

Partially Splitting Rings for Faster Lattice-Based Zero-Knowledge Proofs

Vadim Lyubashevsky and Gregor Seiler

IBM Research – Zurich

Abstract. When constructing practical zero-knowledge proofs based on the hardness of the Ring-LWE or the Ring-SIS problems over polynomial rings $\mathbb{Z}_p[X]/(X^n + 1)$, it is often necessary that the challenges come from a set \mathcal{C} that satisfies three properties: the set should be large (around 2^{256}), the elements in it should have small norms, and all the non-zero elements in the difference set $\mathcal{C} - \mathcal{C}$ should be invertible. The first two properties are straightforward to satisfy, while the third one requires us to make efficiency compromises. We can either work over rings where the polynomial $X^n + 1$ only splits into two irreducible factors modulo p , which makes the speed of the multiplication operation in the ring sub-optimal; or we can limit our challenge set to polynomials of smaller degree, which requires them to have (much) larger norms.

In this work we show that one can use the optimal challenge sets \mathcal{C} and still have the polynomial $X^n + 1$ split into more than two factors. This comes as a direct application of our more general result that states that all non-zero polynomials with “small” coefficients in the cyclotomic ring $\mathbb{Z}_p[X]/(\Phi_m(X))$ are invertible modulo p (where “small” depends on the size of p and how many irreducible factors the m^{th} cyclotomic polynomial $\Phi_m(X)$ splits into). We furthermore establish sufficient conditions for p under which $\Phi_m(X)$ will split in such fashion.

For the purposes of implementation, if the polynomial $X^n + 1$ splits into k factors, we can run FFT for $\log k$ levels until switching to Karatsuba multiplication. Experimentally, we show that increasing the number of levels from one to three or four results in a speedup by a factor of $\approx 2 - 3$. We point out that this improvement comes completely for free simply by choosing a modulus p that has certain algebraic properties. In addition to the speed improvement, having the polynomial split into many factors has other applications – e.g. when one embeds information into the Chinese Remainder representation of the ring elements, the more the polynomial splits, the more information one can embed into an element.

1 Introduction

Cryptography based on the presumed hardness of the Ring / Module-SIS and Ring / Module-LWE problems [20, 21, 16, 18, 13] is seen as a very likely replacement of traditional cryptography after the eventual coming of quantum computing. There already exist very efficient basic public key primitives, such as encryption schemes and digital signatures, based on the hardness of these problems. For added efficiency, most practical lattice-based constructions work over polynomial rings $\mathbb{Z}_p[X]/(f(X))$ where $f(X)$ is the cyclotomic polynomial $f(X) = X^n + 1$ and p is chosen in such a way that the $X^n + 1$ splits into n linear factors modulo p . With such a choice of parameters, multiplication in the ring can be performed very efficiently via the Number Theoretic Transform, which is an analogue of the Fast Fourier Transform that works over a finite field. Some examples of practical implementations that utilize NTT implementations of digital signatures and public key encryption based on the Ring-LWE problem can be found in [11, 22, 1, 7, 10].

Constructions of more advanced lattice-based primitives sometimes require that the underlying ring has additional properties. In particular, *practical* protocols that utilize zero-knowledge proofs (e.g. [4, 2, 17, 3, 9]) require that elements with small coefficients are invertible. This restriction,

which precludes using rings where $X^n + 1$ splits completely modulo p , stems from the structure of *approximate* zero-knowledge proofs, and we sketch this intuition below.

1.1 Approximate Zero-Knowledge Proofs

Abstractly, in a zero-knowledge proof the prover wants to prove the knowledge of s that satisfies the relation $f(s) = t$, where f and t are public. In the lattice setting, the function

$$f(s) := As \tag{1}$$

where A is a random matrix over some ring (the ring is commonly \mathbb{Z}_p or $\mathbb{Z}_p[X]/(X^n + 1)$) and s is a vector over that same ring, where the coefficients of all (or almost all) the elements comprising s are bounded by some small value $\ll p$.

The function f in (1) satisfies the property that $f(s_1) + f(s_2) = f(s_1 + s_2)$ and for any c in the ring and any vector s over the ring we have $f(sc) = c \cdot f(s)$. The zero-knowledge proof for attempting to prove the knowledge of s proceeds as follows:

The Prover first chooses a “masking parameter” y and sends $w := f(y)$ to the Verifier. The Verifier picks a random challenge c from a subset of the ring and sends it to the prover (in a non-interactive proof, the Prover himself would generate $c := H(t, w)$, where H is a cryptographic hash function). The Prover then computes $z := sc + y$ and sends it to the Verifier.¹

The Verifier checks that $f(z) = ct + w$ and, crucially, it also checks to make sure that the coefficients of z are small. If these checks pass, then the Verifier accepts the proof. To show that the protocol is a proof of knowledge, one can rewind the Prover to just after his first move and send a different challenge c' , and get a response z' such that $f(z') = c't + w$. Combined with the first response, we extract the equation

$$f(\bar{s}) = \bar{c}t \tag{2}$$

where $\bar{s} = z - z'$ and $\bar{c} = c - c'$.

Notice that while the prover started with the knowledge of an s with small coefficients such that $f(s) = t$, he only ends up proving the knowledge of an \bar{s} with larger coefficients such that $f(\bar{s}) = \bar{c}t$. If \bar{c} also has small coefficients, then this type of proof is good enough in many (but not all) situations.

Applications of Approximate Zero-Knowledge Proofs. As a simple example of the utility of approximate zero-knowledge proofs, we consider commitment schemes where a commitment to a message m involves choosing some randomness r , and outputting $f(s) = t$, where s is defined as $\begin{bmatrix} r \\ m \end{bmatrix}$ where r and m have small coefficients.² Using the zero-knowledge proof from Section 1.1, one can prove the knowledge of an \bar{s} and \bar{c} such that $f(\bar{s}) = \bar{c}t$. If \bar{c} is invertible in the ring, then we can argue that this implies that if t is later opened to any valid commitment s' where $f(s') = t$, then it must be $s' = \bar{s}/\bar{c}$.

¹ In lattice-based schemes, it is important to keep the coefficients of z small, and so y must be chosen to have small coefficients as well. This can lead to the distribution of z being dependent on sc , which leaks some information about s . This problem is solved in [14, 15] via various rejection-sampling procedures. How this is done is not important to this paper, and so we ignore this step.

² It was shown in [4, 2] that one actually does not need the message m to have small coefficients, but for simplicity we assume here that it still has them.

The sketch of the argument is as follows: If we extract \bar{s}, \bar{c} and the commitment is opened with s' such that $f(s') = t$, then multiplying both sides by \bar{c} results in $f(\bar{c}s') = \bar{c}t$. Combining this with what was extracted from the zero-knowledge proof, we obtain that $f(\bar{c}s') = f(\bar{s})$. If $s' \neq \bar{s}/\bar{c}$, then $\bar{c}s' \neq \bar{s}$ and we found a collision (with small coefficients) for the function f . Such a collision implies a solution to the (Ring-)SIS problem, or, depending on the parameters, may simply not exist (and the scheme can thus be based on (Ring-)LWE).

