting documentation* (2021). URL: <https://web.archive.org/web/20220119174856/https://frodokem.org/files/FrodoKEM-specification-20210604.pdf>. Citations in this document: §1.2, §2, §6, §6.2, §6.2, §6.3, §6.3, §6.3, §6.4, §6.5, §6.5, §6.5, §6.5.
- [13] Erdem Alkim, Léo Ducas, Thomas Poppelmann, Peter Schwabe, *Post-quantum key exchange—a new hope*, in USENIX Security 2016 [59] (2016), 327–343. URL:

- <https://eprint.iacr.org/2015/1092>. Citations in this document: §2, §2.1, §3.1, §4.5, §4.5, §5.1, §19, §5.2, §5.2, §5.2.
- [14] Daniel Apon, *Re: Looseness, security risks, and LWR vs. LWE* (2021). URL: <https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/YxOwZuZP6ag/m/CqxTSn3TCAAJ>. Citations in this document: §6.5.
- [15] Benny Applebaum, David Cash, Chris Peikert, Amit Sahai, *Fast cryptographic primitives and circular-secure encryption based on hard learning problems*, in *Crypto 2009* [54] (2009), 595–618. URL: <https://www.wisdom.weizmann.ac.il/~bennyap/pubs/kdm-learning.pdf>. Citations in this document: §5.4, §5.4.
- [16] Jean-Philippe Aumasson, *Too much crypto* (2019). URL: <https://eprint.iacr.org/2019/1492>. Citations in this document: §1.3.
- [17] Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, Gregor Seiler, Damien Stehlé, *CRYSTALS-Kyber: Algorithm specifications and supporting documentation* (2020). URL: <https://web.archive.org/web/20211007045636/https://pq-crystals.org/kyber/data/kyber-specification-round3.pdf>. Citations in this document: §1.1, §1.1, §1.1, §1.1, §2.9.
- [18] Roberto Avanzi, Howard M. Heys (editors), *Selected areas in cryptography—SAC 2016—23rd international conference, St. John’s, NL, Canada, August 10–12, 2016, revised selected papers*, 10532, Springer, 2017. ISBN 978-3-319-69452-8. See [71].
- [19] László Babai, *On Lovász’ lattice reduction and the nearest lattice point problem*, *Combinatorica* **6** (1986), 1–13. Citations in this document: §3.1.
- [20] Razvan Barbulescu, *Algorithms of discrete logarithm in finite fields* (2013). URL: <https://tel.archives-ouvertes.fr/tel-00925228/file/these%5Favec%5Fresume.pdf>. Citations in this document: §1.
- [21] Anja Becker, Léo Ducas, Nicolas Gama, Thijs Laarhoven, *New directions in nearest neighbor searching with applications to lattice sieving*, in *SODA 2016* [70] (2016), 10–24. URL: <https://eprint.iacr.org/2015/1128>. Citations in this document: §2.2, §2.2, §2.2, §3.2.
- [22] Mihir Bellare, Hannah Davis, Felix Günther, *Separate your domains: NIST PQC KEMs, oracle cloning and read-only indifferenciability*, in *Eurocrypt 2020* [35] (2020), 3–32. URL: <https://eprint.iacr.org/2020/241>. Citations in this document: §2.9.
- [23] Mihir Bellare, Phillip Rogaway, *The exact security of digital signatures: how to sign with RSA and Rabin*, in *Eurocrypt 1996* [78] (1996), 399–416. URL: <https://cseweb.ucsd.edu/~mihir/papers/exactsigs.html>. Citations in this document: §6, §6.5, §23.
- [24] Daniel J. Bernstein, *Understanding brute force* (2005). ECRYPT STVL Workshop on Symmetric Key Encryption. URL: <http://cr.yp.to/papers.html#bruteforce>. Citations in this document: §1.
- [25] Daniel J. Bernstein, *Break a dozen secret keys, get a million more for free* (2015). URL: <https://blog.cr.yp.to/20151120-batchattacks.html>. Citations in this document: §1.
- [26] Daniel J. Bernstein, *OFFICIAL COMMENT: Frodo* (2018). URL: <https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/rJYnyTEi92E/m/15xBpeTpBQAJ>. Citations in this document: §6.5.
