



Faster Homomorphic Linear Transformations in HELib

Shai Halevi¹(✉) and Victor Shoup^{1,2}

¹ IBM Research, Yorktown Heights, NY, USA
shaih@alum.mit.edu

² New York University, New York, USA
shoup@cs.nyu.edu

Abstract. HELib is a software library that implements homomorphic encryption (HE), with a focus on effective use of “packed” ciphertexts. An important operation is applying a known linear map to a vector of encrypted data. In this paper, we describe several algorithmic improvements that significantly speed up this operation: in our experiments, our new algorithms are 30–75 times faster than those previously implemented in HELib for typical parameters.

One application that can benefit from faster linear transformations is bootstrapping (in particular, “thin bootstrapping” as described in [Chen and Han, Eurocrypt 2018]). In some settings, our new algorithms for linear transformations result in a $6\times$ speedup for the entire thin bootstrapping operation.

Our techniques also reduce the size of the large public evaluation key, often using 33%–50% less space than the previous HELib implementation. We also implemented a new tradeoff that enables a drastic reduction in size, resulting in a $25\times$ factor or more for some parameters, paying only a penalty of a $2\text{--}4\times$ times slowdown in running time (and giving up some parallelization opportunities).

Keywords: Homomorphic encryption · Implementation
Linear transformations

1 Introduction

Homomorphic encryption (HE) [5, 13] enables performing arithmetic operations on encrypted data even without knowing the secret key. All contemporary HE schemes roughly follow the outline of Gentry’s first candidate, where fresh ciphertexts are “noisy” to ensure security. This noise grows with every operation, until it becomes so large so as to cause decryption errors. This results in a “somewhat homomorphic” encryption scheme (SWHE) that can only evaluate low-depth circuits, such a scheme can be converted to a “fully homomorphic”

Supported by the Defense Advanced Research Projects Agency (DARPA) and Army Research Office (ARO) under Contract No. W911NF-15-C-0236.

encryption scheme (FHE) using bootstrapping. The most asymptotically efficient SWHE schemes are based on the hardness of ring-LWE. Most of these schemes use $R_p = \mathbb{Z}[X]/(F(X), p)$ as their native plaintext space, with F a cyclotomic polynomial and p an integer (usually a prime or prime power).

Smart and Vercauteren observed [15] that (for a prime p) an element in this native plaintext space can be used to encode (via Chinese Remaindering) a vector of values from a finite field \mathbb{F}_{p^d} , for some integer d that depends on F and p , and that operations on elements in R_p induce the corresponding entry-wise operation on the encoded vectors. This technique of encoding many plaintext elements from \mathbb{F}_{p^d} in a single R_p element, which is then encrypted and manipulated homomorphically, is called “ciphertext packing”, and the entries in the vector are called “plaintext slots.” Gentry, Halevi, and Smart showed in [6] how to use special automorphisms on R_p (which were used for different purposes in [2, 10]) to enable data movement between the slots.

HElib [7–9] is an open-source C++ library that implements the ring variant of the scheme due to Brakerski-Gentry-Vaikuntanathan [2], focusing on effective use of ciphertext packing. It includes an implementation of the BGV scheme itself with all its basic homomorphic operations, as well as higher-level procedures for data-movement, simple linear algebra, bootstrapping, etc. One can think of the lower levels of HElib as providing a “hardware platform”, defining a set of operations that can be applied homomorphically. These operations include entry-wise addition and multiplication operations on the vector of plaintext values, as well as data movement, making this “platform” a SIMD environment.

Our Results. In this work, we improve performance of core linear algebra algorithms in HElib that apply publicly known linear transformations to encrypted vectors. These improvements are now integrated into HElib. For typical, realistic parameter settings, our new algorithms can run 30–75 times faster than those in the previous implementation of HElib, where the exact speedup depends on myriad details.¹ Our implementation also exploits multiple cores, when available, to get even further speedups.

Our techniques also reduce the size of the large public evaluation key. In the old HElib implementation, the evaluation key typically consists of a large number of large “key switching matrices”: Each of these “matrices” can take 1–4 MB of space, and the implementation uses close to a hundred of them. Our new implementation reduces the number of key-switching matrices by 33–50% in some parameter settings (that arise fairly often in practice), while at the same time improves the running time. Moreover, a new tradeoff that we implemented enables a drastic reduction in the number of matrices (sometimes as few as four or six matrices overall), for a small price of only 2–4× in performance. This space efficient variation, however, is inherently sequential, as opposed to our other procedure than can be easily parallelized.

¹ One could also consider algorithms that apply encrypted linear transformations to encrypted vectors; some of our new algorithmic techniques may apply to that problem as well; however, we have not yet implemented this in HElib.

