EXTENDED ABSTRACT APPEARS IN THE PROCEEDINGS OF SCN 2012.

Field Switching in BGV-Style Homomorphic Encryption

Craig Gentry Shai Halevi Chris Peikert Nigel P. Smart

January 9, 2013

Abstract

The security of contemporary homomorphic encryption schemes over cyclotomic number field relies on fields of very large dimension. This large dimension is needed because of the large modulus-to-noise ratio in the key-switching matrices that are used for the top few levels of the evaluated circuit. However, larger noise (and hence smaller modulus-to-noise ratio) is used in lower levels of the circuit, so from a security standpoint it is permissible to switch to lower-dimension fields, thus speeding up the homomorphic operations for the lower levels of the circuit. However, implementing such field-switching is nontrivial, since these schemes rely on the field algebraic structure for their homomorphic properties.

A basic field-switching operation was used by Brakerski, Gentry and Vaikuntanathan, over number fields of the form $\mathbb{Z}[X]/(X^{2^n}+1)$, in the context of bootstrapping. In this work we generalize and extend this technique to work over any cyclotomic number field, and show how it can be used not only for bootstrapping but also during the computation itself (in conjunction with the "packed ciphertext" techniques of Gentry, Halevi and Smart).

1 Introduction

The last few years have seen a rapid advance in the state of fully homomorphic encryption, yet despite these advances, the existing schemes are still too expensive for many practical purposes. In this paper we make another step forward in making such schemes more efficient. In particular, we present a technique for reducing the dimension of the ciphertexts involved in the homomorphic computation of the lower levels of a circuit. Our techniques apply to homomorphic encryption schemes over number fields, such as the schemes of Brakerski et al. [4, 5, 3], as well as the variants due to López-Alt et al. [14] and Brakerski [2].

The most efficient variants of these schemes work over number fields of the form $\mathbb{Q}(\zeta) \cong \mathbb{Q}[X]/F(X)$, and in all of them the field dimension n, which is the degree of $F(X)$, must be set large enough to ensure security: to support homomorphic evaluation of depth-L circuits with security parameter λ, the schemes require $n = \tilde{\Omega}(L \cdot \mathrm{polylog}(\lambda))$, even under the strongest plausible hardness assumptions for their underlying computational problems (e.g., ring-LWE [15]).[1] In practice, the field dimension for moderately deep circuits can easily be many thousands. For example, to be able to evaluate AES homomorphically, Gentry et al. [13] used circuits of depth $L \geq 50$, with a corresponding field dimension of over $50{,}000$.

As homomorphic operations are performed, the ratio of noise to modulus in the ciphertexts grows. Consequently, it becomes permissible to use lower-dimension fields, which can speed up further homomorphic

[1] The schemes from [3, 2] can also obtain security by using high-dimensional vectors over low-dimensional number fields. But their most efficient variants use low-dimensional vectors over high-dimensional fields, since the runtime of certain operations is cubic in the dimension of the vectors.

computations. However, since we must start with ciphertexts from a high-dimensional field, we need a method for transforming them into small-field ciphertexts that encrypt the same (or related) messages. Such a "field switching" procedure was described by Brakerski et al. [3], in the context of reducing the ciphertext size prior to bootstrapping. The procedure in [3], however, is specific to number fields of the form $K_{2^k} = \mathbb{Q}[X]/(X^{2^{k-1}}+1)$, i.e., cyclotomic number fields with power-of-2 index. Moreover, by itself it cannot be combined with the "packed evaluation" techniques from [18, 11]. (These techniques use Chinese-remainder encoding to "pack" many plaintext values into each ciphertext, and then each homomorphic operation is applied to all these values at once. For our purposes, we must consider the effect of the field-switching operation on all these plaintext values.) Extending and improving the field switching procedure is the goal of our work.

1.1 Our Contribution

We present a general field-switching transformation that can be applied to any *cyclotomic* number field $K = \mathbb{Q}(\zeta_m) \cong \mathbb{Q}[X]/\Phi_m(X)$ for arbitrary m (where $\Phi_m(X) \in \mathbb{Z}[X]$ is the mth cyclotomic polynomial), and works well in conjunction with packed ciphertexts. For any divisor m' of m, our procedure takes as input a "big-field ciphertext" c over K that encrypts many plaintext values, and outputs a "small-field ciphertext" c over $K' = \mathbb{Q}(\zeta_{m'}) \cong \mathbb{Q}[X]/\Phi_{m'}(X) \subseteq K$ that encrypts a certain subset of the input plaintext values.[2]

Our transformation relies heavily on the algebraic properties of the cyclotomic number fields K, K' and their respective rings of (algebraic) integers R, R'. In particular, we use the interpretation of K as an extension field of K', and relationships between their various embeddings into the complex numbers $\mathbb{C}$; the factorization of integer primes in R and R'; and the *trace function* $\mathrm{Tr}_{K/K'}$ that maps elements in K to the subfield K'. With these tools in hand, the transformation itself is quite simple, and consists of the following three steps:

1. We first apply a key-switching operation to obtain a big-field ciphertext over K with respect to a small-field secret key $s' \in K' \subset K$. Proving the security of this operation relies on a novel way of embedding the ring-LWE problem over K' into K, which may be of independent interest.

2. Next, we multiply the resulting ciphertext by a certain element of the ring $R \subset K$, which depends only on the subset (or other function) of the plaintext values that we want to include in the output ciphertext.

3. Finally, we take the trace of all the K-elements in the ciphertext, thus obtaining an output ciphertext over the subfield K', which decrypts under the secret key $s' \in K'$ to the desired plaintext values.

We note that in addition to being simpler and more general than the transformation from [3], our transformation is also more efficient even when applied in the special case of K_{2^k}: when switching from K_{2^k} to $K_{2^{k'}}$, the transformation from [3] includes a step where the size of the ciphertext (and hence the time that it takes to perform operations) is expanded by a factor of $2^{k-k'}$. Our transformation does not need that extra step, hence saving this extra factor in performance.

In Section 2 below we recall the algebraic concepts needed for our transformation, and then the transformation itself it described in Section 3.

[2] More generally, the output ciphertext can even encrypt certain linear functions of the input plaintext values.

2 Preliminaries

For any positive integer u we let $[u] = \{0, \ldots, u-1\}$. Throughout this work, for a coset $z \in \mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$ we let $[z]_q \in \mathbb{Z}$ denote its canonical representative in $\mathbb{Z} \cap [-q/2, q/2)$. One can also view $[\cdot]_q$ as the operation that takes an arbitrary integer z and reduces it modulo q into the interval $[-q/2, q/2)$.

2.1 Algebraic Background

Recall that an *ideal* I in a commutative ring R is a nontrivial (i.e., $I \neq \emptyset$ and $I \neq \{0\}$) additive subgroup which is closed under multiplication by R. For ideals I, J, their sum is the ideal $I + J = \{a + b : a \in I, b \in J\}$, and their product IJ is the ideal consisting of all *sums* of terms ab for $a \in I, b \in J$. An R-ideal $\mathfrak{p}$ is prime if $ab \in \mathfrak{p}$ (for some $a, b \in R$) implies $a \in \mathfrak{p}$ or $b \in \mathfrak{p}$ (or both). All the rings we work with have unique factorization of ideals into powers of prime ideals, and a Chinese Remainder Theorem. A *fractional ideal* is, informally, an ideal with a denominator. Formally, letting K be the field of fractions of R, a fractional ideal of R is a subset $I \subseteq K$ for which there exists a denominator $d \in R$ such that $dI \subseteq R$ is an ideal in R. For an R-ideal I, the quotient ring $R_I = R/I$ consists of the residue classes $a + I$ for all $a \in R$, with the ring operations induced by R. More generally, for a (possibly fractional) ideal I and an ideal $J \subseteq R$, the quotient $I_J = I/IJ$ is an additive group, and an R-module, with addition and multiplication operations induced by R. We often write $a \bmod I$ instead of $a + I$ to denote the residue classes $a + I$, and we write $a = b \pmod{I}$ to denote that a, b belong to the same residue class, i.e., $a + I = b + I$.

For computational purposes, all of the rings and fields we work with have efficient representations of their elements, and efficient (i.e., polynomial time in the bit length of the arguments) algorithms for all the operations we use. For quotients A/B, cosets are represented using a fixed set of distinguished representatives. In this work we largely ignore the details of concrete representations and algorithms, and refer to [16] for fast, specialized algorithms for working with the cyclotomic fields and rings that we use in this work.

2.1.1 Cyclotomic Fields and Rings

For a positive integer m, let $K = \mathbb{Q}(\zeta_m)$ be the mth *cyclotomic number field*, where ζ_m is an abstract element of order m. (In particular, we do not view ζ_m as any particular root of unity in $\mathbb{C}$.) The minimal polynomial of ζ_m is the mth *cyclotomic polynomial* $\Phi_m(X) = \prod_{i \in \mathbb{Z}_m^*}(X - \omega_m^i) \in \mathbb{Z}[X]$, where $\omega_m = \exp(2\pi\sqrt{-1}/m) \in \mathbb{C}$ is the principal mth complex root of unity, and the roots $\omega_m^i \in \mathbb{C}$ range over all the *primitive* complex mth roots of unity. Therefore, K is a field extension of degree $n = \varphi(m)$ over $\mathbb{Q}$, and is isomorphic to the polynomial ring $\mathbb{Q}[X]/\Phi_m(X)$ by identifying ζ_m with X (there are other representations as well). The *ring of (algebraic) integers* in K, called the mth cyclotomic ring, is $R = \mathbb{Z}[\zeta_m] \cong \mathbb{Z}[X]/\Phi_m(X)$.