There are more intricate examples involving commitment schemes (see e.g. [4, 2]) as well as other applications of such zero knowledge proofs, (e.g. to verifiable encryption [17] and voting protocols [9]) which require that the \bar{c} be invertible.

The Challenge Set and its Effect on the Proof. The challenge c is drawn uniformly from some domain \mathcal{C} which is a subset of $\mathbb{Z}_p[X]/(X^n + 1)$. In order to have small soundness error, we would like \mathcal{C} to be large. When building non-interactive schemes that should remain secure against quantum computers, one should have $|\mathcal{C}|$ be around 2^{256} . On the other hand, we also would like c to have a small norm. The reason for the latter is that the honest prover computes $z := sc + y$ and so the \bar{s} that is extracted from the Prover in (2) is equal to $z - z'$, and must also therefore depend on $\|sc\|$. Thus, the larger the norms of c, c' are, the larger the extracted solution \bar{s} will be, and the easier the corresponding (Ring-)SIS problem will be.

As a running example, suppose that we're working over the polynomial ring $\mathbb{Z}_p[X]/(X^{256} + 1)$. If invertibility were not an issue, then a simple and nearly optimal way to choose \mathcal{C} of size 2^{256} would be to define

$$\mathcal{C} = \{c \in R_p^{256} : \|c\|_\infty = 1, \|c\|_1 = 60\}. \quad (3)$$

In other words, the challenges are ring elements consisting of exactly 60 non-zero coefficients which are ± 1 .³ The l_2 norm of such elements is $\sqrt{60}$.

If we take invertibility into consideration, then we need the difference set $\mathcal{C} - \mathcal{C}$ (excluding 0) to consist only of invertible polynomials. There are some folklore ways of creating such a set. If the polynomial $X^{256} + 1$ splits into k irreducible polynomials modulo p , then all of these polynomials must have degree $256/k$. It is then easy to see, via the Chinese Remainder Theorem that every non-zero polynomial of degree less than $256/k$ is invertible in the ring $\mathbb{Z}_p[X]/(X^{256} + 1)$. We can therefore define the set

$$\mathcal{C}' = \{c \in R_p^{256} : \deg(c) < 256/k, \|c\|_\infty \leq \gamma\},$$

where $\gamma \approx 2^{k-1}$ in order for the size of the set to be greater than 2^{256} . The l_2 norm of elements in this set is $\sqrt{256/k} \cdot \gamma$. If we, for example, take $k = 8$, then this norm becomes $\sqrt{32} \cdot 2^7 \approx 724$, which is around 90 times larger than the norms of the challenges in the set defined in (3). It is therefore certainly not advantageous to increase the norm of the challenge by this much only to decrease the running time of the computation. In particular, the security of the scheme will decrease and one will need to increase the ring dimension to compensate, which will in turn negate any savings in running time. A much more desirable solution would be to have the polynomial $X^n + 1$ split, but still be able to use the optimal challenge set from (3).

³ The size of this set is $\binom{256}{60} \cdot 2^{60} > 2^{256}$.

1.2 Our Contribution

Our main result is a general theorem (Theorem 1.1) about the invertibility of polynomials with small coefficients in polynomial rings $\mathbb{Z}_p[X]/(\Phi_m(X))$, where $\Phi_m(X)$ is the m^{th} cyclotomic polynomial. The theorem states that if a non-zero polynomial has small coefficients (where “small” is related to the prime p and the number of irreducible factors of $\Phi_m(X)$ modulo p), then it’s invertible in the ring $\mathbb{Z}_p[X]/(\Phi_m(X))$. For the particular case of $\Phi_m(X) = X^n + 1$, we show that the polynomial $X^n + 1$ can split into several (in practice up to 8 or 16) irreducible factors and we can still use the optimal challenge sets, like ones of the form from (3). This generalizes and extends a result in [17] which showed that one can use the optimal set when $X^n + 1$ splits into two factors. We also show some methods for creating challenge sets that are slightly sub-optimal, but allow for the polynomial to split further.

Theorem 1.1. *Let $m = \prod p_i^{e_i}$ for $e_i \geq 1$ and $z = \prod p_i^{f_i}$ for any $1 \leq f_i \leq e_i$. If p is a prime such that $p \equiv 1 \pmod{z}$ and $\text{ord}_m(p) = m/z$, then the polynomial $\Phi_m(X)$ factors as*

$$\Phi_m(X) = \prod_{j=1}^{\phi(z)} (X^{m/z} - r_j) \pmod{p}$$

for distinct $r_j \in \mathbb{Z}_p^*$ where $X^{m/z} - r_j$ are irreducible in the ring $\mathbb{Z}_p[X]$. Furthermore, any \mathbf{y} in $\mathbb{Z}_p[X]/(\Phi_m(X))$ that satisfies

$$0 < \|\mathbf{y}\|_\infty < \frac{1}{\sqrt{\tau(z)}} \cdot p^{1/\phi(z)}$$

has an inverse in the ring (where $\tau(z) = z$ if z is odd, and $z/2$ otherwise.).

The above theorem gives sufficient conditions for p so that all polynomials with small coefficients in $\mathbb{Z}_p[X]/(\Phi_m(X))$ are invertible, but it does not state anything about whether there exist such p . In Theorem 2.4, we show that if we additionally put the condition on m and z that $8|m \Rightarrow 4|z$, then there are indeed infinitely many primes p that satisfy these conditions. In practical lattice constructions involving zero-knowledge proofs, we would normally use a modulus of size at least 2^{20} , and we experimentally confirmed (for various cyclotomic polynomials) that one can indeed find many such primes that are of that size.

Specializing the above to the ring $\mathbb{Z}_p[X]/(X^n + 1)$, we obtain the following corollary:

Corollary 1.2. *Let $n \geq k > 1$ be powers of 2 and $p = 2k + 1 \pmod{4k}$ be a prime. Then the polynomial $X^n + 1$ factors as $X^n + 1 = \prod_{j=1}^k (X^{n/k} - r_j) \pmod{p}$, and any \mathbf{y} in the ring $\mathbb{Z}_p[X]/(X^n + 1)$ that satisfies $0 < \|\mathbf{y}\|_\infty < \frac{1}{\sqrt{k}} \cdot p^{1/k}$ has an inverse in the ring.*

As an application of this result, suppose that we choose $k = 8$ and a prime p congruent to $17 \pmod{32}$ such that $p > 2^{20}$. Furthermore, suppose that we perform our zero-knowledge proofs over the ring $\mathbb{Z}_p[X]/(X^n + 1)$ (where n is a power of 2 greater than 8), and prove the knowledge of \bar{s}, \bar{c} such that $f(\bar{s}) = \bar{c}t$ where $\|\bar{c}\|_\infty \leq 2$ (i.e. the challenges c are taken such that $\|c\|_\infty = 1$). Then the above theorem states that $X^n + 1$ factors into 8 polynomials and \bar{c} will be invertible in the ring since $\frac{1}{\sqrt{8}} \cdot p^{1/8} > 2$.

