EXTENDED ABSTRACT APPEARS IN THE PROCEEDINGS OF SCN 2012.

Ring Switching in BGV-Style Homomorphic Encryption

Craig Gentry Shai Halevi Chris Peikert Nigel P. Smart

June 30, 2012

Abstract

The security of BGV-style homomorphic encryption schemes over polynomial rings relies on rings 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 rings, thus speeding up the homomorphic operations for the lower levels of the circuit. However, implementing such ring-switching is nontrivial, since these schemes rely on the ring algebraic structure for their homomorphic properties.

A basic ring-switching operation was used by Brakerski, Gentry and Vaikuntanathan, over polynomial rings 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 ring 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).

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.

Contents

1 Introduction

The last year has seen a rapid advance in the state of fully homomorphic encryption; yet despite these advances the existing schemes are still too inefficient for most practical purposes. In this paper we make another step forward in making such schemes more efficient. In particular we present a technique to reduce the dimension of the ring needed for homomorphic computation of the lower levels of a circuit. Our techniques apply to homomorphic encryption schemes over polynomial rings, such as the scheme of Brakerski et al. [4, 5, 3], as well as the variants due to Lòpez-Alt et al. [14] and Brakerski [2].

The most efficient variants of all these schemes work over polynomial rings of the form $\mathbb{Z}[X]/F(X)$, and in all of them the ring dimension (which is the degree of $F(X)$) must be set high enough to ensure security: to be able to handle depth-L circuits, these schemes must use key-switching matrices with modulus-to-noise ratio of $2^{\tilde{\Omega}(L\cdot\text{polylog}(\lambda))}$, hence the ring dimension must also be $\tilde{\Omega}(L \cdot \text{polylog}(\lambda))$ (even if we assume that ring-LWE [15] is hard to within fully exponential factors).[1] In practice, the ring dimension for moderately deep circuits can easily be many thousands. For example, to be able to evaluate AES homomorphically, Gentry et al. used in [13] circuits of depth $L \geq 50$, with corresponding ring-dimension of over 50000.

As homomorphic operations proceed, the noise in the ciphertext grows (or the modulus shrinks, if we use the modulus-switching technique from [5, 3]), hence reducing the modulus-to-noise ratio. Consequently, it becomes permissible to start using lower-dimension rings in order to speed up further homomorphic computation. However, in the middle of the computation we already have evaluated ciphertexts over the big ring, and so we need a method for transforming these into small-ring ciphertexts that encrypt the same thing. Such a "ring 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 polynomial rings of the form $R_{2^n} = \mathbb{Z}[X]/(X^{2^{n-1}}+1)$, and moreover by itself it cannot be combined with the "packed evaluation" techniques of Gentry et al. [11]. Extending this procedure is the focus of this work.

1.1 Our Contribution

In this work we present two complementary techniques:

- We extend the procedure from [3] to any cyclotomic ring $R = \mathbb{Z}[X]/\Phi_m(X)$ for a composite m. This is important, since the tools from [11] for working with "packed" ciphertexts require that we work with an odd integer m. For $m = u \cdot w$, we show how to break a ciphertext over the big ring R into a collection of $u' = \varphi(m)/\varphi(w)$ ciphertexts over the smaller ring $R' = \mathbb{Z}[X]/\Phi_w(X)$, such that the plaintext encrypted in the original big-ring ciphertext can be recovered as a simple linear function of the plaintexts encrypted in the smaller-ring ciphertexts.

- We then show how to take a "packed" big-ring ciphertext that contains many plaintext values in its plaintext slots, and distribute these plaintext values among the plaintext slots of several small-ring ciphertexts. If the original big-ring ciphertext was "sparse" (i.e., if only few of its plaintext slots were used), then our technique yields just a small number of small-ring ciphertexts, only as many as needed to fit all the used plaintext slots.

The first technique on its own may be useful in the context of bootstrapping, but it is not enough to achieve our goal of reducing the computational overhead by switching to small-ring ciphertexts, since we

[1] The schemes from [3, 2] can replace large rings by using higher-dimension vectors over smaller rings. But their most efficient variants use big rings and low-dimension vectors, since the complexity of their key-switching step is quadratic in the dimension of these vectors.

still need to show how to perform homomorphic operations on the resulting small-ring ciphertexts. This is achieved by utilizing the second technique. To demonstrate the usefulness of the second technique, consider the application of homomorphic AES computation [13], where the original big-ring ciphertext contains only 16 plaintext values (corresponding to the 16 bytes of the AES state). If the small-ring ciphertexts has 16 or more plaintext slots, then we can convert the original big-ring ciphertext into a single small-ring ciphertext containing the same 16 bytes in its slots, then continue the computation on this smaller ciphertext.

1.2 An Overview of the Construction

Our starting point is the polynomial composition technique of Brakerski et al. [3]. When $m = u \cdot w$ then a polynomial of degree less than $\varphi(m)$, $a(X) = \sum_{i=0}^{\varphi(m)-1} a_i X^i$, can be broken into u polynomials of degree less than $\varphi(m)/u \leq \varphi(w)$, by splitting the coefficients of a according to their index modulo u. Namely, denoting by $a_{(k)}$ the polynomial with coefficients $a_k, a_{k+u}, a_{k+2u}, \ldots$, we have

$$a(X) \;=\; \sum_{k=0}^{u-1} \sum_{j=0}^{\varphi(w)-1} a_{k+uj} X^{k+uj} \;=\; \sum_{k=0}^{u-1} X^k \sum_{j=0}^{\varphi(w)-1} a_{k+uj} X^{uj} \;=\; \sum_{k=0}^{u-1} X^k a_{(k)}(X^u). \qquad (1)$$

We note that this "very syntactic" transformation of splitting the coefficients of a high-degree polynomial into several low-degree polynomials, has the following crucial algebraic properties:

1. The end result is a collection of "parts" $a_{(k)}$, all from the small ring R' (which is a sub-ring of the big ring R, since $w|m$).
2. Recalling that $f(x) \mapsto f(x^u)$ is an embedding of R' inside R, we have the property that the original a can be recovered as a simple linear combination of (the embedding of) the parts $a_{(k)}$.
3. The transformation $T(a) = (a_{(0)}, \ldots, a_{(u-1)})$ is linear, and as such it commutes with the linear operations inside the decryption formula of BGV-type schemes: If $\mathfrak{s}$ is a big-ring secret key and c is (part of) a big-ring ciphertext, then decryption over the big ring includes computing $a = \mathfrak{s} \cdot c \in R$ (and later reducing a mod q and mod 2). Due to linearity, the parts of a can be expressed in terms of the tensor product between the parts of $\mathfrak{s}$ and c over the small ring. Namely, $T(\mathfrak{s} \cdot c)$ is some linear function of $T(\mathfrak{s}) \otimes T(c)$.

In addition to these algebraic properties, in the case considered in [3] where m, w are powers of two, it turns out that this transformation also possesses the following geometric property:

4. If a is a "short" element of R (in the canonical embedding of R), then all the components $a_{(k)} \in R'$ of $T(a)$ are also short in the canonical embedding of R'.

The importance of this last property stems from the fact that a valid ciphertext in a BGV-type homomorphic encryption scheme must have short noise, namely its inner product with the unknown secret key must be a short ring element. Property 3 above is used to convert a big-ring ciphertext encrypting a (relative to a big-ring secret key $\mathfrak{s}$) into a collection of "syntactically correct" small-ring ciphertexts encrypting the $a_{(k)}$'s (relative to the small-ring secret key $T(\mathfrak{s})$), and Property 4 is used to argue that these small-ring ciphertexts are indeed valid.

Attempting to apply the same transformation in the case where m, w are not powers of two, it turns out that the algebraic properties all still hold, but perhaps the geometric property does not. In this work we therefore describe a different transformation $T(\cdot)$ for breaking a big-ring element into a vector of small-ring elements, that has all the properties 1-4 above,[2] for any integers m, w such that $w|m$. This transformation crucially uses the interpretation of R as a dimension-$\varphi(m)/\varphi(w)$ extension ring of R', and is described in Section 3.2. Another advantage of our transformation over the one from [3] is that it breaks a big-ring element $a \in R$ into only $u' = \varphi(m)/\varphi(w)$ small ring parts $a_{(k)}$, as opposed to $u = m/w$ parts for the transformation from [3].