- [27] Daniel J. Bernstein, *Comparing proofs of security for lattice-based encryption* (2019). Second PQC Standardization Conference. URL: <https://cr.yp.to/papers.html#latticeproofs>. Citations in this document: §5.4.

- [28] Daniel J. Bernstein, *Cryptographic competitions* (2020). URL: <https://cr.yp.to/papers.html#competitions>. Citations in this document: §1.3.
- [29] Daniel J. Bernstein, *Re: Looseness, security risks, and LWR vs. LWE* (2021). URL: <https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/Yx0wZuZP6ag/m/GXIwomNsCAAJ>. Citations in this document: §6.5.
- [30] Daniel J. Bernstein, *Re: Was: eddsa (un)suited for mandatory to implement ciphersuite?* (2022). URL: <https://mailarchive.ietf.org/arch/msg/cfrg/GRigAYvZ8-Z8qmxJ1j0iKR8eLyQ/>. Citations in this document: §22.
- [31] Daniel J. Bernstein, Chitchanok Chuengsatiansup, Tanja Lange, Christine van Vredendaal, *NTRU Prime: round 2* (2019). URL: <https://ntruprime.cr.yp.to/nist.html>. Citations in this document: §2.3, §4.4, §4.6.
- [32] Daniel J. Bernstein, Tanja Lange, *Non-randomness of S -unit lattices* (2021). URL: <https://cr.yp.to/papers.html#spherical>. Citations in this document: §1.3, §11.
- [33] Joe P. Buhler (editor), *Algorithmic number theory, third international symposium, ANTS-III, Portland, Oregon, USA, June 21–25, 1998, proceedings*, Lecture Notes in Computer Science, 1423, Springer, 1998. ISBN 3-540-64657-4. See [56].
- [34] Peter Campbell, Michael Groves, Dan Shepherd, *Soliloquy: a cautionary tale* (2014). URL: https://web.archive.org/web/20211022132624/https://docbox.etsi.org/Workshop/2014/201410_CRYPT0/S07_Systems_and_Attacks/S07_Groves_Annex.pdf. Citations in this document: §5.1.
- [35] Anne Canteaut, Yuval Ishai (editors), *Advances in cryptology—EUROCRYPT 2020—39th annual international conference on the theory and applications of cryptographic techniques, Zagreb, Croatia, May 10–14, 2020, proceedings, part II*, Lecture Notes in Computer Science, 12106, Springer, 2020. ISBN 978-3-030-45723-5. See [22].
- [36] Dario Catalano, Roberto De Prisco (editors), *Security and cryptography for networks—11th international conference, SCN 2018, Amalfi, Italy, September 5–7, 2018, proceedings*, Lecture Notes in Computer Science, 11035, Springer, 2018. See [6].
- [37] André Chailloux, Johanna Loyer, *Lattice sieving via quantum random walks* (2021). URL: <https://eprint.iacr.org/2021/570>. Citations in this document: §5.
- [38] Ming-Shing Chen, Andreas Hülsing, Joost Rijneveld, Simona Samardjiska, Peter Schwabe, *MQDSS specifications* (2019). URL: https://web.archive.org/web/20220120235746/https://mqdss.org/files/MQDSS_Ver2.pdf. Citations in this document: §6.5.
- [39] Cloudflare Research, *Post-quantum key agreement* (2022). URL: <https://pq.cloudflare.com>. Citations in this document: §1.3.
- [40] Henri Cohen, *Advanced topics in computational number theory* **193** (2000). ISBN 0-387-98727-4. Citations in this document: §11.
- [41] Don Coppersmith, Adi Shamir, *Lattice attacks on NTRU*, in Eurocrypt 1997 [50] (1997), 52–61. URL: https://link.springer.com/content/pdf/10.1007/3-540-69053-0_5.pdf. Citations in this document: §2.4.
- [42] Jintai Ding, Rainer Steinwandt (editors), *Post-quantum cryptography—10th international conference, PQCrypto 2019, Chongqing, China, May 8–10, 2019, revised selected papers*, 11505, Springer, 2019. ISBN 978-3-030-25509-1. See [43].