Having $p > 2^{20}$ is quite normal for the regime of zero-knowledge proofs, and therefore having the polynomial $X^n + 1$ split into 8 factors should be possible in virtually every application. If we would like it to split further into 16 or 32 factors, then we would need $p > 2^{48}$ or, respectively, $p > 2^{112}$. In Section 3.3 we describe how our techniques used to derive Theorem 1.1 can also be used in a somewhat “ad-hoc” fashion to create different challenge sets \mathcal{C} that are nearly-optimal (in terms of the maximal norm), but allow $X^n + 1$ to split with somewhat smaller moduli than implied by Theorem 1.1.

In Section 4, we describe how one would combine the partially-splitting FFT algorithm with a Karatsuba multiplication algorithm to efficiently multiply in a partially-splitting ring. For primes of size between 2^{20} and 2^{29} , one obtains a speed-up of about a factor of 2 by working over rings where $X^n + 1$ splits into 8 versus just 2 factors.

In addition to the speed improvement, there are applications whose usability can be improved by the fact that we work over rings $\mathbb{Z}_p[X]/(X^n + 1)$ where $X^n + 1$ splits into more factors. For example, [4] constructed a commitment scheme and zero-knowledge proofs of knowledge that allows to prove the fact that $\mathbf{c} = \mathbf{a}\mathbf{b}$ when $\text{Commit}(\mathbf{a})$, $\text{Commit}(\mathbf{b})$, $\text{Commit}(\mathbf{c})$ are public (the same holds for addition). An application of this result is the verifiability of circuits. For this application, one only needs commitments of 0’s and 1’s, thus if we work over a ring where $X^n + 1$ splits into k irreducible factors, one can embed k bits into each Chinese Remainder coefficient of \mathbf{a} and \mathbf{b} , and therefore proving that $\mathbf{c} = \mathbf{a}\mathbf{b}$ implies that all k multiplications of the bits were performed correctly. Thus the larger k is, the more multiplications one can prove in parallel. Unfortunately k cannot be set too large without ruining the necessary property that the difference of any two distinct challenges is invertible or increasing the ℓ_2 -norm of the challenges as described in Section 1.1. Our result therefore allows to prove products of 8 (or 16) commitments in parallel without having to increase the parameters of the scheme to accommodate the larger challenges.

Acknowledgements

We thank Rafaël del Pino for pointing out an improvement to Lemma 3.3. This work is supported by the SNSF ERC Transfer Grant CRETP2-166734 – FELICITY and the H2020 Project Safecrypto.

2 Cyclotomics and Lattices

2.1 Cyclotomic Polynomials

Definition 2.1. For any integer $m > 1$, we write

$$\begin{aligned} \phi(m) &= m \cdot \prod_{p \text{ is prime } \wedge p|m} \frac{p-1}{p} \\ \delta(m) &= \prod_{p \text{ is prime } \wedge p|m} p \\ \tau(m) &= \begin{cases} m, & \text{if } m \text{ is odd} \\ m/2, & \text{if } m \text{ is even} \end{cases} \\ \text{ord}_m(n) &= \min\{k : k > 0 \text{ and } n^k = 1 \pmod{m}\} \end{aligned}$$

The function $\phi(m)$ is the Euler phi function, $\delta(m)$ is sometimes referred to as the *radical* of m , and $\tau(m)$ is a function that sometimes comes into play when working with the geometry of cyclotomic rings. The function $\text{ord}_m(n)$ is the order of an element n in the multiplicative group \mathbb{Z}_m^* . In the special case of $m = 2^k$, we have $\phi(m) = \tau(m) = 2^{k-1}$ and $\delta(m) = 2$.

The m^{th} cyclotomic polynomial, written as $\Phi_m(X)$, is formally defined to be

$$\Phi_m(X) = \prod_{i=1}^{\phi(m)} (X - \omega_i),$$

where ω_i are the m^{th} complex primitive roots of unity (of which there are $\phi(m)$ many). Of particular interest in practical lattice cryptography is the cyclotomic polynomial $\Phi_{2^k}(X) = X^{2^{k-1}} + 1$.

If p is some prime and $r_1, \dots, r_{\phi(m)}$ are elements in \mathbb{Z}_p^* such that $\text{ord}_p(r_j) = \phi(m)$, then one can write

$$\Phi_m(X) = \prod_{j=1}^{\phi(m)} (X - r_j) \pmod{p}.$$

For any $m > 1$, it is known that we can express the cyclotomic polynomial $\Phi_m(X)$ as

$$\Phi_m(X) = \Phi_{\delta(m)}\left(X^{m/\delta(m)}\right), \quad (4)$$

and the below Lemma is a generalization of this statement.

Lemma 2.2. *Let $m = \prod p_i^{e_i}$ for $e_i \geq 1$ and $z = \prod p_i^{f_i}$ for any $1 \leq f_i \leq e_i$. Then*

$$\Phi_m(X) = \Phi_z(X^{m/z}).$$

Proof. By (4), and the fact that $\delta(m) = \delta(z)$, we can rewrite $\Phi_m(X)$ as

$$\begin{aligned} \Phi_m(X) &= \Phi_{\delta(m)}(X^{m/\delta(m)}) = \Phi_{\delta(m)}(X^{z/\delta(m)})(X^{m/z}) \\ &= \Phi_{\delta(z)}(X^{z/\delta(z)})(X^{m/z}) = \Phi_z(X^{m/z}). \end{aligned} \quad (5)$$

□

2.2 The Splitting of Cyclotomic Polynomials

In Theorem 2.3, we give the conditions on the prime p such that the polynomial $\Phi_m(X)$ splits into irreducible factors $X^{m/k} - r$ modulo p . In Theorem 2.4, we then show that when m and k satisfy an additional relation, there are infinitely many p that satisfy the necessary conditions of Theorem 2.3.

Theorem 2.3. *Let $m = \prod p_i^{e_i}$ for $e_i \geq 1$ and $z = \prod p_i^{f_i}$ for any $1 \leq f_i \leq e_i$. If p is a prime such that $p \equiv 1 \pmod{z}$ and $\text{ord}_m(p) = m/z$, then the polynomial $\Phi_m(X)$ factors as*

$$\Phi_m(X) = \prod_{j=1}^{\phi(z)} (X^{m/z} - r_j) \pmod{p}$$

for distinct $r_j \in \mathbb{Z}_p^*$ where $X^{m/z} - r_j$ are irreducible in $\mathbb{Z}_p[X]$.