Applications. Linear transformations of encrypted vectors is a manifestly fundamental operation with many applications. For one example, HELib itself makes critical use of such transformations in its bootstrapping logic. As reported in [8], the bootstrapping routine can typically spend 25–40% of its time performing such transformations. In addition, a new “thin bootstrapping” technique, due to Chen and Han [4], is useful to bootstrap encrypted vectors whose entries are in the base field, rather than an extension field. In practice, this is an important special case of bootstrapping, and our faster algorithms for linear transformations play an even more significant role here. Our timing results in Sect. 9 show that for large vectors, these faster algorithms are essential to make “thin bootstrapping” practical.

As another example, consider a *private information retrieval* protocol in which a client selects one value from a database of values held by a server, while hiding from the server which value was accessed. Using HE, one way to do this is for the server to encode each value as a column vector. The collection of all such values held by the server is thus encoded as a matrix M , where each column in M corresponds to one value. To access the i th value, the client can send to the server an encrypted unit vector v with 1 in the i th entry (or some other encrypted information from which the server can homomorphically compute such an encrypted unit vector). The server then homomorphically computes $M \times v$, which is an encryption of the selected column of M . The server sends the result to the client, who can decrypt it and recover the selected value.

Techniques. In the linear transformation algorithms previously implemented in HELib, the bulk of the time is spent moving data among the slots in the encrypted vector. As mentioned above, this is accomplished by using special automorphisms. The main cost of applying such an automorphism to a ciphertext is actually that of “key switching”: after applying the automorphism to each ring element in the ciphertext (which is actually a very cheap operation), we end up with an encryption relative to the “wrong” secret key; we can recover a ciphertext relative to the “right” secret key by using data in the public key specific to this particular automorphism — a so-called “key switching matrix.”

The main goals in improving performance are therefore to reduce the number of automorphisms, and to reduce the cost of each automorphism.

- To reduce the number of automorphisms, we introduce a “baby-step/giant-step” strategy for computing all of the required automorphisms. This strategy generalizes a similar idea that was used in [8] in the context of bootstrapping. This strategy by itself speeds up the computation by a factor of 15–20 in typical settings. See Sect. 4.1.
- We further reduce the number of automorphisms by refactoring a number of computations, more aggressively exploiting the algebraic properties of the automorphisms that we use. See Sect. 4.4.

- To reduce the cost of each automorphism, we introduce a new technique for “hoisting” the expensive parts of these operations out of the main loop.² Our main observation is that applying many automorphisms to the same ciphertext v can be done faster than applying each one separately. Instead, we can perform an expensive pre-computation that depends only on v (but not the automorphisms themselves), and this pre-computation makes each automorphism much cheaper (typically, 6–8 times faster). See Sects. 4.2 and 5.
- Recall that key switching matrices are a part of the public key and consume quite a lot of space (typically several megabytes per matrix), so keeping their numbers down is desirable. In the previous implementation of HElib, there can easily be several hundred such matrices in the public key. We introduce a new technique that reduces the number of key-switching matrices by 33–50% in some parameter settings (that arise fairly often in practice), while at the same time improves the running time of our algorithms. See Sect. 4.3.
- We introduce yet another technique that drastically reduces the number of key-switching matrices to a very small number (less than 10), but comes at a cost in running time (typically 2–4 times more slowly as our fastest algorithms), and cannot be parallelized.³ Achieving this reduction in key-switching storage without too much degradation in running time requires some new algorithmic ideas. See Sect. 4.5.

Outline. The rest of the paper is organized as follows.

- In Sect. 2, we introduce notation and terminology, and review the basics of the BGV cryptosystem, including ciphertext packing and automorphisms.
- In Sect. 3, we review the basic ideas underlying the previous algorithms in HElib for applying linear transformations homomorphically. We focus on restricted linear transformations, the “one-dimensional” transformations MatMul1D and BlockMatMul1D. It turns out that considering these restricted transformations is sufficient: they can be used directly in applications such as bootstrapping, and can be easily be used to implement more general linear transformations.
- In Sect. 4, we give a more detailed overview of our new techniques.
- In Sect. 5, we give more of the details of our new hoisting technique.
- In Sect. 6, we present all of our new algorithms for MatMul1D and BlockMatMul1D in detail.
- In Sect. 7, we describe how to use algorithms for MatMul1D and BlockMatMul1D for more general linear transformations.

² “Hoisting” is a term used in compiler optimization to describe the action of “hoisting” a computation out of a loop, so that it is only performed once, instead of in every loop iteration.

³ While the “top level” operations in our linear transformations are inherently sequential when using this technique, lower-level routines in HElib will still exploit multiple cores, if available. Such low-level parallelism are usually less effective, however.

- In Sect. 8, we review the bootstrapping procedure from [8], and discuss how those techniques can be adapted to the “thin bootstrapping” technique of Chen and Han [4].
- In Sect. 9, we report on the performance of the implementation of our new algorithms (and their application to bootstrapping).