A Key-Switching Optimization. One source of inefficiency in the ring-switching procedure of Brakerski et al. [3] is that using the tensor product $T(\mathfrak{s}) \otimes T(c)$ amounts essentially to having u small-ring ciphertexts, each of which is a dimension-u vector over the small ring. Brakerski et al. point out that we can use key-switching/dimension-reduction to convert these high dimension ciphertexts into low-dimension ciphertexts over the small ring, but processing u ciphertexts of dimension u inherently requires work quadratic in u. Instead, here we describe an alternative procedure that saves a factor of u in running time.

Before using $T(\cdot)$ to break the ciphertext into pieces, we apply key-switching over the big ring to get a ciphertext with respect to another secret key that happens to belong to the small ring R'. (We again recall that R' is a sub-ring of R). The transformation $T(\cdot)$ has the additional property that when applied to a small-ring element $\mathfrak{s}' \in R' \subset R$, the resulting vector $T(\mathfrak{s}')$ over R' has just a single non-zero entry, namely $\mathfrak{s}'$ itself. Hence $T(\mathfrak{s}') \otimes T(c)$ is the same as just $\mathfrak{s}' \cdot T(c)$, and this lets us work directly with low-dimension ciphertexts over the small ring (as opposed to ciphertexts of dimension u). This is described in Section 3.1, where we prove that key-switching into a key from the small subring is as secure as ring-LWE in that small subring.

Packed Ciphertexts. As sketched so far, the ring-switching procedure lets us convert a big-ring ciphertext encrypting an element $a \in R$ into a collection of u' small-ring ciphertexts encrypting the parts $a_{(k)} \in R'$. However, coming in the middle of homomorphic evaluation, we may need to get small-ring ciphertexts encrypting elements other than the $a_{(k)}$'s. Specifically, if the original a encodes several plaintext values in its plaintext slots (as in [18, 11]), we may want to get encryptions of small-ring elements that encode the very same values in their slots.

We note that the plaintext values encoded in the element $a \in R$ are the evaluations $a(\rho_i)$, where the ρ_i's are primitive m-th roots of unity in some extension field $\mathbb{F}_{2^d}$. (Equivalently, the evaluations $a(\rho_i)$ correspond to the residues $a \bmod \mathfrak{p}_i$, where the $\mathfrak{p}_i = \langle 2, F_i(X) \rangle$ are the distinct prime ideal factors of $\langle 2 \rangle$ in the ring R. Hence the evaluation representation over $\mathbb{F}_{2^d}$ is just Chinese remaindering modulo 2 in R.)

Similarly, the plaintext values encoded in an element $b \in R'$ are the evaluations $b(\tau_j)$, where the τ_j's are primitive w-th roots of unity in $\mathbb{F}_{2^d}$ (equivalently, the residues of b modulo the prime ideal factors of 2 in R'). Our goal, then, is to decompose a big-ring ciphertext encrypting a into small-ring ciphertexts encrypting some b_k's, such that for every i there are some j, k for which $b_k(\tau_j) = a(\rho_i)$.

On a very high level, the approach that we take is to observe that the linear transformation $T(\cdot)$ for break big-ring elements into vectors of small-ring parts, must as a side-effect of induce some linear transformation (over $\mathbb{F}_{2^d}$) on the values in the plaintext slots. Hence after we apply T, we just need to compute homomorphically the inverse linear transformation (e.g., using the techniques from [11] for computing on packed ciphertexts), thereby recovering the original values.

[2] An earlier version of the current work [10] used the same transformation as in [3], and patched the problem with the geometric property by "lifting" everything from the big ring $\mathbb{Z}[X]/\Phi_m(X)$ to the even bigger ring $\mathbb{Z}[X]/(X^m - 1)$, using techniques similar to [11, 7].

2 Notation and Preliminaries

For any positive integer u we let $[u] = \{0, \ldots, u-1\}$.

2.1 Algebraic Background

Recall that an ideal (in an arbitrary commutative ring R) is an additive subgroup which is closed under multiplication by R. Below we typically denote ideals by $\mathfrak{p}, \mathfrak{q}$, etc. 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). When R' is a sub-ring of R and $\mathfrak{p}$ is an R'-ideal, we implicitly identify $\mathfrak{p}$ with its extension to R, namely the R-ideal $\mathfrak{p}R$. For an R-ideal $\mathfrak{p}$, the quotient ring $R_{\mathfrak{p}} = R/\mathfrak{p}R$ is a ring consisting of the residue classes $a + \mathfrak{p}$ for all $a \in R$, with the ring operations induced by R.

For any positive integer $m \geq 2$, let $\mathbb{K} = \mathbb{Q}(\zeta_m) \cong \mathbb{Q}[X]/\Phi_m(X)$ be the mth cyclotomic number field (of degree $\varphi(m)$), and $R = \mathbb{Z}[\zeta_m] \cong \mathbb{Z}[X]/\Phi_m(X)$ its ring of integers, where $\zeta_m = \exp(2\pi\sqrt{-1}/m)$ is the mth principal complex root of unity, and $\Phi_m(X) = \prod_{i \in \mathbb{Z}_m^*}(X - \zeta_m^i) \in \mathbb{Z}[X]$ is the mth cyclotomic polynomial. The elements ζ_m^j (equivalently, X^j) for $j \in [\varphi(m)]$ form a $\mathbb{Q}$-basis of $\mathbb{K}$ and a $\mathbb{Z}$-basis of R, called the "power basis." That is, any $a \in \mathbb{K}$ can be written uniquely as $a = \sum_{j=0}^{\varphi(m)-1} a_j \cdot \zeta_m^j$ for some $a_k \in \mathbb{Q}$, and $a \in R$ if and only if every $a_j \in \mathbb{Z}$.

There are $\varphi(m)$ ring homomorphisms from $\mathbb{K}$ to $\mathbb{C}$ that fix $\mathbb{Q}$ pointwise, called *embeddings*, which are denoted $\sigma_i \colon \mathbb{K} \to \mathbb{C}$ for $i \in \mathbb{Z}_m^*$ and characterized by $\sigma_i(\zeta_m) = \zeta_m^i$. (Equivalently, $\sigma_i(a(X)) = a(\zeta_m^i) \in \mathbb{C}$ when viewing $\mathbb{K}$ as $\mathbb{Q}(X)/\Phi_m(X)$.) We note that the σ_i are automorphisms of $\mathbb{K}$, when viewing it as a sub-field of $\mathbb{C}$. The (field) trace is a $\mathbb{Q}$-linear function $\mathrm{Tr}_{\mathbb{K}/\mathbb{Q}} \colon \mathbb{K} \to \mathbb{Q}$, which can be defined as the sum of the embeddings: $\mathrm{Tr}_{\mathbb{K}/\mathbb{Q}}(a) = \sum_{i \in \mathbb{Z}_m^*} \sigma_i(a)$.

The *canonical embedding* $\sigma \colon \mathbb{K} \to \mathbb{C}^{\varphi(m)}$ is the concatenation of all the embeddings, i.e., $\sigma(a) = (\sigma_i(a))_{i \in \mathbb{Z}_m^*}$, and it endows $\mathbb{K}$ with a canonical geometry. In particular, we define the Euclidean (ℓ_2) and ℓ_∞ norms on $\mathbb{K}$ 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 $\|a \cdot b\| \leq \|a\|_\infty \cdot \|b\|$ for any $a, b \in \mathbb{K}$, because the σ_i are ring homomorphisms.