Proof. Since p is a prime and $p \equiv 1 \pmod{z}$, there exists an element r such that $\text{ord}_p(r) = z$. Furthermore, for all the $\phi(z)$ integers $1 < i < z$ such that $\gcd(i, z) = 1$, we also have $\text{ord}_p(r^i) = z$. We therefore have, by definition of Φ , that

$$\Phi_z(X) = \prod_{j=1}^{\phi(z)} (X - r_j) \pmod{p}.$$

Applying Lemma 2.2, we obtain that

$$\Phi_m(X) = \prod_{j=1}^{\phi(z)} (X^{m/z} - r_j) \pmod{p}.$$

We now need to prove that the terms $X^{m/z} - r_j$ are irreducible modulo p . Suppose they are not and $X^{m/z} - r_j$ has an irreducible divisor f of degree $d < \frac{m}{z}$. Then f defines an extension field of \mathbb{Z}_p of degree d , i.e. a finite field with p^d elements that all satisfy $X^{p^d} = X$. Hence f divides $X^{p^d} - X$. Now, from $\text{ord}_m(p) = \frac{m}{z} > d$ it follows that we can write $p^d = am + b$ where $b \neq 1$. Thus

$$X^{p^d} - X = X^{am+b} - X = X(X^{am+(b-1)} - 1).$$

If we now consider an extension field of \mathbb{Z}_p in which f splits, the roots of f are also roots of $X^{am+(b-1)} - 1$ and therefore have order dividing $am + (b - 1)$. This is a contradiction. As a divisor of $X^{m/z} - r_j$ (and therefore of Φ_m), f has only roots of order m .

Theorem 2.4. *Let $m = \prod p_i^{e_i}$ for $e_i \geq 1$ and $z = \prod p_i^{f_i}$ for any $1 \leq f_i \leq e_i$. Furthermore, assume that if m is divisible by 8, then z is divisible by 4. Then there are infinitely many primes p such that $p \equiv 1 \pmod{z}$ and $\text{ord}_m(p) = m/z$.*

Proof. First we show that an integer not necessarily prime exists that fulfills the two conditions. By the Chinese remainder theorem it suffices to find integers a_i such that $a_i \pmod{p_i^{f_i}} = 1$ and $\text{ord}_{p_i^{e_i}}(a_i) = p_i^{e_i - f_i}$. First consider the odd primes $p_i \neq 2$. It is easy to show that if g is a generator modulo p_i then either g or $g + p_i$, say g' , is a generator modulo every power of p_i (c.f. [8, Lemma 1.4.5]). Define $a_i = (g')^{(p_i-1)p_i^{f_i-1}}$. Then, since g' has order $(p_i - 1)p_i^{f_i-1}$ modulo $p_i^{f_i}$ and order $(p_i - 1)p_i^{e_i-1} \pmod{p_i^{e_i}}$, it follows that $a_i \pmod{p_i^{f_i}} = 1$ and

$$\text{ord}_{p_i^{e_i}}(a_i) = \frac{(p_i - 1)p_i^{e_i-1}}{(p_i - 1)p_i^{f_i-1}} = p_i^{e_i - f_i}$$

as we wanted. Next, consider $p = 2$ and the case where m is divisible by 8; that is, $e_1 \geq 3$. This implies $f_1 \geq 2$. It is a standard fact that 5 is a generator of a cyclic subgroup of \mathbb{Z}_{2^e} of index 2 for every $e \geq 3$, i.e. $\text{ord}_{2^e}(5) = 2^{e-2}$. Therefore, $5^{2^{f_1-2}} \pmod{2^{f_1}} = 1$ and

$$\text{ord}_{2^{e_1}}(5^{2^{f_1-2}}) = \frac{2^{e_1-2}}{2^{f_1-2}} = 2^{e_1 - f_1}.$$

Hence $a_1 = 5^{2^{f_1-2}}$ is a valid choice in this case. If $e_1 = 2$, note that 3 is a generator modulo 4 and $a_1 = 3^{2^{f_1-1}}$ is readily seen to work. When $e_1 = f_1 = 1$, take $e_1 = 1$. So, there exists an integer a that fulfills our two conditions and in fact every integer congruent to $a \pmod{m}$ does. By Dirichlet's theorem on arithmetic progressions, there are infinitely many primes among the $a + lm$ ($l \in \mathbb{Z}$). \square

As an experimental example consider $m = 2^2 3^3 7 = 756$ and $z = 2 \cdot 3 \cdot 7 = 42$. Then Φ_m splits into 12 polynomials modulo primes of the form in Theorem 2.4. There are 2058 primes of this form between 2^{20} and 2^{21} .

2.3 Cyclotomic Rings and Ideal Lattices

Throughout the paper, we will write R_m to be the *cyclotomic ring* $\mathbb{Z}[X]/(\Phi_m(X))$ and $R_{m,p}$ to be the ring $\mathbb{Z}_p[X]/(\Phi_m(X))$, with the usual polynomial addition and multiplication operations. We will denote by normal letters elements in \mathbb{Z} and by bold letters elements in R_m . For an odd p , an element $\mathbf{w} \in R_{m,p}$ can always be written as $\sum_{i=0}^{\phi(m)-1} w_i X^i$ where $|w_i| \leq (p-1)/2$. Using this representation, for $\mathbf{w} \in R_{m,p}$ (and in R_m), we will define the lengths of elements as

$$\|\mathbf{w}\|_\infty = \max_i |w_i| \text{ and } \|\mathbf{w}\| = \sqrt{\sum_i |w_i|^2}.$$

An integer lattice of dimension n is an additive sub-group of \mathbb{Z}^n . For the purposes of this paper, all lattices will be full-rank. The determinant of a full-rank integer lattice Λ of dimension n is the size of the quotient group $|\mathbb{Z}^n/\Lambda|$. If \mathbf{z} is a non-zero vector in \mathbb{Z}^n , then it's easy to see that the lattice

$$\Lambda = \{\mathbf{y} \in \mathbb{Z}^n : \langle \mathbf{y}, \mathbf{z} \rangle \bmod p = 0\}$$

is full-rank and has determinant p . We write $\lambda_1(\Lambda)$ to denote the Euclidean length of the shortest non-zero vector in Λ .

If \mathcal{I} is an ideal in the polynomial ring R_m , then it is also an additive sub-group of $\mathbb{Z}^{\phi(m)}$, and therefore a $\phi(m)$ -dimensional lattice (it can be shown that such lattices are always full-rank). Such lattices are therefore sometimes referred to as *ideal lattices*. For an ideal lattice Λ of the ring R_m , there exists a lower bound on the length of its shortest vector. The below lemma is a simple combination of several facts from [19], and we sketch its proof for completeness.