2 Notations and Background

For a positive modulus $q \in \mathbb{Z}_{>0}$, we identify the ring \mathbb{Z}_q with its representation as integers in $[-q/2, q/2)$ (except for $q = 2$ where we use $\{0, 1\}$). For integer z , we denote by $[z]_q$ the reduction of z modulo q into the same interval. This notation extends to vectors and matrices coordinate-wise, and to elements of other algebraic groups/rings/fields by considering their coefficients in some convenient basis (e.g., the coefficient of polynomials in the power basis when talking about $\mathbb{Z}[X]$). The norm of a ring element $\|a\|$ is defined as the norm of its coefficient vector in that basis.⁴

2.1 The BGV Cryptosystem

The BGV ring-LWE-based scheme [3] is defined over a ring $R \stackrel{\text{def}}{=} \mathbb{Z}[X]/(\Phi_m(X))$, where $\Phi_m(X)$ is the m th cyclotomic polynomial. For an arbitrary integer modulus N (not necessarily prime) we denote the ring $R_N \stackrel{\text{def}}{=} R/NR$.

As implemented in HELib, the native plaintext space of the BGV cryptosystem is R_{p^r} for a prime power p^r . The scheme is parametrized by a sequence of decreasing moduli $q_L \gg q_{L-1} \gg \dots \gg q_0$, and an “ i th level ciphertext” in the scheme is a vector $v \in R_{q_i}^2$. Secret keys are elements $s \in R$ with “small” coefficients (chosen in $\{0, \pm 1\}$ in HELib), and we view s as the second element of the 2-vector $\text{sk} = (1, s) \in R^2$. A level- i ciphertext $v = (p_0, p_1)$ encrypts a plaintext element $\alpha \in R_{p^r}$ with respect to $\text{sk} = (1, s)$ if $[(\text{sk}, v)]_{q_i} = [p_0 + s \cdot p_1]_{q_i} = \alpha + p^r \cdot \epsilon$ (in R) for some “small” error term, $\|\epsilon\| \ll q_i/p^r$.

The error term grows with homomorphic operations of the cryptosystem, and switching from q_{i+1} to q_i is used to decrease the error term roughly by the ratio q_{i+1}/q_i . Once we have a level-0 ciphertext v , we can no longer use that technique to reduce the noise. To enable further computation, we need to use Gentry’s bootstrapping technique [5]. In HELib, each q_i is a product of small (machine-word sized) primes.

2.2 Encoding Vectors in Plaintext Slots

As observed by Smart and Vercauteren [15], an element of the native plaintext space $\alpha \in R_{p^r}$ can be viewed as encoding a vector of “plaintext slots” containing elements from some smaller ring extension of \mathbb{Z}_{p^r} via Chinese remaindering.

⁴ The difference between the norm in the different bases is not very important for the current work.

In this way, a single arithmetic operation on α corresponds to the same operation applied component-wise to all the slots.

Specifically, suppose the factorization of $\Phi_m(X)$ modulo p^r is $\Phi_m(X) \equiv F_1(X) \cdots F_\ell(X) \pmod{p^r}$, where each F_i has the same degree d , which is equal to the order of p modulo m , so that $\ell = \phi(m)/d$. (This factorization can be obtained by factoring $\Phi_m(X)$ modulo p , followed by Hensel lifting.) Then we have the isomorphism $R_{p^r} \cong \bigoplus_{i=1}^{\ell} (\mathbb{Z}[X]/(p^r, F_i(X)))$.

Let us now denote $E \stackrel{\text{def}}{=} \mathbb{Z}[X]/(p^r, F_1(X))$, and let ζ be the residue class of X in E , which is a principal m th root of unity, so that $E = \mathbb{Z}/(p^r)[\zeta]$. The rings $\mathbb{Z}[X]/(p^r, F_i(X))$ for $i = 1, \dots, \ell$ are all isomorphic to E , and their direct product is isomorphic to R_{p^r} , so we get an isomorphism between R_{p^r} and E^ℓ . HELib makes extensive use of this isomorphism, using it to encode an ℓ -vector of elements in E as an element of the native plaintext space R_{p^r} . Addition and multiplication of ciphertexts act on all ℓ slots of the corresponding plaintext in parallel.

2.3 Hypercube Structure and One-Dimensional Rotations

Beyond addition and multiplications, we can also manipulate elements in R_{p^r} using a set of automorphisms on R_{p^r} of the form

$$\theta_t : R_{p^r} \longrightarrow R_{p^r}, \quad a(X) \longmapsto a(X^t) \pmod{(p^r, \Phi_m(X))}.$$

for $t \in \mathbb{Z}_m^*$. Since each θ_t is an automorphism, it distributes over addition and multiplication, i.e., $\theta_t(\alpha + \beta) = \theta_t(\alpha) + \theta_t(\beta)$ and $\theta_t(\alpha\beta) = \theta_t(\alpha)\theta_t(\beta)$. Also, these automorphisms commute with one another, i.e., $\theta_t\theta_{t'} = \theta_{t't} = \theta_{t'}\theta_t$. Moreover, for any integer i , we have $\theta_t^i = \theta_{t^i}$.