- [43] Emmanouil Doulgerakis, Thijs Laarhoven, Benne de Weger, *Finding closest lattice vectors using approximate Voronoi cells*, in PQCrypto 2019 [42] (2019), 3–22. Citations in this document: §3.2.

- [44] Emmanouil Doulgierakis, Thijs Laarhoven, Benne de Weger, *Sieve, enumerate, slice, and lift: hybrid lattice algorithms for SVP via CVPP*, in Africacrypt 2020 [84] (2020), 301–320. URL: <https://eprint.iacr.org/2020/487>. Citations in this document: §2.2, §4, §4, §4.6.
- [45] Léoucas, *Shortest vector from lattice sieving: A few dimensions for free*, in Eurocrypt 2018 [82] (2018), 125–145. URL: <https://eprint.iacr.org/2017/999>. Citations in this document: §2.2, §2.2, §4, §4, §4.6.
- [46] Léoucas, Thijs Laarhoven, Wessel P. J. van Woerden, *The randomized slicer for CVPP: sharper, faster, smaller, batchier*, in PKC 2020 [66] (2020), 3–36. URL: <https://eprint.iacr.org/2020/120>. Citations in this document: §3, §3.2, §12, §12, §3.2, §3.3, §3.3, §3.3, §3.3, §4.4, §4.4.
- [47] Julien Duman, Kathrin Hövelmanns, Eike Kiltz, Vadim Lyubashevsky, Gregor Seiler, *Faster lattice-based KEMs via a generic Fujisaki-Okamoto transform using prefix hashing*, in CCS 2021 [67] (2021), 2722–2737. URL: <https://eprint.iacr.org/2021/1351>. Citations in this document: §1.1, §1.1, §7, §5, §5.4, §5.4, §5.4, §5.4, §5.4, §5.4, §5.4, §5.4, §5.4, §5.4, §5.4, §5.4, §5.4, §5.4, §5.4.
- [48] Thomas Espitau, Paul Kirchner, *The nearest-colattice algorithm: time-approximation tradeoff for approx-CVP*, in ANTS 2020 [51] (2020), 251–266. URL: <https://eprint.iacr.org/2020/694>. Citations in this document: §3, §3.1, §3.1, §3.1, §3.1, §3.1, §3.1, §3.1, §3.1, §10, §3.1, §3.1, §12, §3.3, §3.3, §3.3, §3.3.
- [49] Philippe Flajolet, Robert Sedgewick, *Analytic combinatorics*, Cambridge University Press, 2009. ISBN 978-0-521-89806-5. URL: <https://ac.cs.princeton.edu/home/>. Citations in this document: §4.2.
- [50] Walter Fumy (editor), *Advances in cryptology—EUROCRYPT '97, international conference on the theory and application of cryptographic techniques, Konstanz, Germany, May 11–15, 1997*, Lecture Notes in Computer Science, 1233, Springer, 1997. See [41].
- [51] Steven Galbraith (editor), *ANTS XIV: proceedings of the fourteenth algorithmic number theory symposium, Auckland 2020*, Open Book Series, 4, Mathematical Sciences Publishers, 2020. ISBN 978-1-935107-07-1. See [48].
- [52] Michael T. Goodrich, *Zig-zag sort: a simple deterministic data-oblivious sorting algorithm running in $O(n \log n)$ time*, in STOC 2014 [92] (2014), 684–693. URL: <https://arxiv.org/abs/1403.2777>. Citations in this document: §1.3.
- [53] Qian Guo, Thomas Johansson, *Faster dual lattice attacks for solving LWE with applications to CRYSTALS*, in Asiacrypt 2021 [97] (2021), 33–62. Citations in this document: §1.1.
- [54] Shai Halevi (editor), *Advances in cryptology—CRYPTO 2009, 29th annual international cryptology conference, Santa Barbara, CA, USA, August 16–20, 2009. proceedings*, 5677, Springer, 2009. ISBN 978-3-642-03355-1. See [15].
- [55] Max Heiser, *Improved quantum hypercone locality sensitive filtering in lattice sieving* (2021). URL: <https://eprint.iacr.org/2021/1295>. Citations in this document: §5.