For some $w|m$, let $u = m/w$ (so $\zeta_w = \zeta_m^u$) and $u' = \varphi(m)/\varphi(w)$, and let $\mathbb{K}' = \mathbb{Q}(\zeta_w) \subseteq \mathbb{K}$ and $R' = \mathbb{Z}[\zeta_w] \subseteq R$ be the wth cyclotomic number field and ring (respectively), with $\varphi(w)$ embeddings $\sigma_i' \colon \mathbb{K}' \to \mathbb{C}$ for $i \in \mathbb{Z}_w^*$ defining the canonical embedding $\sigma' \colon \mathbb{K}' \to \mathbb{C}^{\varphi(w)}$. Notice that when we restrict to the subfield $\mathbb{K}' = \mathbb{Q}(\zeta_m^u)$ of $\mathbb{K}$, for any $i \in \mathbb{Z}_m^*$ we have $\sigma_i = \sigma'_{i \bmod w}$, because $\sigma_i(\zeta_m^u) = \zeta_m^{u \cdot (i \bmod w)} = \sigma'_{i \bmod w}(\zeta_w)$.

Observe that using the polynomial representation in the small ring $R' \cong \mathbb{Z}[X]/\Phi_w(X)$, the element ζ_w is represented by the indeterminate X. However, using polynomial representation in the big ring, $R \cong \mathbb{Z}[X]/\Phi_m(X)$, the same ring element $\zeta_w = \zeta_m^u \in R' \subset R$ is represented by the monomial X^u. In general, if $r' \in R'$ is a small-ring element represented by the polynomial $b(X) \in \mathbb{Z}[X]/\Phi_w(X)$, then the same small-ring element is represented by the polynomial $a(X) = b(X^u) \bmod \Phi_m(X) \in \mathbb{Z}[X]/\Phi_m(X)$, when viewed as an element in the sub-ring R' of R. In other words, the mapping $f(X) \mapsto f(X^u) \bmod \Phi_m(X)$, mapping polynomials of degree less than $\varphi(w)$ into a subset of the polynomials of degree less than $\varphi(m)$, is a ring embedding of $\mathbb{Z}[X]/\Phi_w(X) \cong R'$ as a subring of $\mathbb{Z}[X]/\Phi_m(X) \cong R$. Similarly, this mapping is also a field embedding of $\mathbb{Q}[X]/\Phi_w(X) \cong \mathbb{K}'$ as a subfield of $\mathbb{Q}[X]/\Phi_m(X) \cong \mathbb{K}$.

We will use extensively the fact that $\mathbb{K}$ is a degree-u' extension of $\mathbb{K}'$, i.e., $\mathbb{K} = \mathbb{K}'(\zeta_m)$, and similarly $R = R'[\zeta_m]$. The powers ζ_m^k for $k \in [u']$ (also called the "power basis") form a $\mathbb{K}'$-basis of $\mathbb{K}$, and an

R'-basis of R. Looking ahead, our transformation $T(\cdot)$ for breaking a big ring element into small-ring components will just output the vector of coefficients of the big-ring element relative to the power basis.

One can verify that among all the embeddings σ_i of $\mathbb{K}$, exactly u' of them fix $\mathbb{K}'$ (not just $\mathbb{Q}$) pointwise. Specifically, these are the embeddings σ_i indexed by each $i = 1 \bmod w$. The intermediate trace function $\mathrm{Tr}_{\mathbb{K}/\mathbb{K}'} \colon \mathbb{K} \to \mathbb{K}'$ is a $\mathbb{K}'$-linear function, defined as the sum of all those $\mathbb{K}'$-fixing embeddings, i.e.,

$$\mathrm{Tr}_{\mathbb{K}/\mathbb{K}'}(a) = \sum_{i \in I} \sigma_i(a), \quad I = \{i \in \mathbb{Z}_m^* : i = 1 \bmod w\}.$$

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

The following lemma relates the intermediate trace to the embeddings of $\mathbb{K}$ and $\mathbb{K}'$, and will be used later to show that our ciphertext decomposition from R to R' produces component ciphertexts having short error terms.

Lemma 1. *For any $a \in \mathbb{K}$ and $i \in \mathbb{Z}_w^*$,*

$$\sigma'_i(\mathrm{Tr}_{\mathbb{K}/\mathbb{K}'}(a)) = \sum_{j = i \bmod w} \sigma_j(a).$$

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

Proof. Recall that for any $i' \in \mathbb{Z}_m^*$ such that $i' = i \bmod w$, the $\mathbb{K}'$-embedding σ'_i and the $\mathbb{K}$-embedding $\sigma_{i'}$ coincide on $\mathbb{K}'$. In particular, $\sigma'_i(\mathrm{Tr}_{\mathbb{K}/\mathbb{K}'}(a)) = \sigma_{i'}(\mathrm{Tr}_{\mathbb{K}/\mathbb{K}'}(a))$ because $\mathrm{Tr}_{\mathbb{K}/\mathbb{K}'}(a) \in \mathbb{K}'$. Then by definition of $\mathrm{Tr}_{\mathbb{K}/\mathbb{K}'}$ and linearity of $\sigma_{i'}$, we have

$$\sigma'_i(\mathrm{Tr}_{\mathbb{K}/\mathbb{K}'}(a)) = \sigma_{i'}\Bigg(\sum_{j = 1 \bmod w} \sigma_j(a)\Bigg) = \sum_j \sigma_{i'}(\sigma_j(a)) = \sum_{j' = i \bmod w} \sigma_{j'}(a),$$

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

2.2 The Big Ring-to-Small Ring Decomposition

As sketched in the introduction, our approach is rooted in the technique of decomposing an element of the "big" ring $R = \mathbb{Z}[\zeta_m]$ (or field $\mathbb{K}$) into several elements of the "small" ring $R' = \mathbb{Z}[\zeta_w]$ (or field $\mathbb{K}'$). Recall from Section 2.1 that $\mathbb{K} = \mathbb{K}'[\zeta_m]$ is a field extension of degree $u' = \varphi(m)/\varphi(w)$ over $\mathbb{K}'$, having power $\mathbb{K}'$-basis $\zeta_m^0, \ldots, \zeta_m^{u'-1}$. That is, any $a \in \mathbb{K}$ can be written uniquely as $a = \sum_{k=0}^{u'-1} a_k \cdot \zeta_m^k$ for some "coefficients" $a_k \in \mathbb{K}'$, and $a \in R$ if and only if every $a_k \in R'$. We define the decomposition map $T \colon \mathbb{K} \to (\mathbb{K}')^{u'}$ (which also maps R to $(R')^{u'}$) to simply output the vector of these coefficients:[3]

$$T(a) = (a_0, \ldots, a_{u'-1}). \tag{2}$$

We note a few simple but important properties of T:

[3] Alternatively, we could define T to output coefficients with respect to the "dual power" $\mathbb{K}'$-basis of $\mathbb{K}$, which would map the (fractional) dual ideal $R^\vee$ of R to $(R'^\vee)^{u'}$. That decomposition has better geometric properties and is more consonant with the ring-LWE problem as defined in [15], but it is more technically involved. We defer the details to the full version.

1. It is $\mathbb{K}'$-linear (and hence also R'-linear): for any $a, b \in \mathbb{K}$ and $r' \in \mathbb{K}'$, $T(a+b) = T(a) + T(b)$ (i.e., T is an additive homomorphism), and $T(r' \cdot a) = r' \cdot T(a)$.
2. Any ideal $\mathfrak{p}$ in R' induces a bijective $R'_{\mathfrak{p}}$-linear map $T_{\mathfrak{p}}\colon R_{\mathfrak{p}} \to (R'_{\mathfrak{p}})^{u'}$, namely, $T_{\mathfrak{p}}(a + \mathfrak{p}R) = T_{\mathfrak{p}}(a) + (\mathfrak{p})^{u'} = (a_0 + \mathfrak{p}, \ldots, a_{u'-1} + \mathfrak{p})$.

When using polynomial representations, the $\mathbb{K}'$-linearity of T must be interpreted relative to the embedding $f(X) \mapsto f(X^u)$ that maps the polynomial representation of $\mathbb{K}'$ into that of $\mathbb{K}$. Specifically, it means that for any polynomials $b(X) \in \mathbb{Q}[X]/\Phi_w(X)$ and $a(X) \in \mathbb{Q}[X]/\Phi_m(X)$, it holds that

$$T\big(\, b(X^u) \cdot a(X) \bmod \Phi_m(X) \,\big) \;=\; b(X) \cdot T\big(a(X)\big) \bmod \Phi_w(X). \tag{3}$$

Another important property is that T maps short elements in R to vectors of relatively short elements in R' (where as always, “short” is with respect to the canonical embeddings).