We can homomorphically apply such an automorphism by applying it to the individual ciphertext components and then performing “key switching” (see [3, 6]). In somewhat more detail, a ciphertext in HELib consists of two “parts,” each an element of R_q for some q . Applying the same automorphism (defined in R_q) to the two parts, we get a ciphertext with respect to a different secret key. In order to do anything more with this ciphertext, we usually have to convert it back to a ciphertext with respect to the original secret key. In order to do this, the public-key must contain data specific to the automorphism θ_t , called a “key switching matrix”.⁵ We will discuss this key-switching operation in more detail below in Sect. 5.

As discussed in [6], these automorphisms induce a hypercube structure on the plaintext slots, that depends on the structure of the group $\mathbb{Z}_m^*/\langle p \rangle$. Specifically, HELib keeps a hypercube basis $g_1, \dots, g_n \in \mathbb{Z}_m^*$ with orders $D_1, \dots, D_n \in \mathbb{Z}_{>0}$, and then defines the set of representatives for $\mathbb{Z}_m^*/\langle p \rangle$ as

$$\{g_1^{e_1} \cdots g_n^{e_n} : 0 \leq e_s < D_s, \quad s = 1, \dots, n\}.$$

⁵ Note that this “key switching” technique is a generalization of that used to allow multiplication of ciphertexts.

More precisely, D_s is the order of g_s in $\mathbb{Z}_m^*/\langle p, g_1, \dots, g_{s-1} \rangle$. Thus, the slots are in one-to-one correspondence with tuples (e_1, \dots, e_n) with $0 \leq e_s < D_s$. This induces an n -dimensional hypercube structure on the plaintext space. If we fix $e_1, \dots, e_{s-1}, e_{s+1}, \dots, e_n$, and let e_s range over $0, \dots, D_s - 1$, we get a set of D_s slots, which we refer to as a *hypercolumn in dimension s* (and there are ℓ/D_s such hypercolumns).

Using automorphisms, we can efficiently perform rotations in any dimension; a *rotation by i in dimension s* maps a slot corresponding to $(e_1, \dots, e_s, \dots, e_n)$ to the slot corresponding to $(e_1, \dots, e_s + i \bmod D_s, \dots, e_n)$. In other words, it rotates each hypercolumn in dimension s by i . We denote by ρ_s the rotation-by-1 operation in dimension s . Observe that ρ_s^i is the rotation-by- i operation in dimension s .

We can implement ρ_s^i by applying either one or two of the automorphisms $\{\theta_t\}_{t \in \mathbb{Z}_m^*}$ defined above. If the order of g_s in \mathbb{Z}_m^* is D_s , then we get by with just a single automorphism, since

$$\rho_s^i(\alpha) = \theta_{g_s^i}(\alpha). \quad (1)$$

In this case, we call s a “good dimension”.

If the order of g_s in \mathbb{Z}_m^* is different from D_s , then we call s a “bad dimension”, and we need to implement this rotation using two automorphisms. Specifically, we use a constant “0–1 mask value” μ that selects some slots and zeros-out the others, and use the two automorphisms $\psi \stackrel{\text{def}}{=} \theta_{g_s^i}$ and $\psi^* \stackrel{\text{def}}{=} \theta_{g_s^{i-D}}$. Then we have

$$\rho_s^i(\alpha) = \psi(\mu \cdot \alpha) + \psi^*((1 - \mu) \cdot \alpha). \quad (2)$$

The idea is roughly as follows. Even though ψ does not act as a rotation by i in dimension s , it does act as the desired rotation if we restrict it to inputs with zeros in each slot whose coordinate in dimension s is at least $D - i$. Similarly, ψ^* acts as the desired rotation if we restrict it to inputs with zeros in each slot whose coordinate in dimension s is less than $D - i$. This tells us that μ should have a 1 in all slots whose coordinate in dimension s is less than $D - i$, and a 0 in all other slots. Note also that

$$\rho_s^i(\alpha) = \mu' \cdot \psi(\alpha) + (1 - \mu') \cdot \psi^*(\alpha), \quad (3)$$

where $\mu' = \psi(\mu)$ is a mask with a 1 in all slots whose coordinate in dimension s is at least i , and a 0 in all other slots. This formulation will be convenient in some of the algorithms we present.

2.4 Frobenius and Linearized Polynomials

We define the automorphism $\sigma \stackrel{\text{def}}{=} \theta_p$, which is the Frobenius map on R_{p^r} (where θ_p is one of the automorphisms defined in Sect. 2.3). It acts on each slot independently as the Frobenius map σ_E on E , which sends ζ to ζ^p and leaves elements of \mathbb{Z}_{p^r} fixed. (When $r = 1$, σ is the same as the p th power map on E .)