Lemma 2.5. *If Λ is an ideal lattice in R_m , then*

$$\lambda_1(\Lambda) \geq \sqrt{\frac{\phi(m)}{\tau(m)}} \cdot \det(\Lambda)^{1/\phi(m)}.$$

Proof. Let $\omega_1, \dots, \omega_{\phi(m)}$ be the complex roots of $\Phi_m(X)$, and define the Vandermonde matrix $\mathbf{V} \in \mathbb{C}^{\phi(m) \times \phi(m)}$ where coefficients $\mathbf{V}_{i,j} = \omega_i^j$ for $1 \leq i \leq \phi(m)$ and $0 \leq j \leq \phi(m) - 1$. Let us also define the shortest vector in Λ relative to the *embedding norm* as

$$\lambda'_1(\Lambda) = \min_{\mathbf{w} \in \Lambda \setminus \{0\}} \|\mathbf{V}\mathbf{w}\|.$$

Then using the notation from the current paper, [19, Lemma 2.14] states that $\lambda'_1(\Lambda) \geq \sqrt{\phi(m)} \cdot \det(\Lambda)^{1/\phi(m)}$. In turn, [19, Lemma 4.3] states that for any $\mathbf{w} \in \mathbb{R}^{\phi(m)}$, we have $\|\mathbf{V}\mathbf{w}\| \leq \sqrt{\tau(m)} \cdot \|\mathbf{w}\|$. Combining the two inequalities implies that

$$\lambda_1(\Lambda) \geq \frac{\lambda'_1(\Lambda)}{\sqrt{\tau(m)}} \geq \sqrt{\frac{\phi(m)}{\tau(m)}} \cdot \det(\Lambda)^{1/\phi(m)}.$$

□

3 Invertible Elements in Cyclotomic Rings

The main goal of this section is to prove Theorem 1.1. To this end, we first prove Lemma 3.1 that can be seen as a special case of the Theorem when the polynomial $\Phi_m(X)$ splits completely modulo p .⁴ In Section 3.2 we consider rings $R_{m,p}$ where $\Phi_m(X)$ only partially splits modulo p and describe how to interpret polynomials $\mathbf{y} \in R_{m,p}$ as a combination of polynomials \mathbf{y}'_i over a smaller, but fully-splitting ring. We then prove in Lemma 3.2 that if any of the \mathbf{y}'_i is invertible in the fully-splitting ring, then the polynomial \mathbf{y} is invertible in $R_{m,p}$. The proof of Theorem 1.1 will follow from these two Lemmas and Theorem 2.3.

3.1 Fully-Splitting Rings

Lemma 3.1. *Let p and m be integers such that $\Phi_m(X) = \prod_{i=1}^{\phi(m)} (X - r_i) \pmod p$ for some distinct $r_i \in \mathbb{Z}_p^*$ and let \mathbf{y} be any element in the ring $R_{m,p}$. If $0 < \|\mathbf{y}\| < \sqrt{\frac{\phi(m)}{\tau(m)}} \cdot p^{1/\phi(m)}$, then \mathbf{y} is invertible in $R_{m,p}$.*

Proof. Suppose that \mathbf{y} is not invertible in $R_{m,p}$. By the Chinese Remainder Theorem, this implies that (for at least) one i , $\mathbf{y} \pmod{(X - r_i)} = \mathbf{y}(r_i)$ is 0 modulo p . For an i for which $\mathbf{y}(r_i) \pmod p = 0$, (if there is more than one such i , pick one of them arbitrarily) define the set

$$\Lambda = \{\mathbf{z} \in R_m : \mathbf{z}(r_i) \pmod p = 0\}.$$

Notice that Λ is an additive group and for any polynomial $\mathbf{z} \in \Lambda$, the polynomial $\mathbf{z} \cdot X^j \in R_m$ is also in Λ for any integer j . This implies that Λ is an ideal of R_m , and so an ideal lattice in the ring R_m . We now want to show that the determinant of Λ is p .

If we consider the polynomials $\mathbf{z} = \sum_{i=0}^{\phi(m)-1} z_i X^i \in R_m$ as vectors

$$\mathbf{z} = (z_0, z_1, \dots, z_{\phi(m)-1}) \in \mathbb{Z}^{\phi(m)},$$

and define the vector $\mathbf{r} = (1, r_i, r_i^2, \dots, r_i^{\phi(m)-1})$, then the lattice Λ can be rewritten as

$$\Lambda = \{\mathbf{z} \in \mathbb{Z}^{\phi(m)} : \langle \mathbf{z}, \mathbf{r} \rangle \pmod p = 0\},$$

which implies that $\det(\Lambda) = p$, and so by Lemma 2.5, $\lambda_1(\Lambda) \geq \sqrt{\frac{\phi(m)}{\tau(m)}} \cdot p^{1/\phi(m)}$.

Since we said that $\mathbf{y}(r_i) \pmod p = 0$ and $0 < \|\mathbf{y}\|$, we know that \mathbf{y} is a non-zero vector in Λ . But we also have that $\|\mathbf{y}\| < \sqrt{\frac{\phi(m)}{\tau(m)}} \cdot p^{1/\phi(m)} \leq \lambda_1(\Lambda)$, which is impossible. □

At this point, one might be tempted to prove Theorem 1.1 by a simple generalization of Lemma 3.1. The proof sketch would proceed as follows: suppose that m, z, p are as in the statement of Theorem 2.3 and so

$$\Phi_m(X) = \prod_{j=1}^{\phi(z)} (X^{m/z} - r_j) \pmod p.$$

⁴ This Lemma was already implicit in [23, Lemma 8] for $\Phi_m(X) = X^n + 1$.

Then one can define a lattice

$$\Lambda = \{\mathbf{z} \in R_m : \mathbf{z} \bmod (X^{m/z} - r_j) \bmod p = 0\},$$

and similarly conclude that Λ is an ideal lattice in R_m with $\det(\Lambda) = p^{m/z} = p^{\phi(m)/\phi(z)}$ and $\lambda_1(\Lambda) \geq \sqrt{\frac{\phi(m)}{\tau(m)}} \cdot p^{1/\phi(z)}$. This would in turn imply that any polynomial $\mathbf{y} \in R_{m,p}$ such that $0 < \|\mathbf{y}\| < \sqrt{\frac{\phi(m)}{\tau(m)}} \cdot p^{1/\phi(z)}$ is invertible. This gives a weaker bound in the ℓ_∞ norm than what is claimed in Theorem 1.1 – we can only conclude that all vectors \mathbf{y} such that $\|\mathbf{y}\|_\infty < \frac{1}{\sqrt{\tau(m)}} \cdot p^{1/\phi(z)}$ are invertible.

Since $z \ll m$, generalizing Lemma 3.1 to rings $R_{m,p}$ where $\Phi_m(X)$ only “partially splits” is therefore a sub-optimal approach for achieving the tightest bounds. In Section 3.2, we instead prove a lemma showing that only some parts of \mathbf{y} , which happen to correspond to elements of the smaller ring $R_{z,p}$, need to be invertible in $R_{z,p}$ in order for the entire element \mathbf{y} to be invertible in $R_{m,p}$.