Lemma 2. *For $a \in R$, let $T(a) = (a_0, \ldots, a_{u'-1})$. Then for any $k \in [u']$, we have $\|a_k\| \le c_{m,w} \cdot \|a\| / \sqrt{u'}$, where $c_{m,w} \ge 1$ is a constant that depends only on m and w.*

Note that the $\sqrt{u'}$ term appearing above is merely a normalization factor associated with the fact that the power basis elements of R' are a $\sqrt{u'}$ factor shorter than those of R under the canonical embeddings, so the decomposition does *not* actually shrink the elements in any effective way.

Te constant $c_{m,w}$ turns out to depend only on the ratio $r = \operatorname{rad}(m)/\operatorname{rad}(w)$, where $\operatorname{rad}(n)$ denotes the radical of n, i.e., the product of all the prime divisors of n (without multiplicities). Hence we hereafter denote it by c_r rather than $c_{m,w}$. For typical values of $r = \operatorname{rad}(m)/\operatorname{rad}(w)$, the constant c_r is (somewhat) small, e.g., $c_1 = 1$ and $c_p = \sqrt{2 - 2/p}$ when p is a prime. (Hence if m and w share all the same prime divisors, the relevant constant is 1, and if m has only one additional prime divisor then the constant is smaller than $\sqrt{2}$.) Some other examples are $c_{3\cdot 5\cdot 7} \approx 17.4$ and $c_{5\cdot 7\cdot 11} \approx 155$. We also note that the constant factor c_r can actually be removed entirely, by following the framework of [15, 16] and defining T to work with the fractional ideals $R^\vee$ and $R'^\vee$ (as mentioned in Footnote 3); see the discussion after the proof of Lemma 2.

sketch. We first express T in terms of the intermediate field trace $\operatorname{Tr}_{\mathbb{K}/\mathbb{K}'}$, then use Lemma 1 to bound $\|a_k\|$. Recall that every $\mathbb{K}'$-linear map from $\mathbb{K}$ to $\mathbb{K}'$ can be expressed as $L(x) = \operatorname{Tr}_{\mathbb{K}/\mathbb{K}'}(d \cdot x)$ for some fixed $d \in \mathbb{K}$. Since T is $\mathbb{K}'$-linear, then for every $k \in [u']$ there exists $d_k \in \mathbb{K}$ such that $a_k = \operatorname{Tr}_{\mathbb{K}/\mathbb{K}'}(d_k \cdot a)$ (for all $a \in \mathbb{K}$). The elements d_k are “dual” to the power $\mathbb{K}'$-basis elements ζ_m^k: for every $j, k \in [u']$ we have $\operatorname{Tr}_{\mathbb{K}/\mathbb{K}'}(\zeta_m^j \cdot d_k) = 1$ if $j = k$ and $\operatorname{Tr}_{\mathbb{K}/\mathbb{K}'}(\zeta_m^j \cdot d_k) = 0$ if $j \ne k$.

Now by Lemma 1 and the fact that the σ_j are ring homomorphisms, we have

$$\sigma'(a_k) = P \cdot \sigma(d_k \cdot a) = P \cdot D \cdot \sigma(a),$$

where $D = \operatorname{diag}(\sigma(d_k))$. Notice that the rows of $P \cdot D$ are orthogonal (since each column has exactly one nonzero entry). The Euclidean norm of row $i \in \mathbb{Z}_w^*$ is $\|\mathbf{d}_{k,i}\|$, where $\mathbf{d}_{k,i} = (\sigma_j(d_k))_{j = i \bmod w} \in \mathbb{C}^{u'}$. Therefore, $\|a_k\| \le \|a\| \cdot \max_i \|\mathbf{d}_{k,i}\|$.

It remains to bound $\max_{k,i} \|\mathbf{d}_{k,i}\|$. For each $i \in \mathbb{Z}_w^*$, denote by Z_i the matrix (of dimension $u' \times u'$) defined as $Z_i = (\sigma_j(\zeta_m^k))_{j = i \bmod w, k \in [u']}$. Then the $\mathbf{d}_{k,i}$'s are determined by the linear constraints $Z_i^t \cdot \mathbf{d}_{k,i} = \mathbf{e}_k$ (where $\mathbf{e}_k \in \mathbb{Z}^{u'}$ is k'th standard basis vector). From Galois theory it follows that $\|\mathbf{d}_{k,i}\|$ is actually the same for every $i \in \mathbb{Z}_w^*$. It can also be shown that $\max_k \|\mathbf{d}_{k,1}\| \cdot \sqrt{u'}$ depends only on $\operatorname{rad}(m)/\operatorname{rad}(w)$; we omit the details. □

We note again that the constant c_r in Lemma 2 can be eliminated by defining the transformation T relative to a different basis, specifically the "dual" of the power basis, consisting of the vectors $d_0, d_1, \ldots, d_{u'-1}$ from the proof above. The proof then proceeds in the same way, but with the roles of d_k and ζ_m^k reversed. The tighter bound then follows by observing that the magnitude of each $\sigma_j(\zeta^k)$ is exactly one. One technical issue with using the dual basis, however, is that T no longer maps R to vectors over R'. Instead, it maps the dual ideal $R^\vee$ to vectors over $R'^\vee$, which introduces some additional algebraic subtleties but also turns out to have certain other advantages, as described in [15, 16]. We defer further details to the full version.

2.3 RLWE-Based Cryptosystems

Below and throughout this work, for a residue class $z + q\mathbb{Z} \in \mathbb{Z}_q$ we let $[z]_q \in \mathbb{Z}$ denote its canonical representative in the interval $[-q/2, q/2)$. (One can think of $[\cdot]_q$ as an operation that takes an arbitrary integer z and reduces it modulo q into the interval $[-q/2, q/2)$, so as to get the canonical representative of $z + q\mathbb{Z}$.) We extend this to a map from $R_q = R/qR$ to R, by applying the operation coefficient-wise to the input (viewed as a polynomial in coefficient representation). I.e., for $z = \sum_i z_i X^i \in R_q$ we get $[z]_q = \sum_i [z_i]_q \cdot X^i$. A standard fact is that if $z \in a + qR$ for some $a \in R$ that is sufficiently short relative to q and the dimension of R, then $[z]_q = a$. Throughout the paper we implicitly assume that q is chosen large enough to ensure that all of the operations we describe produce valid ciphertexts.

In a basic ring-LWE-based cryptosystem [15], secret keys and ciphertexts are elements of $(R_q)^2$ for some odd integer q, and moreover the secret key has the form $\mathbf{s} = (1, \mathfrak{s}) \bmod q$, where $\mathfrak{s} \in R$ is short. The plaintext space is the quotient ring $R_2 = R/2R$. A valid ciphertext $\mathbf{c} = (c_0, c_1) \in (R_q)^2$ that encrypts a plaintext $a \in R_2$ with respect to $\mathbf{s} = (1, \mathfrak{s})$ satisfies

$$\langle \mathbf{c}, \mathbf{s} \rangle = c_0 + \mathfrak{s} \cdot c_1 \in (a + 2e) + qR \tag{4}$$

for some sufficiently short $a + 2e \in R$. To decrypt, one just computes $[c_0 + \mathfrak{s} \cdot c_1]_q = a + 2e$ and reduces modulo 2 to recover the plaintext a. Additionally, Brakerski et al. [4, 3] showed that this system (with certain additions to the public key) supports additive and multiplicative homomorphisms.

Our ring-switching procedure will be given a ciphertext where Equation (4) holds over R (for some $\mathfrak{s} \in R$), and will output ciphertexts for which the equality holds over R' (for a different secret $\mathfrak{s}' \in R'$).

2.4 Plaintext Arithmetic

Following [18, 3, 11, 12, 13], we recall how to encode vectors over a certain finite field into the message spaces. A summary is provided in Figure 1 below.