For any \mathbb{Z}_p -linear transformation on E , denoted M , there exist unique constants $\lambda_0, \dots, \lambda_{d-1} \in E$ such that $M(\eta) = \sum_{j=0}^{d-1} \lambda_j \sigma_E^j(\eta)$ for all $\eta \in E$. When $r = 1$, this follows from the general theory of linearized polynomials (see, e.g., Theorem 10.4.4 on p. 237 of [14]), but the same results are easily seen to hold for $r > 1$ as well. These constants are readily computable by solving a system of equations mod p^r .

Using linearized polynomials, we may effectively apply a fixed linear map to each slot of a plaintext element $\alpha \in R_{p^r}$ (either the same or different maps in each slot) by computing $\sum_{j=0}^{d-1} \kappa_j \sigma^j(\alpha)$, where the κ_j 's are R_{p^r} -constants obtained by embedding appropriate E -constants in the slots.

2.5 Key Switching Strategies

The total number of automorphisms is $\phi(m)$, which is typically many thousands, so it is not very practical to store all possible key switching matrices in the public key: each such matrix typically occupies a few megabytes of storage, and storing all of them will consume hundreds of gigabytes. Therefore, we consider strategies that trade off space for time with respect to key switching matrices.

For almost all applications, we only need the key switching matrices for one-dimensional rotations in each dimension, as well as for the Frobenius map (and its powers). For a fixed dimension $s = 1, \dots, n$ of size $D \stackrel{\text{def}}{=} D_s$ with generator $g \stackrel{\text{def}}{=} g_s$, consider the automorphism $\theta \stackrel{\text{def}}{=} \theta_{g_s}$. In the original implementation of HElib, one of two key switching strategies for dimension s are used.

Full: We store key switching matrices for θ^i for $i = 0, \dots, D - 1$. If s is a “bad dimension”, we additionally store key switching matrices for θ^{-i} for $i = 1, \dots, D - 1$.

Baby-step/giant-step: We store key switching matrices for θ^j with $j = 1, \dots, g - 1$, where $g \stackrel{\text{def}}{=} \lceil \sqrt{D} \rceil$ (the “baby steps”), as well as for θ^{gk} with $k = 1, \dots, h - 1$, where $h \stackrel{\text{def}}{=} \lceil D/g \rceil$ (the “giant steps”). If s is a “bad dimension”, we additionally store key switching matrices for θ^{-gk} with $k = 1, \dots, h$ (negative “giant steps”).

Using the full strategy, any rotation in dimension s can be implemented using a single automorphism and key switching if s is a good dimension, and using two automorphisms and key switchings if s is a bad dimension.

Using the baby-step/giant-step strategy, any rotation in dimension s can be implemented using at most two automorphisms and key switchings if s is a good dimension, and using at most four automorphisms and key switchings if s is a bad dimension. The idea is that to compute $\theta^i(v)$, for a given $i = 0, \dots, D - 1$, we can write $i = j + gk$, so that to compute $\theta^i(v)$, we first compute $w = \theta^{gk}(v)$, which takes one automorphism and a key switching, and then compute $\theta^j(w)$, which takes another automorphism and key switching.

These two strategies give us a time/space trade-off: although it slows down the computation time by a factor of two, the baby-step/giant-step strategy

requires space for just $O(\sqrt{D})$ key switching matrices, rather than the $O(D)$ key switching matrices required by the full strategy.

The same two strategies can be used to store key switching matrices for powers of the Frobenius map, so that any power of the Frobenius map can be computed using either one or two automorphisms. Indeed, it is convenient to think of the powers of the Frobenius map as defining an additional (effectively “good”) dimension.

The default behavior of `HElib` is to use the full key-switching strategy for “small” dimensions (of size at most 50), and the baby-step/giant-step strategy for larger dimensions.

3 Matrix Multiplication — Basic Ideas

In [7], it is observed that we can multiply a matrix $M \in E^{\ell \times \ell}$ by a column vector $v \in E^{\ell \times 1}$ by computing

$$Mv = M_0v_0 + \cdots + M_{\ell-1}v_{\ell-1}, \quad (4)$$

where each v_i is the vector obtained by rotating the entries of v by i positions, and each M_i is a diagonal matrix containing one diagonal of M .

3.1 `MatMul1D`: One-Dimensional E -Linear Transformations

In many applications, such as the decryption procedure in [8], instead of a general E -linear transformation on R_{p^r} , we only need to work with a *one-dimensional* E -linear transformation that acts independently on the individual hypercolumns of a single dimension $s = 1, \dots, n$. We can adapt the diagonal decomposition of Eq. (4) to this setting using appropriate rotation maps on the slots of R_{p^r} . Let $\rho \stackrel{\text{def}}{=} \rho_s$ be the rotation-by-1 map in dimension s , and let $D \stackrel{\text{def}}{=} D_s$ be the size of dimension s . If T is a one-dimensional E -linear transformation on R_{p^r} , then for every $v \in R_{p^r}$, we have

$$T(v) = \sum_{i=0}^{D-1} \kappa_i \cdot \rho^i(v), \quad (5)$$

where the κ_i 's are constants in R_{p^r} determined by T , obtained by embedding appropriate constants in E in each slot. Equation (5) translates directly into a simple homomorphic evaluation algorithm, just by applying the same operations to a ciphertext encrypting v . In a straightforward implementation, in a good dimension, the computational cost is about D automorphisms and D constant-ciphertext multiplications, and the noise cost is a single constant-ciphertext multiplication. In bad dimensions, all of these costs would essentially double. In practice, if the constants have been pre-computed, the computation cost of the constant-ciphertext multiplications is negligible compared to that of the automorphisms.