3.2 Partially-Splitting Rings

In this section, we will be working with rings $R_{m,p}$ where p is chosen such that the polynomial $\Phi_m(X)$ factors into k irreducible polynomials of the form $X^{\phi(m)/k} - r_i$. Theorem 2.3 states the sufficient conditions on m, k, p in order to obtain such a factorization. Throughout this section, we will use the following notation: suppose that

$$\mathbf{y} = \sum_{j=0}^{\phi(m)-1} y_j X^j$$

is an element of the ring $R_{m,p}$, where the value p is chosen as above. Then for all integers $0 \leq i < \phi(m)/k - 1$, we define the polynomials \mathbf{y}'_i as

$$\mathbf{y}'_i = \sum_{j=0}^{k-1} y_{j\phi(m)/k+i} X^j. \quad (6)$$

For example, if $\phi(m) = 8$ and $k = 4$, then for $\mathbf{y} = \sum_{i=0}^7 y_i X^i$, we have $\mathbf{y}'_0 = y_0 + y_2 X + y_4 X^2 + y_6 X^3$ and $\mathbf{y}'_1 = y_1 + y_3 X + y_5 X^2 + y_7 X^3$.

The intuition behind the definition in (6) is that one can write \mathbf{y} in terms of the \mathbf{y}'_i as

$$\mathbf{y} = \sum_{i=0}^{\phi(m)/k-1} \mathbf{y}'_i (X^{\phi(m)/k}) \cdot X^i.$$

Then to calculate $\mathbf{y} \bmod (X^{\phi(m)/k} - r_j)$ where $(X^{\phi(m)/k} - r_j)$ is one of the irreducible factors of $\Phi_m(X)$ modulo p , we have

$$\mathbf{y} \bmod (X^{\phi(m)/k} - r_j) = \sum_{i=0}^{\phi(m)/k-1} \mathbf{y}'_i(r_j) \cdot X^i \quad (7)$$

simply because we plug in r_j for every $X^{\phi(m)/k}$.

Lemma 3.2. Let $m = \prod p_i^{e_i}$ for $e_i \geq 1$ and $z = \prod p_i^{f_i}$ for any $1 \leq f_i \leq e_i$, and suppose that we can write

$$\Phi_m(X) = \prod_{j=1}^{\phi(z)} (X^{m/z} - r_j) \pmod{p} \quad (8)$$

for distinct $r_j \in \mathbb{Z}_p^*$ where $(X^{m/z} - r_j)$ are irreducible in $\mathbb{Z}_p[X]$. Let \mathbf{y} be a polynomial in $R_{m,p}$ and define the associated \mathbf{y}'_i as in (6), where $k = \phi(z)$. If some \mathbf{y}'_i is invertible in $R_{z,p}$, then \mathbf{y} is invertible in $R_{m,p}$.

Proof. By the Chinese Remainder Theorem, the polynomial \mathbf{y} is invertible in $R_{m,p}$ if and only if $\mathbf{y} \pmod{(X^{m/z} - r_j)} \neq 0$ for all r_1, \dots, r_k . When we use $k = \phi(z)$, (7) can be rewritten as

$$\mathbf{y} \pmod{(X^{m/z} - r_j)} = \sum_{i=0}^{m/z-1} \mathbf{y}'_i(r_j) \cdot X^i.$$

To show that \mathbf{y} is invertible, it is therefore sufficient to show that

$$\exists i \text{ s.t. } \forall j, \mathbf{y}'_i(r_j) \pmod{p} \neq 0.$$

Let i be such that \mathbf{y}'_i is invertible in the ring $R_{z,p}$. From (8) and Lemma 2.2 we have that

$$\Phi_z(X) = \prod_{j=1}^{\phi(z)} (X - r_j) \pmod{p},$$

and so the ring $R_{z,p}$ is fully-splitting. Since \mathbf{y}'_i is invertible in $R_{z,p}$, the Chinese Remainder Theorem implies that for all $1 \leq j \leq \phi(z)$, $\mathbf{y}'_i(r_j) \pmod{p} \neq 0$, and therefore \mathbf{y} is invertible in $R_{m,p}$. \square

Theorem 1.1 now follows from the combination of Theorem 2.3, and Lemmas 3.1 and 3.2.

Proof (Theorem 1.1). For the conditions on m, z , and p , it follows from Theorem 2.3 that the polynomial $\Phi_m(X)$ can be factored into irreducible factors modulo p as $\prod_{j=1}^{\phi(z)} (X^{m/z} - r_j)$. Lemma

2.2 then states that $\Phi_z(X) = \prod_{j=1}^{\phi(z)} (X - r_j) \pmod{p}$.

For any $\mathbf{y} \in R_{m,p}$, let the \mathbf{y}'_i be defined as in (6) where $k = \phi(z)$. If $0 < \|\mathbf{y}\|_\infty < \frac{1}{\sqrt{\tau(z)}} \cdot p^{1/\phi(z)}$, then because each \mathbf{y}'_i consists of $\phi(z)$ coefficients, we have that for all i , $\|\mathbf{y}'_i\| < \sqrt{\frac{\phi(z)}{\tau(z)}} \cdot p^{1/\phi(z)}$. Since $\mathbf{y} \neq 0$, it must be that for some i , $\mathbf{y}'_i \neq 0$.

Lemma 3.1 therefore implies that the non-zero \mathbf{y}'_i is invertible in $R_{z,p}$. In turn, Lemma 3.2 implies that \mathbf{y} is invertible in $R_{m,p}$. \square

3.3 Example of “Ad-hoc” Applications of Lemma 3.2

Using Lemma 3.2, as we did in the proof of Theorem 1.1 above, gives a clean statement as to a sufficient condition under which polynomials are invertible in a partially-splitting ring. One thing to note is that putting a bound on the ℓ_∞ norm does not take into account the other properties that our challenge space may have. For example, our challenge space in (3) is also sparse, in addition to having the ℓ_∞ norm bounded by 1. Yet we do not know how to use this sparseness to show that one can let $\Phi_m(X)$ split further while still maintaining the invertibility of the set $\mathcal{C} - \mathcal{C}$.

In some cases, however, there are ways to construct challenge sets that are more in line with Lemma 3.2 and will allow further splitting. We do not see a simple way in which to systematize these ideas, and so one would have to work out the details on a case-by-case basis. Below, we give such an example for the case in which we are working over the ring $\mathbb{Z}_p[X]/(X^{256} + 1)$ and would like to have the polynomial $X^{256} + 1$ split into 16 irreducible factors. If we would like to have $X^n + 1$ split into 16 factors modulo p and the set $\mathcal{C} - \mathcal{C}$ to have elements whose infinity norm is bounded by 2, then applying Theorem 1.1 directly implies that we need to have $2 < \frac{1}{\sqrt{16}} \cdot p^{1/16}$, which implies $p > 2^{48}$.