For concreteness, we focus first on R', viewing it as $\mathbb{Z}[X]/\Phi_w(X)$. Let d' be the order of 2 in the multiplicative group $\mathbb{Z}_w^*$. Then $\Phi_w(X)$ factors modulo 2 into $\ell' = \varphi(w)/d'$ distinct irreducible (over $\mathbb{F}_2$) polynomials $F_i(X)$, each of degree d'. The ideal $2R'$ has factorization $2R' = \prod_{i=1}^{\ell'} \mathfrak{p}_i$, where $\mathfrak{p}_i = \langle 2, F_i(X) \rangle$ are distinct prime ideals. Since each $F_i(X)$ is irreducible modulo 2, each $R'/\mathfrak{p}_i R' = \mathbb{F}_2[X]/\langle F_i(X) \rangle$ is isomorphic to the finite field $\mathbb{F}_{2^{d'}}$. By the Chinese remainder theorem, we can therefore identify elements of R'_2 with elements of $(\mathbb{F}_{2^{d'}})^{\ell'}$, as summarized by the following diagram of ring isomorphisms.

$$\begin{tikzcd} R'/2R' \arrow[r, leftrightarrow, "\text{CRT}"] & \bigoplus_i (R'/\mathfrak{p}_i R') \arrow[r, leftrightarrow] & (\mathbb{F}_{2^{d'}})^{\ell'} \end{tikzcd}$$

For our ring-switching application we use a particular ring isomorphism between $R'/2R'$ and $(\mathbb{F}_{2^{d'}})^{\ell'}$, for some fixed representation of $\mathbb{F}_{2^{d'}}$. Consider the quotient group $\mathbb{Z}_w^* / \langle 2 \rangle$ (which has cardinality ℓ'), and fix a specific set of representatives for this quotient group, $U_w = \{j_0, j_1, \ldots, j_{\ell'-1}\} \subseteq \mathbb{Z}_w^*$, containing of exactly one member from every conjugacy class in $\mathbb{Z}_w^* / \langle 2 \rangle$.[4] Also fix a specific primitive w-th root of unity $\tau \in \mathbb{F}_{2^{d'}}$, and identify each element $a \in R_2'$ with the ℓ'-vector consisting of $a(\tau^j) \in \mathbb{F}_{2^{d'}}$ for all $j \in U_w$:

$$a \in R_2' \;\longleftrightarrow\; \left\langle a(\tau^{j_1}), \ldots, a(\tau^{j_{\ell'}}) \right\rangle \in \left(\mathbb{F}_{2^{d'}}\right)^{\ell'}.$$

Showing that this is indeed a bijection is standard. In one direction, from a we can compute all the values $a(\tau^{j_k})$. In the other direction we have the following simple claim:

Claim 1. *For every vector $(\alpha_0, \alpha_1, \ldots, \alpha_{\ell'-1}) \in \mathbb{F}_{2^{d'}}^{\ell'}$, there is a unique polynomial $a \in R_2'$ such that over $\mathbb{F}_{2^{d'}}$ it holds that $a(\tau^{j_k}) = \alpha_k$ for all $k \in [\ell']$.*

Proof. We identify R_2' with $\mathbb{F}_2[X]/\Phi_w(X) \subset \mathbb{F}_{2^{d'}}[X]/\Phi_w(X)$, and recall that a polynomial $a \in \mathbb{F}_{2^{d'}}[X]/\Phi_w(X)$ belongs to the subring R_2' if and only if $a(X^2) = a(X)^2$ (as an identity in R_2'). Given a vector of values $(\alpha_0, \alpha_1, \ldots, \alpha_{\ell'-1}) \in \mathbb{F}_{2^{d'}}^{\ell'}$, we can therefore deduce from $a(\tau^{j_k}) = \alpha_k$ the evaluations of a on the other members of the same conjugacy class, namely $a(\tau^{2j_k}) = \alpha_k^2,\ a(\tau^{4j_k}) = \alpha_k^4,\ a(\tau^{8j_k}) = \alpha_k^8$, etc. Since U_w is a complete set of representatives for the quotient group $\mathbb{Z}_w^* / \langle 2 \rangle$, we can get in this way the evaluations of $a(\tau^j)$ for all the indices $j \in \mathbb{Z}_w^*$. This gives us the evaluation of a at $\varphi(w)$ different points, from which a is uniquely defined (because $\mathbb{F}_{2^{d'}}$ is a field and a has degree less than $\varphi(w)$). □

We thus view the evaluation of the plaintext element at τ^{j_k} as the k'th "plaintext slot," and note that arithmetic operations in the ring R_2' act on the plaintext slots in a componentwise manner.

For $R \cong \mathbb{Z}[X]/\Phi_m(X)$ the analysis proceeds similarly. Let d be the order of 2 in the multiplicative group $\mathbb{Z}_m^*$, so $d'|d$, and let $\ell = \varphi(m)/d$. Recalling that $F_i(X^u)$ is the embedding of $F_i(X) \in R'$ into R, we denote the factorizaton of $F_i(X^u)$ into irreducible factors modulo 2 by $F_i(X^u) = \prod_j F_{i,j}(X)$. We note that each $F_i(X)$ factors into exactly ℓ/ℓ' distinct irreducible (mod 2) factors, each of degree d, and that the factorization of $\Phi_m(X)$ into irreducible factors mod 2 is $\Phi_m(X) = \prod_{i,j} F_{i,j}(X)$. Therefore, each prime ideal $\mathfrak{p}_i$ in R' factors further in R, into the product of the ℓ/ℓ' prime ideals $\mathfrak{p}_{i,j} = \langle 2, F_{i,j}(X) \rangle$, where each $R/\mathfrak{p}_{i,j}R$ is isomorphic to $\mathbb{F}_{2^d}$.

We use a concrete ring isomorphism between $R/2R$ and $(\mathbb{F}_{2^d})^\ell$ analogous to the one described above, using some representative set U_m of the quotient group $\mathbb{Z}_m^* / \langle 2 \rangle$ and a primitive m-th root of unity ρ, and considering the "plaintext slots" of $a \in R_2$ as the evaluations $a(\rho^i)$ for all $i \in U_m$. Of course, the analog of Claim 1 holds here too.

$$\begin{tikzcd}
R/2R \arrow[r, leftrightarrow, "\text{CRT}"] \arrow[d, leftrightarrow, dashed, "T_2"'] & \bigoplus_i (R/\mathfrak{p}_i R) \arrow[r, leftrightarrow, "\text{CRT}"] \arrow[d, leftrightarrow, dashed, "\bigoplus_i T_{\mathfrak{p}_i}"'] & \bigoplus_{i,j} (R/\mathfrak{p}_{i,j} R) \arrow[r, leftrightarrow] & (\mathbb{F}_{2^d})^\ell \\
(R'/2R')^{u'} \arrow[r, leftrightarrow, "\text{CRT}"] & \bigoplus_i (R'/\mathfrak{p}_i R')^{u'} \arrow[rr, leftrightarrow] & & (\mathbb{F}_{2^{d'}})^{\ell' \cdot u'}
\end{tikzcd}$$

Figure 1: Commutative diagram of various representations of the plaintext spaces, and morphisms between them. Solid lines are ring isomorphisms, and dashed lines are R'-linear homomorphisms (i.e., satisfying $T(x+y) = T(x) + T(y)$ and $T(rx) = rT(x)$ for all $r \in R'$).

[4] In other words, the sets $U_w, 2U_w, 4U_w, \ldots 2^{d'-1}U_w$ are all disjoint, and their union is the entire group $\mathbb{Z}_w^*$.

3 The Ring-Switching Procedure

Given a big-ring ciphertext $\mathbf{c} \in (R_q)^2$ that encrypts a plaintext $a \in R_2$ relative to a big-ring secret key $\mathfrak{s} \in R$, our goal is to output u' small-ring ciphertexts $\mathbf{c}_k \in (R'_q)^2$ for $k \in [u']$, where each $\mathbf{c}_k$ encrypts $a_k \in R'_2$, namely the kth component $T(a)$, all relative to some small-ring secret key $\mathfrak{s}' \in R'$. The procedure consists of the following two steps:

1. **Key-switch.** We use the key-switching method from [5, 3] to switch to a ciphertext that is still over the big ring R, but which has a secret key $\mathfrak{s}' \in R'$ belonging to the small subring $R' \subseteq R$.