One of our main goals in this paper is to dramatically improve upon the computational cost for performing such a `MatMul1D` operation.

3.2 BlockMatMul1D: One-Dimensional \mathbb{Z}_{p^r} -Linear Transformations

In some applications (again, including the reryption procedure in [8]), instead of applying an E -linear transformation, we need to apply a \mathbb{Z}_{p^r} -linear map. Again, we focus on *one-dimensional* \mathbb{Z}_{p^r} -linear maps that act independently on the hypercolumns of a single dimension.

We can still use the same diagonal decomposition as in Eq. (4), except that the entries in the diagonal matrices are no longer elements of E , but rather, \mathbb{Z}_{p^r} -linear maps on E . These maps may be encoded using linearized polynomials, as in Sect. 2.4. Therefore, if T is a one-dimensional \mathbb{Z}_{p^r} -linear transformation on R_{p^r} , then for every $v \in R_{p^r}$, we have

$$T(v) = \sum_{i=0}^{D-1} \sum_{j=0}^{d-1} \kappa_{i,j} \cdot \sigma^j(\rho^i(v)), \quad (6)$$

where the $\kappa_{i,j}$'s are constants in R_{p^r} determined by T .

A naive homomorphic implementation of the formula from Eq. (6) takes $O(dD)$ automorphisms, but as shown in [8], this can be reduced to $O(d + D)$ automorphisms. In this paper, we will also present significant improvements to the BlockMatMul1D algorithm in [8], although they are not as dramatic as our improvements to the MatMul1D algorithm.

4 Overview of Algorithmic Improvements

4.1 Baby-Step/Giant-Step Multiplication

As already mentioned, [8] introduces a technique that reduces the number of automorphisms needed to implement BlockMatMul1D in dimension s from $O(dD)$ to $O(d + D)$, where $D \stackrel{\text{def}}{=} D_s$ is the size of the dimension, and d is the order of p mod m . A very similar idea, essentially a baby-step/giant-step technique, can be used to reduce the number of automorphisms needed to implement MatMul1D in dimension s from $O(D)$ to $O(\sqrt{D})$. See Sect. 6 for details.

This technique is distinct from the baby-step/giant-step key switching strategy discussed above in Sect. 2.5. However, for best results, the two techniques should be combined in a way that harmonizes the baby-step/giant-step thresholds.

4.2 Hoisting

As we have seen, in many situations, we want to compute $\psi(v)$ for a fixed ciphertext v and many automorphisms ψ . Assuming we have key switching matrices for each automorphism ψ , the dominant cost of computing all of these values is that of performing one key-switching operation for each ψ . Our “hoisting” technique is a method that refactors the computation, performing a pre-computation that only depends on v , and whose computational cost is roughly equivalent to a single key-switching operation. After performing this pre-computation, computing

$\psi(v)$ for any individual ψ is much faster than a single key-switching operation (typically, around 6–8 times faster). We describe this idea in more detail below in Sect. 5.

4.3 Better Key Switching Strategies in Bad Dimensions

Recall from Sect. 2.5 that with the “full” key-switching strategy, in a bad dimension, we stored key-switching matrices for the automorphisms θ^i , with $i = -(D - 1), \dots, -1, 1, \dots, D - 1$. To perform a rotation by i on v in the given dimension, we need to compute $\theta^i(v)$ and $\theta^{i-D}(v)$, and so with these key-switching matrices available, we need to perform two automorphisms and key switchings. However, we do not really need all of these negative-power key switching matrices. In fact, we can get by with key-switching matrices just for θ^i , with $i = 1, \dots, D - 1$, and for θ^{-D} . To perform a rotation by i on v in the given dimension, we can compute $w = \theta^i(v)$ and $\theta^{-D}(w) = \theta^{i-D}(v)$. So again, we need to perform two automorphisms and key switchings. This cuts the number of key-switching matrices in half without a significant increase in running time. Moreover, this key-switching strategy aligns well with the strategy discussed below for decoupling rotations and automorphisms in bad dimensions.

Similarly, for the baby-step/giant-step key-switching strategy in a bad dimension, we just store a key-switching matrix for θ^{-D} , rather than for all the negative “giant steps”. This cuts down the number of key-switching matrices by a third. Moreover, the number of key switchings we need to perform per rotation is only 3 (instead of 4).