The field extension $K/\mathbb{Q}$ has n automorphisms $\tau_i \colon K \to K$ that fix $\mathbb{Q}$ pointwise, which are characterized by $\tau_i(\zeta_m) = \zeta_m^i$ for $i \in \mathbb{Z}_m^*$. (Equivalently, $\tau_i(a(X)) = a(X^i) \bmod \Phi_m(X)$ when viewing K as $\mathbb{Q}[X]/\Phi_m(X)$.) Because $K/\mathbb{Q}$ is Galois (i.e., the number of automorphisms equals the dimension of the extension), the $\mathbb{Q}$-linear[3] *(field) trace* $\mathrm{Tr}_{K/\mathbb{Q}} \colon K \to \mathbb{Q}$ can be defined as the sum of the automorphisms: $\mathrm{Tr}_{K/\mathbb{Q}}(a) = \sum_{i \in \mathbb{Z}_m^*} \tau_i(a) \in \mathbb{Q}$. (See below for another formulation.)

Similarly to the automorphisms τ_i (which map K to itself), there are n concrete ways of viewing K as a subfield of the complex numbers $\mathbb{C}$. Namely, there are n injective ring homomorphisms from K to $\mathbb{C}$ that fix $\mathbb{Q}$ pointwise, called *embeddings*, which are denoted $\sigma_i \colon K \to \mathbb{C}$ for $i \in \mathbb{Z}_m^*$ and characterized by $\sigma_i(\zeta_m) = \omega_m^i$. The embeddings may be seen as the compositions of the abstract automorphisms τ_i with the

[3] A function f is D-linear if $f(a + b) = f(a) + f(b)$ and $f(d \cdot a) = d \cdot f(a)$ for all $d \in D$ and all a, b.

complex embedding that identifies $\zeta_m \in K$ with $\omega_m \in \mathbb{C}$. Therefore, the field trace can also be written as the sum of the embeddings, as $\mathrm{Tr}_{K/\mathbb{Q}}(a) = \sum_{i \in \mathbb{Z}_m^*} \sigma_i(a) \in \mathbb{Q}$. The *canonical embedding* $\sigma \colon K \to \mathbb{C}^n$ is the concatenation of all the complex embeddings, i.e., $\sigma(a) = (\sigma_i(a))_{i \in \mathbb{Z}_m^*}$, and it endows K with a canonical geometry. In particular, define the Euclidean (ℓ_2) and ℓ_∞ norms on K as

$$\|a\| := \|\sigma(a)\| = \sqrt{\sum_i |\sigma_i(a)|^2} \quad \text{and} \quad \|a\|_\infty := \|\sigma(a)\|_\infty = \max_i |\sigma_i(a)|,$$

respectively. Note that $\|a \cdot b\| \leq \|a\|_\infty \cdot \|b\|$ and $\|a \cdot b\|_\infty \leq \|a\|_\infty \cdot \|b\|_\infty$ for any $a, b \in K$, because the σ_i are ring homomorphisms.

2.1.2 Towers of Cyclotomics

For any positive integer m' dividing m, let $K' = \mathbb{Q}(\zeta_{m'})$ and $R' = \mathbb{Z}[\zeta_{m'}]$ be the m'th cyclotomic field and ring (of dimension $n' = \varphi(m')$ over $\mathbb{Q}$ and $\mathbb{Z}$), respectively. As above, the field extension $K'/\mathbb{Q}$ has $n' = \varphi(m')$ automorphisms $\tau'_{i'} \colon K' \to K'$ and n' complex embeddings $\sigma'_{i'} \colon K' \to \mathbb{C}$ (for $i' \in \mathbb{Z}_{m'}^*$), the latter of which define the canonical embedding $\sigma' \colon K' \to \mathbb{C}^{n'}$.

We will use extensively the fact that K is a field extension of K', and R is a ring extension of R', both of dimension n/n' (because $K/\mathbb{Q}$ and $K'/\mathbb{Q}$ have dimensions n and n', respectively). That is, K' and R' may respectively be seen as a subfield of $K = K'(\zeta_m)$ and a subring of $R = R'[\zeta_m]$, under the ring embedding that identifies $\zeta_{m'}$ with $\zeta_m^{m/m'}$. Moreover, the field extension K/K' is Galois, i.e., it has n/n' automorphisms that fix K' pointwise, which are precisely those τ_i for which $i = 1 \pmod{m'}$. This follows from the fact that

$$\tau_i(\zeta_{m'}) = \tau_i(\zeta_m^{m/m'}) = \zeta_m^{(m/m')i \bmod m} = \zeta_{m'}^{i \bmod m'}, \tag{2.1}$$

and that reducing modulo m' induces an (n/n')-to-1 mapping from $\mathbb{Z}_m^*$ to $\mathbb{Z}_{m'}^*$. The K'-linear (intermediate) trace function $\mathrm{Tr}_{K/K'} \colon K \to K'$ may be defined as the sum of these automorphisms:

$$\mathrm{Tr}_{K/K'}(a) = \sum_{i = 1 \pmod{m'}} \tau_i(a).$$

A standard fact from field theory is that the intermediate trace satisfies $\mathrm{Tr}_{K/\mathbb{Q}} = \mathrm{Tr}_{K'/\mathbb{Q}} \circ \mathrm{Tr}_{K/K'}$. Another standard fact is that $\mathrm{Tr}_{K/K'}$ is a "universal" K'-linear function, in that such function $L \colon K \to K'$ can be expressed as $L(a) = \mathrm{Tr}_{K/K'}(d \cdot a)$ for some fixed $d \in K$.

Similarly to Equation (2.1), for any $i \in \mathbb{Z}_m^*$ the embedding σ_i coincides with $\sigma'_{i \bmod m'}$ on the subfield K'. Using this fact we get the following relation between the intermediate trace and the complex embeddings of K and K'.

Lemma 2.1. *For any $a \in K$ and $i' \in \mathbb{Z}_{m'}^*$,*

$$\sigma'_{i'}(\mathrm{Tr}_{K/K'}(a)) = \sum_{i = i' \pmod{m'}} \sigma_i(a).$$

In matrix form, $\sigma'(\mathrm{Tr}_{K/K'}(a)) = P \cdot \sigma(a)$, where P is the $\varphi(m')$-by-$\varphi(m)$ matrix (with rows indexed by $i' \in \mathbb{Z}_{m'}^*$ and columns by $i \in \mathbb{Z}_m^*$) whose (i', i)th entry is 1 if $i = i' \pmod{m'}$, and is 0 otherwise.

Proof. Fix an arbitrary $k \in \mathbb{Z}_m^*$ such that $k = i' \pmod{m'}$. Then because $\sigma'_{i'}$ coincides with σ_k on K', and by definition of $\mathrm{Tr}_{K/K'}$ and linearity of σ_k, we have

$$\sigma'_{i'}(\mathrm{Tr}_{K/K'}(a)) = \sigma_k\bigg(\sum_{j=1 \pmod{m'}} \tau_j(a)\bigg) = \sum_{j=1 \pmod{m'}} \sigma_k(\tau_j(a)) = \sum_{i=i' \pmod{m'}} \sigma_i(a),$$

where for the last equality we have used $\sigma_k \circ \tau_j = \sigma_{k \cdot j}$ and $k \in \mathbb{Z}_m^*$, so $i = k \cdot j \in \mathbb{Z}_m^*$ runs over all indices congruent to i' modulo m' when $j \in \mathbb{Z}_m^*$ runs over all indexes congruent to 1 modulo m'. □

An immediate corollary is that the intermediate trace maps short elements of K to short elements of K'.

Corollary 2.2. *For any $a \in K$, we have $\|\mathrm{Tr}_{K/K'}(a)\| \leq \|a\| \cdot \sqrt{n/n'}$.*

Proof. By Lemma 2.1, we have $\sigma'(\mathrm{Tr}_{K/K'}(a)) = P \cdot \sigma(a)$. The rows of P are orthogonal (since each column of P has exactly one nonzero entry), and each has Euclidean norm exactly $\sqrt{n/n'}$. □

2.1.3 Prime Splitting and Plaintext Arithmetic

We now describe the factorization ("splitting") of prime integers in cyclotomic rings, how it allows for encoding and operating on several finite-field elements, and the particular functions induced by the (intermediate) trace function $\mathrm{Tr}_{K/K'}$. Further details and proofs can be found in many texts on algebraic number theory, e.g., [19].

Prime splitting. Let $p \in \mathbb{Z}$ be a prime integer. In the mth cyclotomic ring $R = \mathbb{Z}[\zeta_m]$ (which has degree $n = \varphi(m)$ over $\mathbb{Z}$), pR is often not a prime ideal, but instead factors into prime ideals as follows. First divide out all the factors of p from m, writing $m = \bar{m} \cdot p^k$ where $p \nmid \bar{m}$. Let $e = \varphi(p^k)$, and let d be the multiplicative order of p modulo $\bar{m}$ (i.e., in $\mathbb{Z}_{\bar{m}}^*$); note that d divides $\varphi(\bar{m}) = n/e$. (The values d, e are respectively called the *inertial degree* and *ramification index* of p in R.) Let $G = \mathbb{Z}_{\bar{m}}^*/\langle p \rangle$, the (multiplicative) quotient group $\mathbb{Z}_{\bar{m}}^*$ modulo the order-d subgroup generated by p, so G has order $f = \varphi(\bar{m})/d = n/(de)$. The ideal pR then factors as

$$pR = \prod_{i \in G} \mathfrak{p}_i^e,$$

where the $\mathfrak{p}_i$ are distinct prime ideals in R, all having norm $|R/\mathfrak{p}_i| = p^d$. These are called the prime ideals *lying over* p in R. Each quotient ring $R/\mathfrak{p}_i$ is therefore isomorphic to the finite field $\mathbb{F}_{p^d}$. (In fact there are exactly d isomorphisms between them, because $\mathbb{F}_{p^d}$ has d automorphisms.)