2. **Decompose.** We break the resulting big-ring ciphertext (over R_q) into u' small-ring ciphertexts (over R'_q) using the decomposition T_q. These ciphertexts will be valid with respect to the small-ring secret key $\mathfrak{s}' \in R'$, and will encrypt the components of $T(a)_k$ as desired (see Lemma 4).

3.1 Switching to a Small-Ring Secret Key

To enable this transformation, we include in the public key a "key switching hint," essentially encrypting the old big-ring key $\mathfrak{s}$ under the new small-ring key $\mathfrak{s}'$. Note that using such a small-ring secret key has security implications, since it severely reduces the dimension of the underlying LWE problem. In our case, however, the whole point of switching to a smaller ring is to get ciphertexts over a smaller dimension, so we are not actually losing any additional security by giving out the hint. Indeed, we show below that assuming the hardness of the decision-ring-LWE problem [15] over the small ring R'_q, the key-switching hint is indistinguishable from uniformly random over R_q (even for a distinguisher that knows the old secret key $\mathfrak{s}$).

Ring-LWE. The ring-LWE (RLWE) problem [15] over R'_q is parameterized by an error distribution χ' over R', typically derived from a Gaussian and so highly concentrated on short elements.[5] For a "secret" element $\mathfrak{s}' \in R'$, a sample from the RLWE distribution $A_{\mathfrak{s}',\chi'}$ is generated by choosing $\alpha \in R'_q$ uniformly at random and $\epsilon \leftarrow \chi'$, computing $\beta \leftarrow \alpha \cdot \mathfrak{s}' + \epsilon$ in R'_q, then outputting the pair $(\alpha, \beta) \in (R'_q)^2$. The decision RLWE problem in R'_q is: given arbitrarily many pairs $(\alpha_i, \beta_i) \in (R'_q)^2$, distinguish the case where the samples are chosen independently from $A_{\mathfrak{s}',\chi'}$ (for a single $\mathfrak{s}' \leftarrow \chi'$) from the case where they are uniformly random and independent.

To set up the key-switching technique, we first prove a lemma of independent interest about the hardness of RLWE over the big ring R_q when the secret is chosen according to χ' from the subring R'. Define an error distribution χ over R as $\chi = T^{-1}((\chi')^{u'})$, i.e., a sample from χ is generated by choosing independent $\epsilon_i \leftarrow \chi'$ for $i \in [u']$, and outputting $\epsilon = T^{-1}(\epsilon_0, \ldots, \epsilon_{u'-1}) = \sum_i \epsilon_i \cdot \zeta_m^i \in R$. Note that elements drawn from χ are short: because $\|\sigma(\zeta_m^i)\|_\infty = 1$ for all i, we have

$$\|\sigma(\epsilon_i \cdot \zeta_m^i)\| = \|\sigma(\epsilon_i)\| = \sqrt{u'} \cdot \|\sigma'(\epsilon_i)\|$$

(where as usual, the $\sqrt{u'}$ term is effectively a normalization factor between R' and R). Then by the triangle inequality, $\|\sigma(\epsilon)\| \leq (u')^{3/2} \cdot B$, where B is an upper bound on every $\|\sigma'(\epsilon_i)\|$. (Tighter bounds can also be obtained when χ' is Gaussian, as is typical with RLWE.)

Lemma 3. *If the decision RLWE problem over R' with error distribution χ' is hard, then so is the decision RLWE problem over R with error distribution $\chi = T^{-1}((\chi')^{u'})$, but where the secret is chosen from χ', and in particular is in subring R'.*

[5] Or, following [15] more closely, χ' would be a distribution over the dual $R'^{\vee}$.

Proof. It suffices to give a reduction that maps small-ring samples over $(R'_q)^2$, drawn from $A_{\mathfrak{s}',\chi'}$ (respectively, the uniform distribution), to big-ring samples over $(R_q)^2$ with distribution $A_{\mathfrak{s}',\chi}$ (resp., the uniform distribution). To generate each output sample, the reduction takes u' fresh input samples $(\alpha_i, \beta_i) \in (R'_q)^2$ for $i \in [u']$, defines $\alpha' = (\alpha_i)_i, \beta' = (\beta_i)_i \in (R'_q)^{u'}$, and outputs $(\alpha, \beta) = (T_q^{-1}(\alpha'), T_q^{-1}(\beta')) \in R_q$.

Since T_q is a bijection, it is clear that the reduction maps the uniform distribution to the uniform distribution. On the other hand, if the samples $(\alpha_i, \beta_i = \alpha_i \cdot \mathfrak{s}' + \epsilon_i)$ are drawn from $A_{\mathfrak{s}',\chi'}$ for some $\mathfrak{s}' \leftarrow \chi'$, then α is still uniformly random, and moreover, letting $\epsilon' = (\epsilon_i)_i \in (R')^{u'}$ and by R'-linearity of T_q^{-1}, we have (over R_q)

$$\beta = T_q^{-1}(\alpha' \cdot \mathfrak{s}' + \epsilon') = T_q^{-1}(\alpha') \cdot \mathfrak{s}' + T^{-1}(\epsilon') = \alpha \cdot \mathfrak{s}' + \epsilon,$$

where $\epsilon = T^{-1}(\epsilon')$ is distributed according to χ by construction. So (α, β) is distributed according to $A_{\mathfrak{s}',\chi}$, as desired. □

The key-switching hint. Let $\mathfrak{s} \in R$ be the big-ring secret key, and $\mathfrak{s}' \in R' \subset R$ be the small-ring secret key that we want to switch to. To construct the key-switching hint, we independently draw $l = \lceil \log_2 q \rceil$ error terms $\epsilon_i \leftarrow \chi$ and uniformly random elements $\alpha_i \in R_q$, for $i \in [l]$. The hint consists of the all the pairs[6]

$$(\alpha_i, \beta_i = 2^i \cdot \mathfrak{s} - \alpha_i \cdot \mathfrak{s}' + 2\epsilon_i) \in (R_q)^2.$$

For security, note that by the form of the hint, it is immediate from Lemma 3 that for any big-ring secret key $\mathfrak{s} \in R$, the hint (even along with $\mathfrak{s}$) is computationally indistinguishable from uniform.

Since the errors ϵ_i are short, the hint is functional for key-switching, as described in [5]. Specifically, suppose we are given a valid ciphertext $\mathbf{c} = (c_0, c_1)$ relative to $\mathfrak{s}$, for which $c_0 + \mathfrak{s} \cdot c_1 = (a + 2e) \bmod q$ for some short $(a + 2e) \in R$. We decompose c_1 into its bitwise representation as $c_1 = \sum_{i \in [l]} 2^i d_i \bmod q$ for short elements $d_i \in R$ having 0-1 coefficients in the power basis. We then have the relation (over R_q)

$$\begin{aligned} \underbrace{c_0 + \sum_i d_i \beta_i}_{c'_0} + \underbrace{\sum_i d_i \alpha_i}_{c'_1} \cdot \mathfrak{s}' &= c_0 + \sum_i d_i (2^i \mathfrak{s} + 2\epsilon_i) = c_0 + c_1 \cdot \mathfrak{s} + 2 \sum_i d_i \epsilon_i \\ &= a + 2(e + \sum_i d_i \epsilon_i). \end{aligned}$$

Since $\sum_i d_i \epsilon_i \in R$ is short, (c'_0, c'_1) is a valid ciphertext encrypting a under $\mathfrak{s}'$, as desired.

3.2 Decomposing the Ciphertext

After switching to a small-ring secret key $\mathfrak{s}' \in R'$ in the previous step, the ciphertext is a pair $\mathbf{c} = (c_0, c_1) \in (R_q)^2$ such that

$$c_0 + \mathfrak{s}' \cdot c_1 \in (a + 2e) + qR,$$

where $a + 2e \in R$ is sufficiently short. We decompose this ciphertext into u' ciphertexts $\mathbf{c}_k = (c_{k,0}, c_{k,1}) \in (R'_q)^2$ for $k \in [u']$, where for $b \in \{0, 1\}$, $T_q(c_b) = (c_{0,b}, \ldots, c_{u'-1,b})$. (Recall that $T_q \colon R_q \to (R'_q)^{u'}$ is the R'-linear bijection induced by the decomposition T defined in Section 2.2.)