- [56] Jeffrey Hoffstein, Jill Pipher, Joseph H. Silverman, *NTRU: a ring-based public key cryptosystem*, in ANTS 1998 [33] (1998), 267–288. URL: <https://ntru.org/f/hps98.pdf>. Citations in this document: §2.4, §2.4, §2.4, §4.6.
- [57] Jeffrey Hoffstein, Jill Pipher, Joseph H. Silverman, *NTRU: a new high speed public key cryptosystem* (2016). Circulated at Crypto 1996, put online in 2016. URL: <https://ntru.org/f/hps96.pdf>. Citations in this document: §2.4.

- [58] Dennis Hofheinz, Kathrin Hövelmanns, Eike Kiltz, *A modular analysis of the Fujisaki-Okamoto transformation*, in TCC 2017-1 [63] (2017), 341–371. URL: <https://eprint.iacr.org/2017/604>. Citations in this document: §6.1.
- [59] Thorsten Holz, Stefan Savage (editors), *25th USENIX security symposium, USENIX Security 16, Austin, TX, USA, August 10–12, 2016*, USENIX Association, 2016. URL: <https://www.usenix.org/conference/usenixsecurity16>. See [13].
- [60] Nick Howgrave-Graham, *A hybrid lattice-reduction and meet-in-the-middle attack against NTRU*, in Crypto 2007 [80] (2007), 150–169. URL: <https://www.iacr.org/archive/crypto2007/46220150/46220150.pdf>. Citations in this document: §3.1, §3.1, §8, §3.1.
- [61] Yuval Ishai, Vincent Rijmen (editors), *Advances in cryptology—EUROCRYPT 2019—38th annual international conference on the theory and applications of cryptographic techniques, Darmstadt, Germany, May 19–23, 2019, proceedings, part II*, Springer, 2019. ISBN 978-3-030-17655-6. See [86].
- [62] Thomas Johansson, Fredrik Jönsson, *On the complexity of some cryptographic problems based on the general decoding problem*, IEEE Transactions on Information Theory **48** (2002), 2669–2678. Citations in this document: §5.3.
- [63] Yael Kalai, Leonid Reyzin (editors), *Theory of cryptography—15th international conference, TCC 2017, Baltimore, MD, USA, November 12–15, 2017, proceedings, part I*, Lecture Notes in Computer Science, 10677, Springer, 2017. ISBN 978-3-319-70499-9. See [58].
- [64] Daniel Kales, Greg Zaverucha, *Forgery attacks on MQDSSv2.0* (2019). URL: <https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/round-2/official-comments/MQDSS-round2-official-comment.pdf>. Citations in this document: §6.5.
- [65] Aggelos Kiayias (editor), *Topics in cryptology—CT-RSA 2011—the cryptographers’ track at the RSA Conference 2011, San Francisco, CA, USA, February 14–18, 2011, proceedings*, Lecture Notes in Computer Science, 6558, Springer, 2011. ISBN 978-3-642-19073-5. See [73].
- [66] Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, Vassilis Zikas (editors), *Public-key cryptography—PKC 2020—23rd IACR international conference on practice and theory of public-key cryptography, Edinburgh, UK, May 4–7, 2020, proceedings, part II*, 12111, Springer, 2020. ISBN 978-3-030-45387-9. See [46].
- [67] Yongdae Kim, Jong Kim, Giovanni Vigna, Elaine Shi (editors), *CCS ’21: 2021 ACM SIGSAC conference on computer and communications security, virtual event, Republic of Korea, November 15–19, 2021*, ACM, 2021. ISBN 978-1-4503-8454-4. See [47].
- [68] Neal Koblitz, Alfred Menezes, *Critical perspectives on provable security: fifteen years of “another look” papers*, Advances in Mathematics of Communications **13** (2019), 517–558. URL: <https://eprint.iacr.org/2019/1336>. Citations in this document: §6.5, §23.
- [69] Neal Koblitz, Subhabrata Samajder, Palash Sarkar, Subhadip Singha, *Concrete analysis of approximate Ideal-SIVP to decision Ring-LWE reduction* (2022). URL: <https://eprint.iacr.org/2022/275>. Citations in this document: §4.6.