One useful explicit representation of the prime ideals $\mathfrak{p}_i$, and the isomorphisms between $R/\mathfrak{p}_i$ and (some canonical representation of) $\mathbb{F}_{p^d}$, is as follows.[4] Modulo p (i.e., in $\mathbb{F}_p[X]$), the $\bar{m}$th cyclotomic polynomial factors into f distinct irreducible degree-d polynomials, as $\Phi_{\bar{m}}(X) = \prod_{i \in G} F_i(X) \pmod{p}$, and the mth cyclotomic polynomial factors as

$$\Phi_m(X) = \Phi_{\bar{m}}(X)^e = \prod_{i \in G} F_i(X)^e \pmod{p}.$$

Then each $\mathfrak{p}_i$ is generated by p and $F_i(\zeta_m)$, i.e., $\mathfrak{p}_i = \langle p, F_i(\zeta_m) \rangle = pR + F_i(\zeta_m)R$.

[4] Other (sometimes more computationally efficient) representations are possible, e.g., by expressing $\mathbb{Z}[\zeta_m]$ as a multi-variate polynomial ring; see [16].

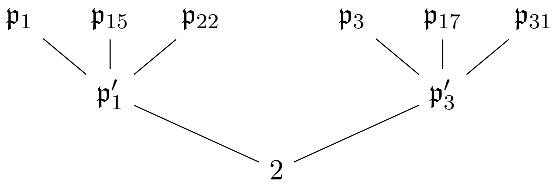

Figure 1: Factorization of $2 \in \mathbb{Z}$ into distinct prime ideals $\mathfrak{p}'_{i'}$ in $R' = \mathbb{Z}[\zeta_7]$, and $\mathfrak{p}_i$ in $R = \mathbb{Z}[\zeta_{91}]$. The displayed subscripts indicate a choice of representatives from the cosets of the multiplicative subgroups $\langle 2 \rangle \subseteq \mathbb{Z}_7^*$ and $\langle 2 \rangle \subseteq \mathbb{Z}_{91}^*$, which have orders $d' = 3$ and $d = 12$, respectively.

Let $\omega_{\bar{m}}$ denote some arbitrary element of order $\bar{m}$ in $\mathbb{F}_{p^d}$. (Such an element exists because the multiplicative group $\mathbb{F}_{p^d}^*$ is cyclic and has order $p^d - 1 = 0 \pmod{\bar{m}}$.) Then for any $i \in G$, the polynomial $F_i(X)$ splits into linear factors over the extension field $\mathbb{F}_{p^d}$ as

$$F_i(X) = \prod_{j \in i\langle p \rangle \subseteq \mathbb{Z}_{\bar{m}}^*} (X - \omega_{\bar{m}}^j) = \prod_{k=0}^{d-1} (X - \omega_{\bar{m}}^{i \cdot p^k}).$$

Therefore, fixing for each $i \in G$ some arbitrary representative $u_i \in i\langle p \rangle \subseteq \mathbb{Z}_{\bar{m}}^*$, we get a concrete isomorphism between the quotient $R/\mathfrak{p}_i$ and $\mathbb{F}_{p^d}$, which is characterized by $\zeta_m \mapsto \omega_{\bar{m}}^{u_i}$. (Note that ζ_m has order $\bar{m}$ modulo p.) Looking ahead, these isomorphisms will be used to define several "plaintext slots" in a homomorphic cryptosystem, i.e., an encoding of f plaintext elements of $\mathbb{F}_{p^d}$ in a single element of $R/2R$, which is the plaintext space of the cryptosystem.

Splitting in cyclotomic towers. Of course, the above derivation also applies to the ideals that lie over p in $R' = \mathbb{Z}[\zeta_{m'}] \subseteq R$. For each such ideal $\mathfrak{p}'$, we next describe the factorization of $\mathfrak{p}'R$ into prime ideals in R. These are the prime ideals that lie over $\mathfrak{p}'$ in R, and since "lying over" is an associative property, they also lie over p (as illustrated in Figure 1).

Let $\bar{m}, d, e, f, G$ and the prime ideals $\mathfrak{p}_i$ for $i \in G$ be as above for R, and define $\bar{m}', d', e', f', G' = \mathbb{Z}_{\bar{m}'}^*/\langle p \rangle$ and prime ideals $\mathfrak{p}'_{i'}$ for $i' \in G'$ similarly for R'. Note that $d'|d$, $e'|e$, and $f'|f$, and that the natural homomorphism $g \colon G \to G'$ defined as $g(i) = i \bmod \bar{m}$ is surjective and (f/f')-to-1. Then for every $i' \in G'$, the factorization of $\mathfrak{p}'_{i'}R$ is

$$\mathfrak{p}'_{i'}R = \prod_{i \in g^{-1}(i')} \mathfrak{p}_i^{e/e'} = \prod_{i = i' \pmod{\bar{m}}} \mathfrak{p}_i^{e/e'}.$$

Therefore, there are f/f' prime ideals of R lying over each $\mathfrak{p}'_{i'}$, and taken over all $i' \in G'$ they partition the prime ideals of R lying over p.

Plaintext encoding. Let $\mathbb{F} = \mathbb{F}_{p^d}$ and $\mathbb{F}' = \mathbb{F}_{p^{d'}} \subseteq \mathbb{F}$. By the above and the Chinese remainder theorem, the natural ring homomorphisms yield the following (where $\cong$ denotes a ring isomorphism, and multiplication $\odot$ in $\mathbb{F}'^{f'}$ and $\mathbb{F}^f$ is coordinate-wise):

$$R'/pR' \longrightarrow R'/\big(\prod_{i' \in G'} \mathfrak{p}'_{i'}\big) \cong \bigoplus_{i' \in G'} R'/\mathfrak{p}'_{i'} \cong \mathbb{F}'^{f'}$$

$$R/pR \longrightarrow R/\big(\prod_{i \in G} \mathfrak{p}_i\big) = R/\big(\prod_{i' \in G'} \prod_{i \in g^{-1}(i')} \mathfrak{p}_i\big) \cong \bigoplus_{i' \in G'} \bigoplus_{i \in g^{-1}(i')} R/\mathfrak{p}_i \cong (\mathbb{F}^{f/f'})^{f'} = \mathbb{F}^f.$$

(Note that the first homomorphism in each line is surjective, but not necessarily an isomorphism, due to possible ramification.) Following [18, 3, 11, 12, 13], in the context of homomorphic encryption the above morphisms allow for encoding a vector of f' individual elements of $\mathbb{F}'$ (respectively, f elements of $\mathbb{F}$) into the plaintext ring $R'_p = R/pR'$ (resp., $R_p = R/pR$), so that a single homomorphic addition and multiplication acts component-wise on the underlying vectors of field elements.

Trace operations. As mentioned in the introduction, our field-switching technique is built around applying the trace function $\mathrm{Tr}_{K/K'}$ to the elements of a big-field ciphertext, thus obtaining a related small-field ciphertext. Since we use "packed" ciphertexts that encrypt arrays of elements in $\mathbb{F}$ via the above isomorphisms, we need to understand the effect of the trace function on those $\mathbb{F}$-elements.

The remainder of this subsection is therefore devoted to characterizing the functions $(\mathbb{F}^{f/f'})^{f'} \to \mathbb{F}'^{f'}$ that can be induced by $\mathrm{Tr}_{K/K'}$. More specifically, we determine exactly which functions $L\colon R/(\prod_{i \in G} \mathfrak{p}_i) \to R'/(\prod_{i' \in G'} \mathfrak{p}'_{i'})$ can be expressed as $L(a) = \mathrm{Tr}_{K/K'}(d \cdot a)$ for some fixed $d \in K$. It turns out that by fixing an appropriate choice of isomorphisms between the quotient rings and finite fields above, we can obtain the *concatenation* of any f' individual $\mathbb{F}'$-linear functions $\mathbb{F}^{f/f'} \to \mathbb{F}'$ (see Corollary 2.5 for a precise statement).[5]