Lemma 4. *If $\mathbf{c}$ is a valid encryption of plaintext $a \in R_2$ under secret key $\mathfrak{s}' \in R'$, then each $\mathbf{c}_k$ is a valid encryption of the kth component of $T_2(a) \in (R'_2)^{u'}$.*

[6] We could alternatively use the key-switching variant from [13], where the hint consists of a single pair (β, α), but with respect to a large modulus $Q \approx q^2 \cdot m$. The proof of security would then depend on the hardness of ring-LWE in R'_Q rather than in R'_q.

Proof. Below we identify $\mathfrak{s}' \in R'$ with its mod-q equivalence class $\mathfrak{s}' + qR' \in R'_q$. Because T_q is R'_q-linear, we have

$$T_q(c_0) + \mathfrak{s}' \cdot T_q(c_1) = T_q(c_0 + \mathfrak{s}' \cdot c_1) = T(a + 2e) + (qR')^{u'},$$

where the multiplication of scalar $\mathfrak{s}' \in R'$ with $T_q(c_1) \in (R'_q)^{u'}$ is coordinate-wise. By Lemma 2, each component of $T(a + 2e)$ has length bounded by $c_r \cdot \|a + 2e\|/\sqrt{u'}$ (where the $\sqrt{u'}$ term is a normalization factor), so the "effective" lengths (relative to q and the dimension of R') grow by at most a fixed constant factor c_r, and are sufficiently small. Moreover, $T(a + 2e) \in T_2(a) \in (R'_2)^{u'}$, so the message encrypted by $\mathbf{c}_k$ is the kth component of $T_2(a)$. □

4 Homomorphic Computation in the Small Ring

So far we have shown how to break a big-ring ciphertext, encrypting some big-ring element $a \in R_2$, into a collection of u' small-ring ciphertexts encrypting the small-ring elements $T(a) = (a_0, a_1, \ldots, a_{u'-1}) \in R'_2$. This, however, still falls short of our goal of speeding-up homomorphic computation by switching to small-ring ciphertexts. Indeed we have not shown how to use the encryption of the a_k's for further homomorphic computation.

Following the narrative of SIMD homomorphic computation from [18, 11, 12, 13], we view the big-ring plaintext element $a \in R_2$ as an encoding of a vector of plaintext values from the extension field $\mathbb{F}_{2^d}$ (with d the order of 2 in $\mathbb{Z}_m^*$). We therefore wish to obtain small-ring ciphertexts encrypting small-ring elements that encode of the same underlying $\mathbb{F}_{2^d}$ values.

One potential "algebraic issue" with this goal, is that it is not always possible to embed $\mathbb{F}_{2^d}$ values inside small-ring elements from R'_2. Recall that the extension degree d is determined by the order of 2 in $\mathbb{Z}_m^*$. But the order of 2 in $\mathbb{Z}_w^*$ may be smaller than d, in general it will be some d' that divides d. If $d' < d$ then we can only embed values from the sub-field $\mathbb{F}_{2^{d'}}$ in small-ring element from R'_2, and not the $\mathbb{F}_{2^d}$ values that we have encoded in the big-ring element a. For the rest of this section we only consider the special case where the order of 2 in both $\mathbb{Z}_m^*$ and $\mathbb{Z}_w^*$ is the same d, leaving the general case to the full version.

Even for the special case where the order of 2 in $\mathbb{Z}_m^*$ and $\mathbb{Z}_w^*$ is the same (and hence the "plaintext slots" in the small ring contain values from the same extension field as those in the big ring), we still need to tackle the issue that big ring elements have more plaintext slots than small ring elements. Specifically, big-ring elements have $\ell = \varphi(m)/d$ slots, whereas small-ring elements only have $\ell' = \varphi(w)/d$ slots. The solution here is obvious: we just use more small-ring elements to hold all the plaintext slots that we need.

Note that if the original plaintext element a was "sparsely populated", holding only a few plaintext values in its slots, then we would like to generate only as many small-ring ciphertexts as needed to hold these few plaintext slots. A good example is the computation of the AES circuit in [13]: Since there are only 16 bytes in the AES state, we only use 16 slots in the plaintext element a. In this case, as long as we have at least 16 slots in small-ring elements, we can continue working with a single small-ring ciphertext (as opposed to the u' ciphertexts that the technique of the previous section gives us).

4.1 Ring-Switching with Plaintext Encoding

Below we describe our method for converting the plaintext encoding between the different rings, for the special case where the order of 2 is the same in $\mathbb{Z}_m^*$ and $\mathbb{Z}_w^*$. As sketched in the introduction, the basic observation underlying our approach is that the transformation $T(a) = (a_0, a_1, \ldots, a_{u'-1})$ that we apply to our plaintext when breaking a big-ring ciphertext into its small-ring parts, induces a linear transformation over the values in the plaintext slots. We then just finish-up the process by homomorphically computing the inverse linear transformation over the resulting small-ring ciphertexts (using "general purpose" techniques for computing on packed ciphertexts, such as in [11]), thereby restoring the plaintext slots to their original values.

As explained in Section 2.4, each plaintext slot in the big-ring element is associated with a member of the quotient group $\mathcal{Q}_m = \mathbb{Z}_m^* / \langle 2 \rangle$, and similar association holds between plaintext slots in small-ring elements and members of the quotient group $\mathcal{Q}_w = \mathbb{Z}_w^* / \langle 2 \rangle$. We thus begin by relating the structures and representations of these two quotient groups.

Below let $U_w \subseteq \mathbb{Z}_w^*$ be a representative set for $\mathcal{Q}_w$. i.e., a set containing exactly one index from each conjugacy class in $\mathbb{Z}_w^* / \langle 2 \rangle$. It is easy to see that when the order of 2 is the same in $\mathbb{Z}_m^*$ and $\mathbb{Z}_w^*$, then the set $U_m = \{j \in \mathbb{Z}_m^* : \exists\, i \in U_w \text{ s.t. } j \equiv i \pmod{w}\}$ is a representative set for $\mathcal{Q}_m$. Fixing in addition a primitive m'th root of unity $\rho \in \mathbb{F}_{2^d}$ and the particular primitive w'th root of unity $\tau = \rho^u$, we let the plaintext slots encoded in $a \in R_2$ be the evaluations $a(\rho^j) \in \mathbb{F}_{2^d}$ for $j \in U_m$, and similarly the plaintext slots encoded in $a' \in R_2'$ be the evaluations $a'(\tau^i)$ for $i \in U_w$.

We proceed to prove that under this representation, the transformation T from Section 2.2 induces an $\mathbb{F}_{2^d}$-linear transformation on the values in the $\mathbb{F}_{2^d}$ values in the plaintext slots. A key lemma is the following:

Lemma 5. *Let $m = u \cdot w$ for odd integers u, w, such that the order of 2 is the same in $\mathbb{Z}_m^*$ and in $\mathbb{Z}_w^*$. Let U_w be a representative set of $\mathcal{Q}_w = \mathbb{Z}_w^* / \langle 2 \rangle$, and fix the representative set of $\mathcal{Q}_m = \mathbb{Z}_m^* / \langle 2 \rangle$ to be $U_m = \{j \in \mathbb{Z}_m^* : \exists\, i \in U_w \text{ s.t. } j \equiv i \pmod{w}\}$. Denote the order of 2 (in both $\mathbb{Z}_m^*$ and $\mathbb{Z}_w^*$) by d, let $\rho \in \mathbb{F}_{2^d}$ be a primitive m'th root of unity, and fix the particular primitive w'th root of unity $\tau = \rho^u$.*

Finally, fix an arbitrary value $\alpha \in \mathbb{F}_{2^d}$ and let $a(X)$ be the (unique) polynomial in $\mathbb{F}_2[X]/\Phi_m(X)$ that satisfies $a(\rho^j) = \alpha$ for all $j \in U_m$.