- [70] Robert Krauthgamer (editor), *Proceedings of the twenty-seventh annual ACM-SIAM symposium on discrete algorithms, SODA 2016, Arlington, VA, USA, January 10–12, 2016*, SIAM, 2016. ISBN 978-1-61197-433-1. See [21].
- [71] Thijs Laarhoven, *Sieving for closest lattice vectors (with preprocessing)*, in SAC 2016 [18] (2016), 523–542. Citations in this document: §3.2, §3.2.

- [72] Arjen K. Lenstra, Hendrik W. Lenstra, Jr., László Lovász, *Factoring polynomials with rational coefficients*, *Mathematische Annalen* **261** (1982), 515–534. ISSN 0025-5831. MR 84a:12002. URL: https://openaccess.leidenuniv.nl/bitstream/handle/1887/3810/346_050.pdf. Citations in this document: §3.1.
- [73] Richard Lindner, Chris Peikert, *Better key sizes (and attacks) for LWE-based encryption*, in CT-RSA [65] (2011), 319–339. URL: <https://eprint.iacr.org/2010/613>. Citations in this document: §2, §8, §8.
- [74] Xianhui Lu, Yamin Liu, Zhenfei Zhang, Dingding Jia, Haiyang Xue, Jingnan He, Bao Li, *LAC: practical Ring-LWE based public-key encryption with byte-level modulus* (2018). URL: <https://eprint.iacr.org/2018/1009>. Citations in this document: §4.5.
- [75] Vadim Lyubashevsky, Chris Peikert, Oded Regev, *On ideal lattices and learning with errors over rings*, *Journal of the ACM* **60** (2013), Article 43, 35 pages. URL: <https://eprint.iacr.org/2012/230>. Citations in this document: §2, §6, §6, §2.3, §2.3, §2.4, §2.4, §4.6.
- [76] Mitsuru Matsui (editor), *Advances in cryptology—ASIACRYPT 2009, 15th international conference on the theory and application of cryptology and information security, Tokyo, Japan, December 6–10, 2009. proceedings*, 5912, Springer, 2009. ISBN 978-3-642-10365-0. See [95].
- [77] MATZOV, *Report on the security of LWE: improved dual lattice attack* (2022). URL: <https://doi.org/10.5281/zenodo.6412487>. Citations in this document: §1.1, §1.3.
- [78] Ueli M. Maurer (editor), *Advances in cryptology—EUROCRYPT ’96: proceedings of the fifteenth international conference on the theory and application of cryptographic techniques held in Saragossa, May 12–16, 1996*, *Lecture Notes in Computer Science*, 1070, Springer, 1996. ISBN 3-540-61186-X. MR 97g:94002. See [23].
- [79] Alexander May, Joseph H. Silverman, *Dimension reduction methods for convolution modular lattices*, in [93] (2001), 110–15. URL: <https://www.cits.ruhr-uni-bochum.de/personen/may/publications.html>. Citations in this document: §2.4.
- [80] Alfred Menezes (editor), *Advances in cryptology—CRYPTO 2007, 27th annual international cryptology conference, Santa Barbara, CA, USA, August 19–23, 2007, proceedings*, *Lecture Notes in Computer Science*, 4622, Springer, 2007. ISBN 978-3-540-74142-8. See [60].
- [81] Shiho Moriai, Huaxiong Wang (editors), *Advances in cryptology—ASIACRYPT 2020—26th international conference on the theory and application of cryptology and information security, Daejeon, South Korea, December 7–11, 2020, proceedings, part II*, 12492, Springer, 2020. ISBN 978-3-030-64833-6. See [8].
- [82] Jesper Buus Nielsen, Vincent Rijmen (editors), *Advances in cryptology—EUROCRYPT 2018—37th annual international conference on the theory and applications of cryptographic techniques, Tel Aviv, Israel, April 29–May 3, 2018 proceedings, part I*, 10820, Springer, 2018. ISBN 978-3-319-78380-2. See [45].
- [83] NIST PQC team, *Announcement: The end of the 3rd round - the first PQC algorithms to be standardized* (2022). URL: <https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/G0DoD7lkGPK/m/f3Hl0sh3AgAJ>. Citations in this document: §3.