As already noted, the isomorphisms between the quotient rings and finite fields are not necessarily unique; they are determined by the choice of representatives $u_{i'} \in i'\langle p \rangle \subseteq \mathbb{Z}^*_{\bar{m}'}$ and $u_i \in i\langle p \rangle \subseteq \mathbb{Z}^*_{\bar{m}}$, and roots of unity $\omega_{\bar{m}'} \in \mathbb{F}'$ and $\omega_{\bar{m}} \in \mathbb{F}$. For our purposes below, it is important to choose these in a "consistent" fashion, as follows. Given $\omega_{\bar{m}}$, let $\omega_{\bar{m}'} = \omega_{\bar{m}}^{\bar{m}/\bar{m}'} \in \mathbb{F}'$. (Note that all $\varphi(\bar{m}')$ elements of order $\bar{m}'$ in $\mathbb{F}$ are indeed in the subfield $\mathbb{F}'$.) Next, let $\ell \geq 0$ be the integer exponent such that $m/m' = (\bar{m}/\bar{m}') \cdot p^{\ell}$. Then given $u_{i'}$ for $i' \in G'$, choose u_i for each $i \in g^{-1}(i')$ so that $p^{\ell} \cdot u_i = u_{i'} \pmod{\bar{m}'}$. (Note that such u_i always exist, by definition of g and the quotient group G.) We denote the isomorphisms obtained from these choices by

$$h'_{i'}\colon R'/\mathfrak{p}'_{i'} \to \mathbb{F}' \text{ (for } i' \in G'\text{)} \quad \text{and} \quad h_i\colon R/\mathfrak{p}_i \to \mathbb{F} \text{ (for } i \in G\text{)}.$$

Also, for each $i' \in G'$ denote the product of prime ideals lying over $\mathfrak{p}'_{i'}$ in R, called the *radical* of $\mathfrak{p}'_{i'}R$, by $\tilde{\mathfrak{p}}_{i'} = \prod_{i \in g^{-1}(i')} \mathfrak{p}_i$, and define the ring isomorphism

$$\tilde{h}_{i'}\colon R/\tilde{\mathfrak{p}}_{i'} \to \mathbb{F}^{f/f'}, \quad \tilde{h}_{i'}(a) = \big(h_i(a \bmod \mathfrak{p}_i)\big)_{i \in g^{-1}(i')}.$$

In Lemma 2.4 below, we show that under the above isomorphisms, the $\mathbb{F}'$-linear functions $\bar{L}\colon \mathbb{F}^{f/f'} \to \mathbb{F}'$ correspond bijectively with the R'-linear functions $L\colon R/\tilde{\mathfrak{p}}_{i'} \to R'/\mathfrak{p}'_{i'}$ (for any $i' \in G'$). Recall that any function of the latter type can be expressed as $L(a) = \mathrm{Tr}_{K/K'}(d \cdot a)$ for some fixed $d \in K$. Conversely, every function L (with domain and range as above) that can be expressed as $L(a) = \mathrm{Tr}_{K/K'}(d \cdot a)$ is clearly R'-linear, so it always induces an $\mathbb{F}'$-linear function. The heart of Lemma 2.4 is the following fact.

Lemma 2.3. *Let $\mathfrak{p}'_{i'}$ for some $i' \in G'$ be a prime ideal lying over p in R', and let $\tilde{\mathfrak{p}}_{i'}$ be the radical of $\mathfrak{p}_{i'}R$. Let $r' \in R' \subseteq R$ be arbitrary, and let $s = h'_{i'}(r' \bmod \mathfrak{p}'_{i'}) \in \mathbb{F}' \subseteq \mathbb{F}$. Then*

$$\tilde{h}_{i'}(r' \bmod \tilde{\mathfrak{p}}_{i'}) = (s, s, \ldots, s) \in \mathbb{F}'^{f/f'},$$

i.e., every entry of $\tilde{h}_{i'}(r' \bmod \tilde{\mathfrak{p}}_{i'})$ is equal to $h'_{i'}(r' \bmod \mathfrak{p}'_{i'})$.

[5] Note that any $\mathbb{F}'$-linear function $L\colon \mathbb{F}^{f/f'} \to \mathbb{F}'$ can always be expressed as $L(\vec{a}) = \mathrm{Tr}_{\mathbb{F}/\mathbb{F}'}(\langle \vec{d}, \vec{a} \rangle)$ for some fixed $\vec{d} \in \mathbb{F}^{f/f'}$, where $\langle \cdot, \cdot \rangle$ is the usual inner product and $\mathrm{Tr}_{\mathbb{F}/\mathbb{F}'}$ denotes the ($\mathbb{F}'$-linear) trace of the field extension $\mathbb{F}/\mathbb{F}'$.

Proof. Recall that under our choice of isomorphisms, $\omega_{\bar{m}'} = \omega_{\bar{m}}^{\bar{m}/\bar{m}'} \in \mathbb{F}'$ is of order $\bar{m}'$, and $p^\ell \cdot u_i = u_{i'} \bmod \bar{m}'$, where $\ell \geq 0$ is the integer satisfying $m/m' = (\bar{m}/\bar{m}') \cdot p^\ell$. Also recall that

$$\tilde{h}_{i'}(r' \bmod \tilde{\mathfrak{p}}_{i'}) = \big(h_i(r' \bmod \mathfrak{p}_i)\big)_{i \in g^{-1}(i')}.$$

For each $i \in g^{-1}(i')$, the entry $h_i(r' \bmod \mathfrak{p}_i)$ is obtained by mapping ζ_m to $\omega_{\bar{m}}^{u_i}$, and hence $\zeta_{m'} = \zeta_m^{m/m'} = \zeta_m^{(\bar{m}/\bar{m}')\cdot p^\ell}$ to

$$\omega_{\bar{m}}^{(\bar{m}/\bar{m}')\cdot p^\ell u_i} = \omega_{\bar{m}'}^{p^\ell u_i} = \omega_{\bar{m}'}^{u_{i'}} \in \mathbb{F}',$$

which is exactly the mapping done by $h'_{i'}$. Since $r' \in R' = \mathbb{Z}[\zeta_{m'}]$, this proves the claim. □

Lemma 2.4. *Let $i' \in G'$ be arbitrary, and let $\mathfrak{p}' = \mathfrak{p}'_{i'}$ and $\tilde{\mathfrak{p}} = \tilde{\mathfrak{p}}_{i'}$. Then under the isomorphisms $h' = h'_{i'}$ and $\tilde{h} = \tilde{h}_{i'}$ defined above, the $\mathbb{F}'$-linear functions $\bar{L}\colon \mathbb{F}^{f/f'} \to \mathbb{F}'$ are in bijective correspondence with the R'-linear functions $L\colon R/\tilde{\mathfrak{p}} \to R'/\mathfrak{p}'$.*

Proof. For any $\mathbb{F}'$-linear function $\bar{L}$, we claim that $L = h'^{-1} \circ \bar{L} \circ \tilde{h}$ is the corresponding R'-linear function. To see this, note that by Lemma 2.3 and the fact that $\tilde{h}$ is a ring homomorphism, for any $r' \in R'$ and $a \in R/\tilde{\mathfrak{p}}$ we have

$$\tilde{h}(r' \cdot a) = \tilde{h}(r' \bmod \tilde{\mathfrak{p}}) \odot \tilde{h}(a) = h'(r' \bmod \mathfrak{p}') \cdot \tilde{h}(a) \in \mathbb{F}^{f/f'}.$$

So, by $\mathbb{F}'$-linearity of $\bar{L}$ and the fact that h' is a ring homomorphism, we have

$$L(r' \cdot a) = h'^{-1}(\bar{L}(\tilde{h}(r' \cdot a))) = h'^{-1}\big(\, h'(r' \bmod \mathfrak{p}') \cdot \bar{L}(\tilde{h}(a)) \,\big) = r' \cdot L(a) \in R'/\mathfrak{p}',$$

as desired. The other direction proceeds essentially identically, with $\bar{L} = h' \circ L \circ \tilde{h}^{-1}$. □

An application of the Chinese remainder theorem with the prime ideals $\tilde{\mathfrak{p}}_{i'}$ in R, combined with Lemma 2.4, immediately yields the following corollary.

Corollary 2.5. *Let $\mathfrak{p}' = \prod_{i' \in G'} \mathfrak{p}'_{i'}$ and $\mathfrak{p} = \prod_{i' \in G'} \tilde{\mathfrak{p}}_i$ be the radicals of pR' and pR, respectively. Then under the isomorphisms $\{h_{i'}\}_{i' \in G'}$ and $\{\tilde{h}_{i'}\}_{i' \in G'}$ defined above, the R'-linear functions $L\colon R/\mathfrak{p} \to R'/\mathfrak{p}'$ are in bijective correspondence with the functions $\bar{L}\colon (\mathbb{F}^{f/f'})^{f'} \to \mathbb{F}'^{f'}$ of the form*

$$\bar{L}\left((\vec{a}_{i'})_{i' \in G'} \right) = \big(\bar{L}_{i'}(\vec{a}_{i'})\big)_{i' \in G'},$$

where every $\bar{L}_{i'}\colon \mathbb{F}^{f/f'} \to \mathbb{F}'$ is $\mathbb{F}'$-linear.

We note that given a function $\bar{L}$ of the form above, we can efficiently find the corresponding R'-linear function L: choosing sufficiently many linearly independent inputs $\vec{a}_i$ and evaluating $\bar{L}(\vec{a}_i)$, we use the above isomorphisms to translate them to the corresponding pairs $(a_i, L(a_i))$. Recalling that K is a vector space of dimension n/n' over K', we then have a system of linear equations $L(a_i) = \mathrm{Tr}_{K/K'}(d \cdot a_i)$, which we can solve to obtain $d \in K$.