4.4 Decoupling Rotations and Automorphisms in Bad Dimensions

Recall that by Eq. (3), a rotation by i on a ciphertext v in a given bad dimension can be implemented as $\mu\theta^i(v) + (1 - \mu)\theta^{i-D}(v)$, where μ is a “mask” (a constant with a 0 or 1 encoded in each slot). It turns out that in our matrix-vector computations, it is best to work directly with this implementation, and algebraically refactor the computation to improve both running time and noise. This refactoring exploits the fact that θ is an automorphism. See Sect. 6 for details.

4.5 A Horner-Like Rule with Application to a Minimal Key-Switching Strategy

We introduce a new key-switching strategy that reduces the storage requirements even further, to just 1, 2, or 3 key-switching matrices per dimension. This, combined with a simple algorithmic idea, allow us to implement a variant of the baby-step/giant-step multiplication strategy that does not run too much more slowly than when using the full or baby-step/giant-step key-switching strategy. To do this, we observe that if we need to compute $\sum_{i=0}^{h-1} \psi^i(v_i)$, where ψ is some automorphism and the v_i ’s are ciphertexts, we can do this using Horner’s rule, provided we have a key-switching matrix just for ψ . Specifically, we can compute

$$\sum_{i=0}^{h-1} \psi^i(v_i) = \psi(\cdots \psi(\psi(v_{h-1}) + v_{h-2}) + \cdots) + v_0.$$

That is, we set $w_{h-1} \leftarrow v_{h-1}$, then $w_{i-1} \leftarrow \psi(w_i) + v_{h-1}$ for $i = h-1, \dots, 1$, and finally we output w_0 .

4.6 Exploiting Multi-core Platforms

With the exception of the minimal key-switching strategy discussed above, all our other algorithms are very amenable to parallelization. We thus implemented them so as to exploit multiple cores, when available.

5 Hoisting

A ciphertext in `HElib` is a vector $v = (p_0, p_1) \in R_q^2$, with each “part” p_0, p_1 represented in a `DoubleCRT` format (i.e., both integer and polynomial CRT) [9]. We recall the steps in the computation of each $\psi(v)$, as implemented in `HElib`.

1. **Automorphism:** We first apply the automorphism to each part of v , computing $p'_j \leftarrow \psi(p_j)$ for $j = 0, 1$.

Applying an automorphism to a `DoubleCRT` object is a fast, linear time operation, so this step is cheap. If $v = (p_0, p_1)$ decrypts to α under the secret key $\text{sk} = (1, s)$, then $v' = (p'_0, p'_1)$ decrypts to $\psi(\alpha)$ under the secret key $\text{sk}' = (1, \psi(s))$. We next have to perform a “relinearization” operation which converts v' back to a ciphertext that decrypts to $\psi(\alpha)$ under the original secret key sk . This operation can itself be broken down into two steps:

2. **Break into digits:** decompose p'_1 into “small” pieces: $p'_1 = \sum_k q'_k \Delta_k$. Here, the Δ_k ’s are integer constants, and the pieces q'_k are elements of R of small norm. This operation is rather expensive, as it requires conversions between `DoubleCRT` and coefficient representations of elements in R_q .
3. **Key switching:** compute the ciphertext $(p''_0 + p''_0, p''_1)$, where

$$p''_j = \sum_k q'_k A_{jk}, \quad (j = 0, 1).$$

Here, the A_{jk} ’s are the “key switching matrices”, namely, pre-computed elements in R_Q (for some larger Q) which are stored in the public key. The A_{jk} ’s are stored in `DoubleCRT` format, so if we have the q'_k in the same `DoubleCRT` format then this operation is also a fast, linear time operation.

The key observation to our new technique is that we can reverse the order of the first two steps above, without affecting the correctness of the procedure. Namely our new procedure is as follows:

1. **Break into digits:** decompose the original p_1 before applying the automorphism into “small” pieces: $p_1 = \sum_k q_k \Delta_k$.
2. **Automorphism:** compute $p'_0 \leftarrow \psi(p_0)$, and $q'_k \leftarrow \psi(q_k)$ for each q_k . Namely, p'_0 is computed just as before, but we apply the automorphism to the pieces q_k from above rather than to p_1 itself.
3. **Key switching:** compute the ciphertext $(p'_0 + p''_0, p'_1)$, where

$$p''_j = \sum_k q'_k A_{jk}, \quad (j = 0, 1).$$

This is exactly the same computation as before.

The reasons that this works, is that (i) ψ is an automorphism (so it distributes over addition and multiplication), and (ii) applying ψ does not significantly change the norm of an element (cf. [10]). In a little more detail, correctness of the key-switching step depends only on the following two conditions on the q'_k 's:

- (a) $\sum_k q'_k \Delta_k = \psi(p_1)$, and
- (b) the q'_k 's have low norm.