We will now show how one can lower the requirement on p in order to achieve a split into 16 factors by altering the challenge set \mathcal{C} in (3).

For a polynomial $\mathbf{y} \in \mathbb{Z}_p[X]/(X^{256} + 1)$, define the \mathbf{y}'_i as in (6). Define \mathcal{D} as

$$\mathcal{D} = \{\mathbf{y} \in \mathbb{Z}_p[X]/(X^{256} + 1) : \|\mathbf{y}_i\|_\infty = 1 \text{ and } \forall 1 \leq i \leq 16, \|\mathbf{y}'_i\| = 2\} \quad (9)$$

In other words, \mathcal{D} is the set of polynomials \mathbf{y} , such that every \mathbf{y}'_i has exactly 4 non-zero elements that are ± 1 . The size of \mathcal{D} is $\binom{16}{4} \cdot 2^4 \approx 2^{237}$, which should be enough for practical quantum security. The ℓ_2 norm of every element in \mathcal{D} is exactly $\sqrt{64} = 8$. For a fair comparison, we should redefine the set \mathcal{C} so that it also has size 2^{237} . The only change that one must make to the definition in (3) is to lower the ℓ_1 norm to 53 from 60. Thus all elements in \mathcal{C} have ℓ_2 norm $\sqrt{53}$. The elements in set \mathcal{D} therefore have norm that is larger by a factor of about 1.1. It then depends on the application as to whether having $X^n + 1$ split into 16 rather than 8 factors is worth this modest increase. We will now prove that for primes $p > 2^{30.5}$ of a certain form, $X^{256} + 1$ will split into 16 irreducible factors modulo p and all the non-zero elements in $\mathcal{D} - \mathcal{D}$ will be invertible. Therefore if our application calls for a modulus that is larger than $2^{30.5}$ but smaller than 2^{48} , we can use the challenge set \mathcal{D} and the below lemma.

Lemma 3.3. *Suppose that $p > 2^{16 \log_2 \sqrt{14}} \approx 2^{30.5}$ is a prime congruent to $33 \pmod{64}$. Then the polynomial $X^{256} + 1$ splits into 16 irreducible polynomials of the form $X^{16} + r_j$ modulo p , and any non-zero polynomial $\mathbf{y} \in \mathcal{D} - \mathcal{D}$ (as defined in (9)) is invertible in the ring $\mathbb{Z}_p[X]/(X^{256} + 1)$.*

Proof. The fact that $X^{256} + 1$ splits into 16 irreducible factors follows directly from Theorem 2.3. Notice that for any $\mathbf{y} \in \mathcal{D} - \mathcal{D}$, the maximum ℓ_2 norm of \mathbf{y}'_i is bounded by 4. Furthermore, the degree of each \mathbf{y}'_i is $256/16 = 16$. Thus an immediate consequence of Lemmas 3.2 and 3.1 is that if $p > 2^{32}$, then any non-zero element in $\mathcal{D} - \mathcal{D}$ is invertible. To slightly improve the lower bound, we can observe that the \mathbf{y}'_i of norm 4 are polynomials in $\mathbb{Z}_p[X]/(X^{16} + 1)$ with exactly four 2's in them. But such elements can be written as a product of 2 and a polynomial with 4 ± 1 's in it. So if both of those are invertible, so is the product. The maximum norm of these polynomials is 2 and so they are not the elements that set the lower bound. The next largest element in $\mathcal{D} - \mathcal{D}$ is one that has three 2's and two ± 1 's. The norm of such elements is $\sqrt{14}$. Thus for all $p > 2^{16 \cdot \log_2(\sqrt{14})} \approx 2^{30.5}$,

Number of FFT levels	Primes			
	$2^{20} - 2^{14} + 1$	$2^{23} - 2^{13} + 1$	$2^{25} - 2^{12} + 1$	$2^{27} - 2^{11} + 1$
0	123677	123717	134506	144913
1	83820	83778	91775	97641
2	55378	55700	63148	65778
3	38111	38061	43116	43282
4	27374	27626	31782	30836
5	21968	21955	26406	24937
6	17076	17007	21518	19811
7	15149	15144	20483	18026
8	16875	16893	22329	20299

Table 1. CPU cycles of our FFT-accelerated multiplication algorithm for $\mathbb{Z}_p[X]/(X^{256} + 1)$ using Karatsuba multiplication for the base case. Both the FFT and Karatsuba are plain C implementations.

the \mathbf{y}'_i will be invertible in $\mathbb{Z}_p[X]/(X^{16} + 1)$, and thus every non-zero element in $\mathcal{D} - \mathcal{D}$ will be invertible in $\mathbb{Z}_p[X]/(X^{256} + 1)$. \square

4 Polynomial Multiplication Implementation

We now describe in more detail the computational advantage of having the modulus Φ_m split into as many factors as possible and present our experimental results. We focus on the case where m is a power of two and write $n = \phi(m) = m/2$. In this case one can use the standard radix-2 FFT-trick to speed up the multiplication. Note that for other m , one can also exploit the splitting in a divide-and-conquer fashion similar to the radix-2 FFT.

Suppose that \mathbb{Z}_p contains a fourth root of unity r so that we can write

$$X^n + 1 = (X^{n/2} + r)(X^{n/2} - r).$$

Then, in algebraic language, the FFT (or NTT) is based on the Chinese remainder theorem, which says that $R_{m,p} = \mathbb{Z}_p[X]/(X^n + 1)$ is isomorphic to the direct product of $\mathbb{Z}_p[X]/(X^{n/2} + r)$ and $\mathbb{Z}_p[X]/(X^{n/2} - r)$. To multiply two polynomials in $R_{m,p}$ one can first reduce them modulo the two factors of the modulus, then multiply the resulting polynomials in the smaller rings, and finally invert the Chinese remainder map in order to obtain the product of the original polynomials. This is called the (radix-2) FFT-trick (see [5] for a very good survey). Note that reducing a polynomial of degree less than n modulo the two sparse polynomials $X^{n/2} \pm r$ is very easy and takes only $\frac{n}{2}$ multiplications, $\frac{n}{2}$ additions and $\frac{n}{2}$ subtractions. If \mathbb{Z}_p contains higher roots so that $X^n + 1$ splits further, then one can apply the FFT-trick recursively to the smaller rings. What is usually referred to as the number theoretic transform (NTT) is the case where \mathbb{Z}_p contains a $2n$ -th root of unity so that $X^n + 1$ splits completely into linear factors. This reduces multiplication in $R_{m,p}$ to just multiplication in \mathbb{Z}_p .