2.1.4 Duality

An important and useful object in K is the *dual* of R (also known as the *codifferent* of K), defined as

$$R^\vee = \{a \in K : \mathrm{Tr}_{K/\mathbb{Q}}(aR) \subseteq \mathbb{Z}\} \supseteq R.$$

Because $\mathrm{Tr}_{K/\mathbb{Q}} = \mathrm{Tr}_{K'/\mathbb{Q}} \circ \mathrm{Tr}_{K/K'}$, it is easy to verify that also $R^\vee = \{a \in K : \mathrm{Tr}_{K/K'}(aR) \subseteq R'^\vee\}$. Therefore, we have the convenient equation

$$\mathrm{Tr}_{K/K'}(R^\vee) = R'^\vee. \tag{2.2}$$

Note that by contrast, frequently $\mathrm{Tr}_{K/K'}(R)$ does *not* equal R', but is instead some proper ideal of it.[6] Many other algebraic and geometric advantages of working with $R^\vee$ instead of R are discussed in [15, 16].

The codifferent is a principal fractional ideal, i.e., $R^\vee = t^{-1}R$ for some $t \in R$ (which is not unique). Therefore, division by t induces a bijection from R to $R^\vee$, and from any quotient ring $R_\mathfrak{p} = R/\mathfrak{p}$ to $R^\vee_\mathfrak{p} = R^\vee/\mathfrak{p}R^\vee$. Although the target objects are *not* rings (because $R^\vee \cdot R^\vee \not\subseteq R^\vee$), they are R-modules, and the bijections are R-module isomorphisms.

Of course, we also have $R'^\vee = t'^{-1}R'$ for some $t' \in R'$. By Equation (2.2) and K'-linearity of the trace, for any ideal $\mathfrak{p}$ in R', we have

$$\mathrm{Tr}_{K/K'}(R^\vee_\mathfrak{p}) = \mathrm{Tr}_{K/K'}(R^\vee/\mathfrak{p}R^\vee) = R'^\vee/\mathfrak{p}R'^\vee = R'^\vee_\mathfrak{p}.$$

In the previous subsection we considered R'-linear functions $L\colon R \to R'$ (or actually, their induced functions $R_p \to R'_p$), which can always be expressed as $L(a) = \mathrm{Tr}_{K/K'}(d \cdot a)$ for some fixed $d \in K$. Typically, d is *not* in R (because $\mathrm{Tr}_{K/K'}(R) \neq R'$), but it is easy to see that $d \in t'R^\vee$ always (because if not, then $\mathrm{Tr}_{K/K'}(dR) \not\subseteq t'R'^\vee = R'$). For the purposes of our field-switching procedure, it will be more convenient to instead work with corresponding R'-linear functions from $R^\vee$ to $R'^\vee$, which can be represented via the trace by elements in R. Namely, for an R'-linear function $L : R \to R'$, where $L(a) = \mathrm{Tr}_{K/K'}(d^\vee \cdot a)$ for some $d^\vee \in t'R^\vee$, we will consider the corresponding function

$$L^\vee\colon R^\vee \to R'^\vee, \quad L^\vee(a^\vee) = L(t \cdot a^\vee)/t' = \mathrm{Tr}_{K/K'}((t/t')d^\vee \cdot a^\vee) = \mathrm{Tr}_{K/K'}(d \cdot a^\vee),$$

which is represented by $d = (t/t')d^\vee \in R$.

Following [16], we extend the operation $[\cdot]_q$ to $R^\vee_p$ by fixing a particular $\mathbb{Z}$-basis of $R^\vee$ (and $\mathbb{Z}_q$-basis of $R^\vee_q$), called the *decoding basis*, and representing the argument as a $\mathbb{Z}_q$-combination of the basis vectors and applying the $[\cdot]_q$ operation to each of its coefficients. It is shown in [16, Section 5.2] that every sufficiently short (as always, under the canonical embedding) $e \in R^\vee$ is indeed the "canonical" representative of its coset modulo $qR^\vee$. Specifically, if $\|e\| < q/(2\sqrt{n})$ then $[e \bmod qR^\vee]_q = e$.

2.1.5 Good Bases of R and $R^\vee$

In this subsection we construct certain "good" bases of the ring R and its dual $R^\vee$ in terms of R' and $R'^\vee$ (respectively), and prove some of their useful geometrical properties. This (somewhat technical) material is used only in Section 3.1, where we prove the hardness of ring-LWE over K with secret in R', assuming its hardness over K' with secret in R'.

[6] This is easily seen, e.g., for $R = \mathbb{Z}[\zeta_{2^k}]$ and $R' = \mathbb{Z}$, where $\mathrm{Tr}(R) = 2^{k-1}R'$ because $\mathrm{Tr}(1) = 2^{k-1}$ and $\mathrm{Tr}(\zeta_{2^k}^j) = 0$ for $j = 1, \ldots, 2^{k-1} - 1$. More generally, $\mathrm{Tr}_{K/K'}(R)$ is often not even an integer multiple of R'.

Since K is a vector space of dimension n/n' over K', the field K has a K'*-basis* (which is not unique), i.e., a set of n/n' elements of K that are linearly independent over K', so that every element of K can be represented uniquely as a K'-linear combination of the basis elements. Similarly, an R'*-basis* of R is a set of n/n' elements in R, such that every element of R can be represented uniquely as an R'-linear combination of the basis elements. An $R'^{\vee}$-basis of $R^{\vee}$ is defined analogously.

We wish to construct an R'-basis of R, and a corresponding dual $R'^{\vee}$-basis of $R^{\vee}$ (any of which are K'-bases of K), which are "good" in the following sense: for any vector of K'-coefficients (with respect to the basis) which are *short* under σ', the corresponding K-element is also short under σ. More formally, represent an ordered K'-basis of K as a vector $\vec{b} = (b_j) \in K^{n/n'}$, and similarly for an arbitrary vector of K'-coefficients $\vec{a} = (a_j) \in K'^{(n/n')}$, which defines the K-element $a = \langle \vec{a}, \vec{b} \rangle = \sum_j a_j \cdot b_j$. Then by linearity, the basis $\vec{b}$ induces a matrix $B \in \mathbb{C}^{n \times n}$ such that

$$\sigma(a) = B \cdot \sigma'(\vec{a}), \quad \text{where } \sigma'(\vec{a}) = \big(\sigma'(a_j)\big)_j. \tag{2.3}$$

We seek an R'-basis $\vec{b}$ of R for which B (nearly) preserves Euclidean norms up to some scaling factor, i.e., all of its singular values are (nearly) equal.

In addition, for any K'-basis $\vec{b} = (b_j)$ of K, its *dual* K'-basis $\vec{b}^{\vee} = (b_j^{\vee}) \subseteq K$ is uniquely defined by the linear constraints $\mathrm{Tr}_{K/K'}(b_j \cdot b_{j'}^{\vee}) = 1$ if $j = j'$, and 0 otherwise. It is a straightforward exercise to verify that if $\vec{b}$ is an R'-basis of R, then $\vec{b}^{\vee}$ is an $R'^{\vee}$-basis of $R^{\vee}$. Moreover, the matrix $B^{\vee}$ induced by $\vec{b}^{\vee}$ is $B^{\vee} = B^{-T}$, so its singular values are simply the inverses of those of B.

Lemma 2.6. *Let $\hat{m} = m/2$ if m is even and m' is odd, otherwise $\hat{m} = m$, and let $r = \mathrm{rad}(m)/\mathrm{rad}(m')$ be the product of all primes that divide m but not m'. There exists an efficiently computable R'-basis $\vec{b}$ of R, for which the corresponding matrix B has largest and smallest singular values*

$$s_1(B) = \sqrt{\hat{m}/m'} \quad \textit{and} \quad s_n(B) = \sqrt{m/(rm')},$$

respectively. In particular, if $r \in \{1, 2\}$ then B is a unitary matrix scaled by a $\sqrt{\hat{m}/m'}$ factor.

Lemma 2.6 implies that for any $\vec{a} \in K'^{(n/n')}$ defining $a = \langle \vec{a}, \vec{b} \rangle \in K$ and $a^{\vee} = \langle \vec{a}, \vec{b}^{\vee} \rangle \in K$,

$$\|\sigma(a)\| \leq \sqrt{\hat{m}/m'} \cdot \|\sigma'(\vec{a})\| \quad \text{and} \quad \|\sigma(a^{\vee})\| \leq \sqrt{rm'/m} \cdot \|\sigma'(\vec{a})\|. \tag{2.4}$$

More generally, if the a_j are independent and have Gaussian distributions over (the canonical embedding of) K', then a and $a^{\vee}$ also have (possibly non-spherical) Gaussian distributions over K.[7] Since we are not too concerned with the exact distributions, we omit a precise calculation, which is standard. However, one particular case of interest is when the a_j are all i.i.d. according to a spherical Gaussian of parameter s, and $r \in \{1, 2\}$ so that B (respectively, $B^{\vee}$) is a scaled unitary matrix. Then because spherical Gaussians are invariant under unitary transformations, a (resp., $a^{\vee}$) is distributed according to a spherical Gaussian of parameter $s\sqrt{\hat{m}/m'}$ (resp., $s\sqrt{m'/\hat{m}}$).

The remainder of this subsection is devoted to proving Lemma 2.6. We denote the k-dimensional identity matrix by I_k, we use $\otimes$ to denote the Kronecker (or tensor) product of vectors and matrices, and we apply functions to vectors and matrices component-wise.

[7] To be completely formal, the Gaussians should be over the continuous spaces $K'_{\mathbb{R}} = K' \otimes_{\mathbb{Q}} \mathbb{R}$ and $K_{\mathbb{R}}$, which are essentially the "real analogues" of K' and K.