- [84] Abderrahmane Nitaj, Amr M. Youssef (editors), *Progress in cryptology—AFRICACRYPT 2020—12th international conference on cryptology*

- in Africa, Cairo, Egypt, July 20–22, 2020, proceedings*, 12174, Springer, 2020. ISBN 978-3-030-51937-7. See [44].
- [85] NTRU Prime Risk-Management Team, *Risks of lattice KEMs* (2021). URL: <https://ntruprime.cr.yp.to/warnings.html>. Citations in this document: §7, §3.
- [86] Alice Pellet-Mary, Guillaume Hanrot, Damien Stehlé, *Approx-SVP in ideal lattices with pre-processing*, in Eurocrypt 2019 [61] (2019), 685–716. URL: <https://eprint.iacr.org/2019/215>. Citations in this document: §5.1.
- [87] Ray Perlner, *Re: post-quantum benchmarking and RNGs* (2017). URL: <https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/PW6GF-wHGFE/m/nnAnuUAUAAAJ>. Citations in this document: §6.4.
- [88] Ray Perlner, *Requirements for security against multi-target attacks, for McEliece and other code-based cryptosystems?* (2021). URL: <https://web.archive.org/web/20210717065041/https://crypto.stackexchange.com/questions/92073/requirements-for-security-against-multi-target-attacks-for-mceliece-and-other-c>. Citations in this document: §5.3.
- [89] Herbert Robbins, *A remark on Stirling’s formula*, The American Mathematical Monthly **62** (1955), 26–29. Citations in this document: §4.2.
- [90] The Sage Developers (editor), *SageMath, the Sage Mathematics Software System (Version 9.5)*, 2022. URL: <https://www.sagemath.org>. Citations in this document: §4.2.
- [91] Nicolas Sendrier, *Decoding one out of many*, in PQCrypto 2011 [99] (2011), 51–67. URL: <https://eprint.iacr.org/2011/367>. Citations in this document: §13, §5.3.
- [92] David B. Shmoys (editor), *Symposium on theory of computing, STOC 2014, New York, NY, USA, May 31–June 03, 2014*, ACM, 2014. ISBN 978-1-4503-2710-7. See [52].
- [93] Joseph H. Silverman (editor), *Cryptography and lattices: proceedings of the 1st International Conference (CaLC 2001) held in Providence, RI, March 29–30, 2001*, Lecture Notes in Computer Science, 2146, Springer, 2001. ISBN 3-540-42488-1. MR 2002m:11002. See [79].
- [94] Naftali Sommer, Meir Feder, Ofir Shalvi, *Finding the closest lattice point by iterative slicing*, in [1] (2007), 206–210. Citations in this document: §11, §3.2.
- [95] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, Keita Xagawa, *Efficient public key encryption based on ideal lattices*, in Asiacrypt 2009 [76] (2009), 617–635. URL: <https://eprint.iacr.org/2009/285>. Citations in this document: §2.4, §2.4.
- [96] Tsuyoshi Takagi, Thomas Peyrin (editors), *Advances in cryptology—ASIACRYPT 2017—23rd international conference on the theory and applications of cryptology and information security, Hong Kong, China, December 3–7, 2017, proceedings, part II*, Lecture Notes in Computer Science, 10625, Springer, 2017. ISBN 978-3-319-70696-2. See [9].
- [97] Mehdi Tibouchi, Huaxiong Wang (editors), *Advances in cryptology—ASIACRYPT 2021—27th international conference on the theory and application of cryptology and information security, Singapore, December 6–10, 2021, proceedings, part IV*, 13093, Springer, 2021. ISBN 978-3-030-92067-8. See [53].
- [98] Herbert S. Wilf, *generatingfunctionology*, Academic Press, 1994. URL: <https://www2.math.upenn.edu/~wilf/DownldGF.html>. Citations in this document: §4.2.
- [99] Bo-Yin Yang (editor), *Post-quantum cryptography—4th international workshop, PQCrypto 2011, Taipei, Taiwan, November 29–December 2, 2011. proceedings*, 7071, Springer, 2011. ISBN 978-3-642-25404-8. See [91].