As we are interested in the case where the modulus does not split completely, we need to be able to multiply in rings of the form $\mathbb{Z}_p[X]/(X^{n/k} - r_j)$ with $k < n$. As is common in cryptographic applications (see, for example [6]), we will use the Karatsuba multiplication algorithm to perform this operation. For both the FFT and the Karatsuba multiplication, we have written a relatively straight-forward C implementation.

In Table 1 we give the measurements of our experiments. We have performed multiplications in $R_{512,p} = \mathbb{Z}_p[X]/(X^{256} + 1)$ for four completely splitting primes between 2^{20} and 2^{30} . For each

Number of FFT levels	Primes			
	$2^{20} - 2^{14} + 1$	$2^{23} - 2^{13} + 1$	$2^{25} - 2^{12} + 1$	$2^{27} - 2^{11} + 1$
0	28245	31574	33642	35397
1	27168	29343	31419	32613
2	20989	23158	24915	25677
3	20521	22038	23582	23757
4	22543	23695	25016	24628
5	24473	24715	25337	30366
6	13578	13572	14307	13543
7	13981	14020	14522	13986
8	3873	3844	3847	3857

Table 2. CPU cycles of our FFT-accelerated multiplication algorithm for $\mathbb{Z}_p[X]/(X^{256} + 1)$ using FLINT for base case multiplication. The FFT implementation is a highly optimized AVX2-based implementation.

prime we have used between 0 and 8 levels of FFT before switching to Karatsuba multiplication. 0 levels of FFT means that no FFT stage was used at all and the input polynomials were directly multiplied via Karatsuba multiplication. In the other extreme of 8 levels of FFT, no Karatsuba multiplication was used and the corresponding measurements reflect the speed of our full number theoretic transform down to linear factors with pointwise multiplication as the base case. As one more example, when performing 3 levels of FFT, we were multiplying 8 polynomials each of degree less than 32 via Karatsuba multiplication. The listed numbers are numbers of CPU cycles needed for the whole multiplication. They are the medians of 10000 multiplications each. The tests were performed on a laptop equipped with an Intel Skylake i7 CPU running at 3.4 GHz. The cycle counter in this CPU ticks at a constant rate of 2.6GHz. As one can see, being able to use a prime p so that $X^n + 1$ splits into more than two factors is clearly advantageous. For instance, by allowing $X^n + 1$ to split into 8 factors compared to just 2, we achieve a speedup of about a factor of two.

We have also experimented with highly-optimized polynomial multiplication algorithms provided by a popular computer algebra library FLINT [12] and PARI [24]. FLINT employs various forms of Kronecker substitution for the task of polynomial multiplication. For these experiments we used a fast vectorized FFT implementation written in assembler language with AVX2 instructions. For completeness, Table 2 gives the measurements for the tests with FLINT. Unfortunately, each call of the FLINT multiplication function produces additional overhead costs such as deciding on one of several algorithms and computing complex roots for the FFT used in Kronecker substitution. These additional costs are highly significant for our small polynomials. So for every additional stage of our FFT, one needs to multiply twice as many polynomials with FLINT, and hence FLINT spends twice as much time on these auxiliary tasks that one would not have in an actual cryptographic implementation specialized to a particular prime and modulus. This is especially inefficient when the number of FFT levels is large. There nearly all of the time is spent on these tasks as one can see in Table 2 by comparing the cycle counts of 7 and 8 stages of FFT. Note that for 7 stages of FFT, FLINT is used for the trivial task of multiplying polynomials of degree one.

While we were not able to do a meaningful analysis for the combination of our highly-optimized FFT with FLINT, one can see that at level 0 (where the amount of overhead it does is the lowest), FLINT outperforms our un-optimized Karatsuba multiplication by a factor between 4 and 5, while looking at Level 8 shows that our AVX-optimized FFT outperforms the non-optimized version by approximately the same margin. It is then reasonable to assume that one can improve non-FFT multiplication by approximately the same factor as we improved the FFT multiplication, and

therefore the improvement going from level 1 and 3 would still be approximately a factor 2 in a routine where both Karatsuba and FFT multiplication were highly optimized.

References

1. Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - A new hope. In *USENIX*, pages 327–343, 2016.
2. Carsten Baum, Ivan Damgård, Sabine Oechsner, and Chris Peikert. Efficient commitments and zero-knowledge protocols from ring-sis with applications to lattice-based threshold cryptosystems. *IACR Cryptology ePrint Archive*, 2016:997, 2016.
3. Carsten Baum and Vadim Lyubashevsky. Simple amortized proofs of shortness for linear relations over polynomial rings. *IACR Cryptology ePrint Archive*, 2017:759, 2017.
4. Fabrice Benhamouda, Stephan Krenn, Vadim Lyubashevsky, and Krzysztof Pietrzak. Efficient zero-knowledge proofs for commitments from learning with errors over rings. In *ESORICS*, pages 305–325, 2015.
5. Daniel J. Bernstein. Multidigit multiplication for mathematicians, 2001.
6. Daniel J. Bernstein, Chitchanok Chuengsatiansup, Tanja Lange, and Christine van Vredendaal. NTRU prime. *IACR Cryptology ePrint Archive*, 2016.
7. Joppe W. Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, and Damien Stehlé. CRYSTALS - kyber: a cca-secure module-lattice-based KEM. *IACR Cryptology ePrint Archive*, 2017:634, 2017.
8. H. Cohen. *A Course in Computational Algebraic Number Theory*. Graduate Texts in Mathematics. Springer Berlin Heidelberg, 2000.
9. Rafaël Del Pino, Vadim Lyubashevsky, Gregory Neven, and Gregor Seiler. Practical quantum-safe voting from lattices. In *CCS*, 2017.
10. Léo Ducas, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS - dilithium: Digital signatures from module lattices. *IACR Cryptology ePrint Archive*, 2017:633, 2017.
11. Tim Güneysu, Vadim Lyubashevsky, and Thomas Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In *CHES*, pages 530–547, 2012.
12. W. Hart, F. Johansson, and S. Pancratz. FLINT: Fast Library for Number Theory, 2013. Version 2.4.0, <http://flintlib.org>.
13. Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptography*, 75(3):565–599, 2015.
14. Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT*, pages 598–616, 2009.
15. Vadim Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, pages 738–755, 2012.
16. Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP (2)*, pages 144–155, 2006.
17. Vadim Lyubashevsky and Gregory Neven. One-shot verifiable encryption from lattices. In *EUROCRYPT*, pages 293–323, 2017.
18. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT*, pages 1–23, 2010.
19. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-lwe cryptography. In *EUROCRYPT*, pages 35–54, 2013.
20. Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007.
21. Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, pages 145–166, 2006.
22. Thomas Pöppelmann and Tim Güneysu. Towards practical lattice-based public-key encryption on reconfigurable hardware. In *SAC*, pages 68–85, 2013.
23. Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *EUROCRYPT*, pages 27–47, 2011.
24. The PARI Group, Univ. Bordeaux. *PARI/GP version 2.9.0*, 2016.