Following the treatment given in [16], let $m = \prod_\ell m_\ell$ be the prime-power factorization of m, i.e., the $m_\ell > 1$ are powers of distinct primes. The ring $R = \mathbb{Z}[\zeta_m]$ has the following $\mathbb{Z}$-basis $\vec{p}$, which is called the "powerful" basis:

$$\vec{p} = \bigotimes\nolimits_\ell \vec{p}_{m_\ell} \ , \quad \text{where } \vec{p}_{m_\ell} = \big(\zeta_{m_\ell}^j\big)_{j \in [\varphi(m_\ell)]} \ .$$

The set $\vec{p}_{m_\ell}$ is called the "power" $\mathbb{Z}$-basis of $\mathbb{Z}[\zeta_{m_\ell}] = \mathbb{Z}[\zeta_m^{m/m_\ell}] \subseteq R$.

Similarly, let $m' = \prod_\ell m'_\ell$ where each m'_ℓ divides m_ℓ, i.e., they are both powers of the same prime (though possibly $m'_\ell = 1$). Then the powerful $\mathbb{Z}$-basis of R' is defined as $\vec{p}\,' = \bigotimes_\ell \vec{p}_{m'_\ell}$, where the power bases $\vec{p}_{m'_\ell}$ are defined as above. Notice that when $m'_\ell > 1$, there is a bijective correspondence between $j \in [\varphi(m_\ell)]$ and $(j', k) \in [\varphi(m'_\ell)] \times [m_\ell/m'_\ell]$, via $j = (m_\ell/m'_\ell)j' + k$. Therefore, the power bases $\vec{p}_{m_\ell}$ factor as

$$\vec{p}_{m_\ell} = \vec{p}_{m'_\ell} \otimes \vec{b}_\ell, \quad \text{where } \vec{b}_\ell = \begin{cases} \big(\zeta_{m_\ell}^k\big)_{k \in [m_\ell/m'_\ell]} & \text{if } m'_\ell > 1 \\ \vec{p}_{m_\ell} & \text{if } m'_\ell = 1. \end{cases}$$

Hence, using the commutativity of the Kronecker product (up to some permutation) we can factor the powerful basis $\vec{p}$ of R as

$$\vec{p} = \vec{p}\,' \otimes \vec{b}, \quad \text{where } \vec{b} = \bigotimes\nolimits_\ell \vec{b}_\ell. \tag{2.5}$$

Because $\vec{p}\,'$ is a $\mathbb{Z}$-basis of R', it follows that $\vec{b}$ is an R'-basis of R. We next calculate the matrix $B \in \mathbb{C}^{n \times n}$ induced by $\vec{b}$, and verify that it indeed satisfies the claims in the lemma statement.

Following [16, Section 3], for any prime power $\tilde{m}$ we define $\mathrm{CRT}_{\tilde{m}}$ to be the complex $\varphi(\tilde{m})$-by-$\varphi(\tilde{m})$ matrix with $\omega_{\tilde{m}}^{i \cdot j}$ in its ith row and jth column, for $i \in \mathbb{Z}_{\tilde{m}}^*$ and $j \in [\varphi(\tilde{m})]$. Using the prime-power factorizations of our m, m', we define $\mathrm{CRT}_m = \bigotimes_\ell \mathrm{CRT}_{m_\ell}$ and $\mathrm{CRT}_{m'} = \bigotimes_\ell \mathrm{CRT}_{m'_\ell}$. Then up to a permutation of the rows (determined by the CRT correspondence between $\mathbb{Z}_m^*$ and $\prod_\ell \mathbb{Z}_{m_\ell}^*$), we have

$$\sigma(\vec{p}^{\,T}) = \mathrm{CRT}_m,$$

i.e., the columns of CRT_m are $\sigma(p_j)$ for each entry p_j of the row vector $\vec{p}^{\,T}$. In particular, $\sigma(\langle \vec{c}, \vec{p} \rangle) = \mathrm{CRT}_m \cdot \vec{c}$ for any $\vec{c} \in \mathbb{Q}^n$. Similarly, $\sigma'((\vec{p}\,')^T) = \mathrm{CRT}_{m'}$ up to a row permutation.

We now claim that, up to some permutations of B's rows and columns,

$$B = \mathrm{CRT}_m \cdot \big(\mathrm{CRT}_{m'}^{-1} \otimes I_{n/n'}\big) = \bigotimes\nolimits_\ell \Big(\mathrm{CRT}_{m_\ell} \cdot \big(\mathrm{CRT}_{m'_\ell}^{-1} \otimes I_{\varphi(m_\ell)/\varphi(m'_\ell)}\big)\Big), \tag{2.6}$$

where the second equality follows by the mixed-product property and the commutativity (up to row and column permutations) of the Kronecker product. To see the first equality, notice that for any $\vec{a} \in K'^{(n/n')}$ defining $a = \langle \vec{a}, \vec{b} \rangle \in K$, the matrix $(\mathrm{CRT}_{m'}^{-1} \otimes I)$ maps from (a suitable permutation of) the concatenated embeddings $\sigma'(\vec{a})$, to a vector $\vec{c} \in \mathbb{Z}^n$ of coefficients such that $\vec{a} = \langle \vec{c}, \vec{p}\,' \otimes I_{n/n'} \rangle$. In addition,

$$a = \langle \vec{a}, \vec{b} \rangle = \vec{c}^{\,T} \cdot (\vec{p}\,' \otimes I_{n/n'}) \cdot \vec{b} = \langle \vec{c}, \vec{p}\,' \otimes \vec{b} \rangle = \langle \vec{c}, \vec{p} \rangle.$$

Therefore, $\sigma(a) = \mathrm{CRT}_m \cdot \vec{c} = \mathrm{CRT}_m \cdot (\mathrm{CRT}_{m'}^{-1} \otimes I) \cdot \sigma'(\vec{a})$, as desired.

Now, by the last expression in Equation (2.6), and because singular values are multiplicative under the Kronecker product, from now on we drop all the ℓ subscripts, and assume without loss of generality that m and m' are powers of the same prime p (where possibly $m' = 1$). We analyze the singular values of $\mathrm{CRT}_m(\mathrm{CRT}_{m'}^{-1} \otimes I)$, for the cases $m' = 1$ and $m' > 1$. In the first case, clearly $\mathrm{CRT}_{m'} = I_1$, and it is shown in [16, Section 4] that the largest singular value of CRT_m is $\sqrt{m/2}$ if m is even and $\sqrt{m}$ otherwise, and its smallest singular value is $\sqrt{m/p}$.

For the case $m' > 1$, it follows from the decompositions given in [16, Section 3] that, up to some row permutation,

$$\mathrm{CRT}_m = \sqrt{m/p} \cdot Q \cdot (\mathrm{CRT}_p \otimes I_{m/p})$$

for some unitary matrix Q, and similarly for $\mathrm{CRT}_{m'}$. Then a routine calculation using elementary properties of the Kronecker product reveals that $\mathrm{CRT}_m(\mathrm{CRT}_{m'}^{-1} \otimes I)$ is some unitary matrix scaled by a $\sqrt{m/m'}$ factor, so all its singular values are $\sqrt{m/m'}$. This completes the proof of Lemma 2.6.

2.2 Homomorphic Cryptosystems

In ring-LWE-based cryptosystems for arbitrary cyclotomics [16] (generalizing those of [15, 4, 3]), the plaintext space is R_p for some integer $p \geq 2$ that is coprime with all the odd primes dividing m. (We assume that p is prime, which is really without loss of generality since we can always use the Chinese remainder theorem.) $\ll$Shai:I'm not sure what the "wlog" above refers to, if my circuit is mod-6 then p is not a prime. Anyway, I demoted this to parenthesis and made it vague enough so it must have some true interpretation..$\gg$ $\ll$Chris:What you wrote is a better way of saying what I had intended...$\gg$ Ciphertexts are elements of $(R_q^\vee)^2$ for some integer q that is coprime with p, and the secret key is some $s \in R$. A ciphertext $c = (c_0, c_1) \in (R_q^\vee)^2$ that encrypts a plaintext $b \in R_p$ with respect to s satisfies the decryption relation

$$c_0 + c_1 \cdot s = e \pmod{qR^\vee} \tag{2.7}$$

for some sufficiently short $e \in R^\vee$ such that $t \cdot e = b \pmod{pR}$. (Recall that $R^\vee = t^{-1}R$ for some $t \in R$, so $t \cdot e \in R$.) We refer to e as the *noise* of the ciphertext. Throughout this work we implicitly assume that the modulus q is large enough relative to $\|e\|$, so that $[c_0 + c_1 \cdot s]_q = e \in R^\vee$ (see Section 2.1.4 above). Therefore, the decryption algorithm can simply compute e and output $t \cdot e \bmod pR$. As shown in [4, 3, 16], this system (augmented by some additional public values, for greater efficiency) supports additive and multiplicative homomorphisms.

3 The Field-Switching Procedure

Our procedure performs the following operation. Given a big-field ciphertext $c \in (R_q^\vee)^2$ that encrypts a plaintext $b \in R_p$ with respect to a big-ring secret key $s \in R$, and a description of an R'-linear function $L\colon R_p \to R'_p$ to apply to the plaintext, it outputs a small-field ciphertext $c' \in (R_q'^\vee)^2$ that encrypts $b' = L(b) \in R'_p$ with respect to some small-ring secret key $s' \in R'$. (Recall that Corollary 2.5 characterizes how L corresponds to the induced function $\bar{L}\colon \mathbb{F}^f \to \mathbb{F}'^{f'}$ that is applied to the vector of finite field elements encoded by b.)