Then a is of the form $a(X) = b(X^u) \bmod (\Phi_m(X), 2)$ for some polynomial $b(X) \in \mathbb{F}_2[X]/\Phi_2(X)$ satisfying $b(\tau^i) = \alpha$ for all $i \in U_w$. In particular a, b represent the same element $r' \in R_2' \subset R_2$.

Proof. We first note that a polynomial $a(X)$ as above is indeed unique, due to Claim 1. Similarly a polynomial $b(X) \in \mathbb{F}_2[X]/\Phi_2(X)$ satisfying $b(\tau^i) = \alpha$ for all $i \in U_w$ is also uniquely determined. Denoting $c(X) \overset{\text{def}}{=} b(X^u) \bmod (\Phi_m(X), 2) \in R_2$, it is only left to show that $c(X) = a(X)$.

Clearly both c and a are polynomials in $R_2 \cong \mathbb{F}_2[X]/\Phi_m(X) \subset \mathbb{F}_{2^d}[X]/\Phi_m(X)$, so it is sufficient to show that they agree when evaluated on $\rho^j \in \mathbb{F}_{2^d}$ for all $j \in U_m$ (again by Claim 1). By definition of U_m, for every $j \in U_m$ there exists $i \in U_w$ such that $j \equiv i \pmod{w}$, hence we get

$$c(\rho^j) \;=\; b(\rho^{u \cdot j}) \;=\; b(\tau^j) \;=\; b(\tau^i) \;=\; \alpha \;=\; a(\rho^j).$$

□

(We note that the fact that 2 has the same order modulo w and m is used in the assertion that the set U_m as above is a representative set for $\mathcal{Q}_m$.)

Corollary 1. *With notations as in Lemma 5 and the transformation $T\colon R_2 \to (R_2')^{u'}$ from Section 2.2, if a, b are as in Lemma 5 then for any element $x \in R_2$ we have $T_2(a \cdot x \in R_2) = b \cdot T(x) \in (R_2')^{u'}$.*

Proof. Follows immediately from the R_2'-linearity of T_2 and the fact that the polynomials $a \in \mathbb{F}_2[X]/\Phi_m(X)$ and $b \in \mathbb{F}_2[X]/\Phi_w(X)$ represent the same element $r' \in R_2' \subseteq R_2$ (since $a(X) = b(X^u) \bmod (\Phi_m(X), 2)$).
□

Given Corollary 1, the rest of the proof follows quite easily. Consider now the encoding functions that map R_2 elements into the vector of $\mathbb{F}_{2^d}$ values that are encoded in all their slots. Namely, for $a \in R_2$ denote by $\mathsf{Enc}_m(a) \in \mathbb{F}_{2^d}^{\ell}$ the vector of values $a(\rho^j)$ for $j \in U_m$. Similarly consider the encoding of a vector of R'_2 elements into the $\mathbb{F}_{2^d}$ values that are encoded in all the slots of all the elements. That is, for a vector $\vec{a} = (a_0, a_1, \ldots, a_{u'-1}) \in (R'_2)^{u'}$, denote by $\mathsf{Enc}_w(\vec{a})$ the vector of values $a_k(\tau^i)$ for $i \in U_w$ and $k \in [u']$. We note that the dimensions of $\mathsf{Enc}_m(a)$ and $\mathsf{Enc}_w(T_2(a))$ are the same, namely they both have dimension $\ell = \varphi(m)/d = u' \cdot \varphi(w)/d$.

Lemma 6. *There exists an invertible linear transformation L over $\mathbb{F}_{2^d}$ such that for any $a \in R_2$ it holds that $\mathsf{Enc}_w(T_2(a)) = L(\ \mathsf{Enc}_m(a)\ )$.*

Proof. Recalling that the encoding functions are bijections (by Claim 1), we thus define $L(\vec{x}) \stackrel{\text{def}}{=} \mathsf{Enc}_w(T_2(\mathsf{Enc}_m^{-1}(\vec{x})))$, and note that L must be invertible, because $T_2(\cdot)$ is also a bijection.

It remains only to show that L is $\mathbb{F}_{2^d}$-linear. The property $L(\vec{x}) + L(\vec{y}) = L(\vec{x} + \vec{y})$ follows immediately from the facts that the same property holds for each of $T_2(\cdot)$, $\mathsf{Enc}_m(\cdot)$, and $\mathsf{Enc}_w(\cdot)$. We next use Lemma 5 and Corollary 1 to show the property $L(\alpha \cdot \vec{x}) = \alpha \cdot L(\vec{x})$.

Fix a vector $\vec{x} \in \mathbb{F}_{2^d}^{\ell}$ and a value $\alpha \in \mathbb{F}_{2^d}$, and let $a \in \mathbb{F}_2[X]/\Phi_m(X) \cong R_2$ be the element that has α in all of its plaintext slots, $a = \mathsf{Enc}_m^{-1}(\alpha^{\ell})$. Similarly let $x = \mathsf{Enc}_m^{-1}(\vec{x})$. Observe that since multiplication in R_2 implies pointwise multiplication on the slots, then the product $a \cdot x \in R_2$ encodes in its slots exactly α times the slots of x. In other words, we have $\mathsf{Enc}_m^{-1}(\alpha \cdot \vec{x} \in \mathbb{F}_{2^d}^{\ell}) = a \cdot x \in R_2$.

Since a has the same element α in all its slots then it satisfies the condition of Lemma 5 and Corollary 1. Let $b \in \mathbb{F}_2[X]/\Phi_w(X) \cong R'_2$ be the polynomial promised by Lemma 5. Then from Corollary 1 we have that $T_2(a \cdot x) = b \cdot T(x)$. Moreover Lemma 5 tells us that b also have the values α in all its slots. Since multiplication in R'_2 also implies pointwise multiplication on the slots, i.e., $\mathsf{Enc}_w(b \cdot \vec{y}) \ = \ \alpha \cdot \mathsf{Enc}_w(\vec{y})$ for every $\vec{y} \in (R'_2)^{u'}$. In particular,

$$\mathsf{Enc}_w(T_2(a \cdot x)) \ = \ \mathsf{Enc}_w(b \cdot T_2(x)) \ = \ \alpha \cdot \mathsf{Enc}_w(x),$$

or in other words $L(\alpha \cdot \vec{x}) = \alpha \cdot L(\vec{x})$, as needed. □

Our strategy for recovering the original values in the plaintext space after ring-switching is to first use the transformation T to break a big-ring ciphertexts into a collection of small-ring ciphertexts. By Lemma 6 this operation has the side effect of transforming the slots according to the invertible $\mathbb{F}_{2^d}$-linear transformation L, so we compute homomorphically the inverse transformation L^{-1} on the slots, using the tools from [11] for computing on packed ciphertexts.

If we only need a few of the slots in a (as in the AES example), then we can compute only the relevant rows of L^{-1}, thereby getting at the end of the process only as many small-ring ciphertexts as required to encode all the plaintext slots that we are interested in.

Remarks. Note that the *only* properties of T that we used in this work are the properties 1-4 that were described in the introduction. Namely, all we need is a transformation $T\colon R \to (R')^*$ which is injective and R'_2-linear, and that maps small R elements into small R' vectors. There could be many such transformations, and they could offer different tradeoffs in practice. (For example, the transformation in a previous version of this work [10], which was based on the coefficient-splitting technique from [3], turns out to include a very sparse linear transformations L, making the homomorphic computation of L^{-1} at the end must easier.) Also, as mentioned in Section 2.2, in some cases we can use a K'_2-linear transformation $T\colon \mathbb{K} \to (\mathbb{K}')^*$ even if it does not map R-elements to R'-vectors.

We also note that Lemma 5 (and consequently Corollary 1 and Lemma 6) can be extended also to the case where the order of 2 modulo w is smaller than its order modulo m, as long as we only consider elements

$a \in R_2$ that have values from the smaller field $\mathbb{F}_{2^{d'}}$ in their plaintext slots. We defer details of this extension to the full version.

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 Lòpez-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.