The procedure consists of the following three steps:

1. **Switch to a small-ring secret key.** We use the key-switching method from [5, 3, 16] to produce a ciphertext which is still over the big field K and encrypts the same plaintext $b \in R_p$, but with respect to a secret key $s' \in R' \subseteq R$ belonging to the small subring.

2. **Multiply by an appropriate (short) scalar.** We multiply the components of the resulting ciphertext by a short element $d \in R$ that corresponds to the desired R'-linear function to be applied to the input plaintext b.

3. **Map to the small field.** We map the resulting big-field ciphertext (over $R_q^\vee$) to a small-field ciphertext (over $R_q'^\vee$) by simply taking the trace $\mathrm{Tr}_{K/K'}$ of its two components. The resulting ciphertext will still be with respect to the small-ring secret key $s' \in R'$, but will encrypt the plaintext $b' = L(b) \in R'_p$.

Note that Steps 2 and 3 can be repeated multiple times on the same ciphertext (from Step 1), to apply several different R'-linear functions. In this way, the entire input plaintext can be preserved, but in a decomposed form.

3.1 Step 1: Switching to a Small-Ring Secret Key

To switch to a small-field secret key, we publish a "key-switching hint," which essentially encrypts the big-ring secret key $s \in R$ under the small-ring key $s' \in R'$, using ciphertexts over the big field. Note that encrypting under a small-ring secret key has security implications, since the dimension of the underlying RLWE problem is smaller. In our case, however, the whole point of switching to a smaller field is to obtain ciphertexts of a smaller dimension, so we do not actually lose any additional security by publishing the hint. Indeed, we show below that assuming the hardness of the decision RLWE problem in the small field, the key-switching hint reveals nothing about the big-ring secret key. The essence of that claim is Lemma 3.1 below, which says (informally) that RLWE in the big field, with secret chosen in the small ring $R' \subseteq R$, is no easier than RLWE in the small field.

Ring-LWE. The ring-LWE (RLWE) problem [15] (in K) with *continuous* error is parameterized by a modulus q, a "secret distribution" υ over R, and an "error distribution" ψ over K, which is usually a Gaussian (in the canonical embedding) and is therefore concentrated on short elements.[8] For $s \in R$, define the distribution $A_{s,\psi}$ that is sampled by choosing $\alpha \in R_q^\vee$ uniformly at random, choosing $\epsilon \leftarrow \psi$, and outputting the pair $(\alpha, \beta = \alpha \cdot s + \epsilon \bmod qR^\vee) \in R_q^\vee \times K/qR^\vee$. One equivalent form of the (average-case) decision $\mathrm{RLWE}_{q,\psi,\upsilon}$ problem (in K) is, given some ℓ pairs $(\alpha_i, \beta_i) \in R_q^\vee \times K/qR^\vee$, distinguish between the following two cases: in one case, the pairs are chosen independently from $A_{s,\psi}$ for a uniformly random $s \leftarrow \upsilon$ (which remains the same for all samples); in the other case, the pairs are all independent and uniformly random over $R_q^\vee \times K/qR^\vee$. For appropriate parameters q, ψ, υ and ℓ, solving this decision problem with non-negligible distinguishing advantage is as hard as approximating the shortest vector problem on ideal lattices in R, using a quantum algorithm. See [15, 16] for precise statements and further details.

Let $\vec{b}^\vee = (b_j^\vee)_{j\in[n/n']}$ be any $R'^\vee$-basis of $R^\vee$, and hence a K'-basis of K. Then for any error distribution ψ' over K', we can define an error distribution ψ over K as $\psi = \langle \psi'^{(n/n')}, \vec{b}^\vee \rangle$, i.e., a sample from ψ is generated by choosing independent $\epsilon_j \leftarrow \psi'$ (for $j \in [n/n']$) and outputting $\epsilon = \sum_j \epsilon_j b_j^\vee \in K$.

Lemma 3.1. *Let ψ' be an error distribution over K', and let $\psi = \langle \psi'^{(n/n')}, \vec{b}^\vee \rangle$ be the error distribution over K as described above. If the decision $\mathrm{RLWE}_{q,\psi',\upsilon'}$ problem (in K') is hard for some distribution υ' over $R' \subseteq R$, then the decision $\mathrm{RLWE}_{q,\psi,\upsilon'}$ problem (in K) is also hard.*

Although the lemma holds for any $R'^\vee$-basis of $R^\vee$, it is most useful with a basis having "good geometric properties." Specifically, in our case we need the property that if ψ' is concentrated on short elements of K', then ψ is similarly concentrated on short elements of K. Such a basis $\vec{b}^\vee$ is constructed in Lemma 2.6 of Section 2.1.5. For example, if ψ' is a continuous (spherical) Gaussian with parameter s and $r = \mathrm{rad}(m)/\mathrm{rad}(m') = 1$, then ψ' is a spherical Gaussian with parameter $s\sqrt{m'/m} = s\sqrt{n'/n}$.[9]

[8] Again, to be completely formal, a Gaussian should be defined over $K_\mathbb{R}$; see Footnote 7.

[9] Note that the factor $\sqrt{n'/n} \leq 1$ does not really amount to any effective decrease in the noise, because the "sparsity" of $R'^\vee$ versus $R^\vee$ is greater by a corresponding factor.

Proof. It suffices to give an efficient, deterministic reduction that takes n/n' pairs $(\alpha_j, \beta_j) \in R_q'^\vee \times K'/qR'^\vee$ and outputs a single pair $(\alpha, \beta) \in R^\vee \times K/qR^\vee$, with the following properties: if the pairs (α_j, β_j) are i.i.d. according to $A_{s',\psi'}$ for some $s' \in R'$, then (α, β) is distributed according to $A_{s,\psi}$; and if the pairs (α_j, β_j) are independent and uniformly random, then (α, β) is uniformly random. The reduction simply outputs $(\alpha = \langle \vec{\alpha}, \vec{b}^\vee \rangle, \beta = \langle \vec{\beta}, \vec{b}^\vee \rangle)$, where $\vec{\alpha} = (\alpha_j)_j$ and $\vec{\beta} = (\beta_j)_j$.

Since $\vec{b}^\vee$ is an $R'^\vee$-basis of $R^\vee$ and hence an $R_q'^\vee$-basis of $R_q^\vee$, it is immediate that the reduction maps the uniform distribution to the uniform distribution. On the other hand, if the samples $(\alpha_j \beta_j)$ are drawn from $A_{s',\psi'}$, i.e, $\beta_j = \alpha_j \cdot s' + \epsilon_j \bmod qR'^\vee$ for $\epsilon_j \leftarrow \psi$, then α is still uniformly random, and

$$\beta = \langle \vec{\beta}, \vec{b}^\vee \rangle = \langle \vec{\alpha}, \vec{b}^\vee \rangle \cdot s' + \langle \vec{\epsilon}, \vec{b}^\vee \rangle = \alpha \cdot s' + \epsilon \pmod{qR^\vee},$$

where $\vec{\epsilon} = (\epsilon_j)_j$ and ϵ has distribution ψ. This completes the proof. □

Key switching. In [5, 3, 16] it is shown how, given an $s \in R$ and sufficiently many RLWE samples (over K) with short noise and any secret $s' \in R$, it is possible to generate a "key-switching hint" with the following functionality: given the hint and any valid ciphertext c (over K) encrypted under s and with sufficiently short noise, it is possible to efficiently generate a ciphertext c' (also over K) with short noise encrypted under s'. Moreover, the hint is indistinguishable from uniformly random over its domain (even given s), assuming that the RLWE samples are.

For our transformation, we apply Lemma 3.1 using the "good basis" $\vec{b}^\vee$ from Lemma 2.6, thus obtaining RLWE samples over K relative to the secret $s' \in R' \subseteq R$, with noise distribution ψ which is concentrated on short vectors, and with security based on the hardness of RLWE$_{q,\psi',\upsilon'}$ problem in K'. We then construct the key-switching hint from these samples as described in [16, Section 7.3],

3.2 Steps 2 and 3: Mapping to the Small Field

Our goal now is to transform a valid big-field ciphertext $c = (c_0, c_1) \in (R_q^\vee)^2$, which encrypts some $b \in R_p$ with respect to some secret key $s' \in R' \subseteq R$, into a small-field ciphertext $c' = (c_0', c_1') \in (R_q'^\vee)^2$ that encrypts the related plaintext $b' = L(b)$ with respect to the same secret key s', where $L \colon R_p \to R_p'$ is any desired R'-linear function.

The process works as follows:

1. Since L is R'-linear, we can find some $d^\vee \in (t'R^\vee)_p$ such that L has the form $L(a) = \mathrm{Tr}_{K/K'}(d^\vee \cdot a)$.

2. Following the discussion in Section 2.1.4, we then find a *short* representative $d \in R$ such that $d = (t/t')d^\vee \pmod{pR}$, using a "good" basis of pR (i.e., one that has small singular values under σ, e.g., the "powerful" basis as constructed in Section 2.1.5).

   The chosen d defines the R'-linear function $L^\vee \colon R^\vee \to R'^\vee$ of the form $L^\vee(a^\vee) = \mathrm{Tr}_{K/K'}(d \cdot a^\vee)$, whose induced function from $R_p^\vee$ to $R_p'^\vee$ satisfies

$$t' \cdot L^\vee(a^\vee) = L(t \cdot a^\vee) \pmod{pR'}. \tag{3.1}$$

3. We obtain our small-field ciphertext by applying $L^\vee$ (or more precisely, the induced function from $R_q^\vee$ to $R_q'^\vee$) to c_0, c_1, setting

$$c_i' = L^\vee(c_i) = \mathrm{Tr}_{K/K'}(d \cdot c_i) \in R_q'^\vee, \quad i = 0, 1.$$

Lemma 3.2. *The ciphertext $c' = (c'_0, c'_1)$ is an encryption of $b' = L(b) \in R'_p$ under secret key $s' \in R'$, with noise $e' = L^\vee(e) \in R'^\vee$ of length $\|e'\| \leq \|e\| \cdot \|d\|_\infty \cdot \sqrt{n/n'}$, where e is the noise in the original ciphertext c.*

We note that the factor $\sqrt{n/n'}$ in the bound on $\|e'\|$ does not actually amount to any effective increase in the noise, because the dimension has decreased by a corresponding factor, and hence the size of e' relative to $R'^\vee$ remains the same as that of e relative to $R^\vee$. More precisely, the original ciphertext c decrypts correctly if $q > 2\sqrt{n}\|e\|$, whereas c' decrypts correctly if $q > 2\sqrt{n'}\|e'\|$. (See Section 2.1.4.) Therefore, the only real increase in the noise is due solely to $\|d\|_\infty$.

Proof. We need to show three things: that $\|e'\|$ is bounded as claimed, that $c'_0 + c'_1 \cdot s = e' \pmod{qR'^\vee}$, and that $t' \cdot e' = b' = L(b) \pmod{pR'}$.

1. The first claim follows immediately by Corollary 2.2 and the inequality $\|d \cdot e\| \leq \|d\|_\infty \cdot \|e\|$.

2. For the second claim, recall that $c_0 + c_1 \cdot s = e \pmod{qR^\vee}$. Then because the induced function $L^\vee \colon R^\vee_q \to R'^\vee_q$ is R'-linear and $s' \in R'$, we have

$$c'_0 + c'_1 \cdot s' = L^\vee(c_0 + c_1 \cdot s') = L^\vee(e) = e' \pmod{R'^\vee_q}.$$

3. For the last claim, because $t \cdot e = b \bmod pR$ and by Equation (3.1), we have

$$t' \cdot e' = t' \cdot L^\vee(e) = L(t \cdot e) = L(b) \pmod{pR'}.$$

□

3.3 Applying the Field-Switching Procedure

A typical application of the field-switching procedure during homomorphic evaluation of some circuit will begin with a big-field ciphertext that encrypts an array of plaintext values in the subfield $\mathbb{F}'$, as embedded in $\mathbb{F}$.[10] The above procedure is then applied to decompose the ciphertext into a number of small-field ciphertexts, each encrypting a subset of the plaintext values. Since big-field ciphertexts have room for f plaintext elements, but small-field ciphertexts can only hold f' elements, we may need up to f/f' small-ring ciphertexts to hold all the plaintext values that we are interested in. That is, we apply our field-switching transformation using the f'-fold concatenations $\bar{L}_i^{f'}$ of the $\mathbb{F}'$-linear selection functions $\bar{L}_i \colon \mathbb{F}^{f/f'} \to \mathbb{F}'$, $i \in [f/f']$, where $\bar{L}_i$ just selects the ith value (in $\mathbb{F}'$).[11]

Referring to Figure 1 for an example, the big-field ciphertext holds (up to) six plaintext values, and each small-field ciphertext can hold two values, with the big-field plaintext "slots" corresponding to $\mathfrak{p}_1, \mathfrak{p}_{15}, \mathfrak{p}_{22}$ lying over the small-field plaintext slot of $\mathfrak{p}'_1$, and the big-field slots corresponding to $\mathfrak{p}_3, \mathfrak{p}_{17}, \mathfrak{p}_{31}$ lying over the small-field plaintext slot of $\mathfrak{p}'_3$. Then we can produce three small-field ciphertexts, using the three selection functions

$$\begin{aligned}
(x_1, x_{15}, x_{22},\ x_3, x_{17}, x_{31}) &\mapsto (\ x_1\ \ ,\ x_3\ ),\\
(x_1, x_{15}, x_{22},\ x_3, x_{17}, x_{31}) &\mapsto (x_{15}\ , x_{17}),\\
(x_1, x_{15}, x_{22},\ x_3, x_{17}, x_{31}) &\mapsto (x_{22}\ , x_{31}).
\end{aligned}$$

[10] For example, when evaluating AES homomorphically, we would have plaintext values from $\mathbb{F}_{2^8}$ even though $\mathbb{F}$ may be a larger field such as $\mathbb{F}_{2^{16}}$ or $\mathbb{F}_{2^{24}}$, etc.

[11] More precisely, $\bar{L}_i(\vec{a}) = \mathrm{Tr}_{\mathbb{F}/\mathbb{F}'}(d \cdot a_i)$ for some $d \in \mathbb{F}$ such that $\mathrm{Tr}_{\mathbb{F}/\mathbb{F}'}(d) = 1$, so that $\bar{L}_i(\vec{a}) = a_i$ for any $a_i \in \mathbb{F}'$, by $\mathbb{F}'$-linearity.

Acknowledgments

The first and second authors are supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center (DoI/NBC) contract number D11PC20202. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBC, or the U.S. Government.

The third author is supported by the National Science Foundation under CAREER Award CCF-1054495, by the Alfred P. Sloan Foundation, and by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL) under Contract No. FA8750-11-C-0098. The views expressed are those of the authors and do not necessarily reflect the official policy or position of the National Science Foundation, the Sloan Foundation, DARPA or the U.S. Government.

The fourth author is supported by the European Commission through the ICT Programme under Contract ICT-2007-216676 ECRYPT II and via an ERC Advanced Grant ERC-2010-AdG-267188-CRIPTO, by EPSRC via grant COED–EP/I03126X, and by a Royal Society Wolfson Merit Award. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the European Commission or EPSRC.

References

[1] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO'09*, volume 5677 of *Lecture Notes in Computer Science*, pages 595–618. Springer, 2009.

[2] Zvika Brakerski. Fully Homomorphic Encryption without Modulus Switching from Classical GapSVP. In *CRYPTO'12*, volume 7417 of *Lecture Notes in Computer Science*. Springer, 2012. Available at `http://eprint.iacr.org/2012/078`.

[3] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. Fully homomorphic encryption without bootstrapping. In *Innovations in Theoretical Computer Science (ITCS'12)*, 2012. Available at `http://eprint.iacr.org/2011/277`.

[4] Zvika Brakerski and Vinod Vaikuntanathan. Fully homomorphic encryption from ring-LWE and security for key dependent messages. In *Advances in Cryptology - CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 505–524. Springer, 2011.

[5] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *FOCS'11*. IEEE Computer Society, 2011.

[6] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty Computation from Somewhat Homomorphic Encryption. In *CRYPTO'12*, volume 7417 of *Lecture Notes in Computer Science*. Springer, 2012. Available at `http://eprint.iacr.org/2011/535`.

[7] Leo Ducas and Alain Durmus. Ring-LWE in Polynomial Rings. In *PKC'12*, volume 7293 of *Lecture Notes in Computer Science*, pages 34–51. Springer, 2012. Available at `http://eprint.iacr.org/2012/235`

[8] Nicolas Gama and Phong Q. Nguyen. Predicting lattice reduction. In *EUROCRYPT'08*, volume 4965 of *Lecture Notes in Computer Science*, pages 31–51. Springer, 2008.

[9] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *STOC'09*, pages 169–178. ACM, 2009.

[10] Craig Gentry, Shai Halevi, Chris Peikert and Nigel P. Smart. Ring Switching in BGV-Style Homomorphic Encryption (preliminary version). Manuscript, `http://eprint.iacr.org/2012/240`, 2012.

[11] Craig Gentry, Shai Halevi, and Nigel P. Smart. Fully homomorphic encryption with polylog overhead. In *EUROCRYPT'12*, volume 7237 of *Lecture Notes in Computer Science*, pages 446-464, 2012. Available at `http://eprint.iacr.org/2011/566`.

[12] Craig Gentry, Shai Halevi, and Nigel P. Smart. Better bootstrapping for fully homomorphic encryption. In *PKC'12*, volume 7293 of *Lecture Notes in Computer Science*, pages 1–16. Springer, 2012. Available at `http://eprint.iacr.org/2011/680`.

[13] Craig Gentry, Shai Halevi, and Nigel P. Smart. Homomorphic evaluation of the AES circuit. In *CRYPTO'12*, volume 7417 of *Lecture Notes in Computer Science*, pages 850–867. Springer, 2012. Available at `http://eprint.iacr.org/2012/099`.

[14] Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. On-the-Fly Multiparty Computation on the Cloud via Multikey Fully Homomorphic Encryption In *STOC'12*. ACM, 2012.

[15] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT'10*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23, 2010.

[16] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for Ring-LWE cryptography. Manuscript, 2012

[17] Richard Lindner and Chris Peikert. Better key sizes (and attacks) for LWE-based encryption. In *CT-RSA*, volume 6558 of *Lecture Notes in Computer Science*, pages 319–339. Springer, 2011.

[18] Nigel P. Smart and Frederik Vercauteren. Fully homomorphic SIMD operations. Manuscript at http://eprint.iacr.org/2011/133, 2011.

[19] Lawrence C. Washington. *Introduction to Cyclotomic Fields*, volume 83 of *Graduate Texts in Mathematics*. Springer, 1996.