Condition (a) is satisfied in our new procedure since ψ is an automorphism (which acts as the identity on integers), and so

$$\psi(p_1) = \psi\left(\sum_k q_k \Delta_k\right) = \sum_k \psi(q_k) \Delta_k = \sum_k q'_k \Delta_k.$$

Condition (b) is satisfied since the pieces q_k have small norm, and applying ψ to a ring element does not increase its norm significantly.

The new procedure is therefore just as effective as the old one, but now the expensive break-into-digits step can be preformed only once, as a pre-computation that depends only on v , rather than having to perform it for every automorphism ψ . The flip side is that we need to apply ψ to each one of the parts q_k instead of only once to p_1 . But as we mentioned, this is a cheap operation.

5.1 Interaction with Key-Switching Strategy

If we want to compute $\psi(v)$ for various automorphisms ψ , and we have key-switching matrices for all of the ψ 's. Then we can apply the above hoisting strategy directly. In some situations, what we want to do is compute $\theta^i(v)$ for $i = 0, \dots, D - 1$, where $\theta = \theta_{g_s}$ for some dimension s with generator $g_s \in \mathbb{Z}_m^*$, and where $D = D_s$ is the size of the dimension. If we are employing the baby-step/giant-step strategy for storing key-switching matrices, then we do not have all of the requisite key-switching matrices, so we cannot use the hoisting strategy directly. Instead, what we can do is the following. Since we have key-switching

matrices for all of the giant steps θ^{gj} , for $j = 1, \dots, h - 1$, we can use hoisting to compute $\theta^{gj}(v)$ for all of the giant steps, and for each of these values, we perform the pre-computation (i.e., the break-into-digits step). Then, since we have key-switching matrices for all of the baby steps θ^k , for $k = 1, \dots, g - 1$, we can compute any value $\theta^{gj+k}(v)$ as $\theta^k(\theta^{gj}(v))$, using the precomputed data for $\theta^{gj}(v)$ and the key-switching matrix for θ^k .

6 Algorithms for One-Dimensional Linear Transformations

In this section, we describe in detail our algorithms for applying one-dimensional linear transformations to a ciphertext v . We fix a dimension $s = 1, \dots, n$. Recall from Sect. 2.3 that $\rho \stackrel{\text{def}}{=} \rho_s$ is the rotation-by-1 map in dimension s , and that $D \stackrel{\text{def}}{=} D_s$ is the size of dimension s .

6.1 Logic for Basic MatMul1D

Recall from Sect. 3.1 that for the MatMul1D calculation, we need to compute

$$w = \sum_{i \in [D]} \kappa(i) \rho^i(v),$$

where the $\kappa(i)$'s are constants in R_{pr} that depend on the matrix.

If s is a good dimension, then ρ is realized with a single automorphism, $\rho = \theta \stackrel{\text{def}}{=} \theta_{g_s}$ where $g_s \in \mathbb{Z}_m^*$ is the generator for dimension s . We can easily implement this in a number of ways. For example, we can use the hoisting technique from Sect. 5 to compute all of the values $\theta^i(v)$ for $i \in [D]$. Alternatively, if we are using a minimal key-switching strategy (see Sect. 4.5), then with just a key-switching matrix for θ , we can compute the values $\theta^i(v)$ iteratively, computing $\theta^{i+1}(v)$ from $\theta^i(v)$ as $\theta(\theta^i(v))$.

6.2 Revised Logic for Bad Dimensions

From Eq. (3), if s is a bad dimension, then we have

$$\rho^i(v) = \mu(i)\theta^i(v) + \mu'(i)\theta^{i-D}(v), \tag{7}$$

where $\mu(i)$ is a “0–1 mask” and $\mu'(i) = 1 - \mu(i)$. As discussed in Sect. 4.4, it is useful to algebraically decouple the rotations and automorphisms in a bad dimension, which we can do as follows:

$$\begin{aligned} w &= \sum_{i \in [D]} \kappa(i) \rho^i(v) \\ &= \sum_{i \in [D]} \kappa(i) \{ \mu(i)\theta^i(v) + \mu'(i)\theta^{i-D}(v) \} \\ &= \sum_{i \in [D]} \kappa'(i)\theta^i(v) + \theta^{-D} \left[\sum_{i \in [D]} \kappa''(i)\theta^i(v) \right], \end{aligned}$$

where

$$\begin{aligned}\kappa'(i) &= \mu(i)\kappa(i) \quad \text{and} \\ \kappa''(i) &= \theta^D \{ \mu'(i)\kappa(i) \}.\end{aligned}$$

To implement this, we have to compute $\theta^i(v)$ for all $i \in [D]$. This can be done using the same strategies as were discussed above in a good dimension, using either hoisting or iteration. The only other automorphism we need to compute is one evaluation of θ^{-D} . Note that with our new key-switching strategy (see Sect. 4.3), we always have available a key-switching matrix for θ^{-D} .

If we ignore the cost of pre-computing all the constants in DoubleCRT format, we see that the computational cost is roughly the same in both good and bad dimensions. This is because the time needed to perform all the constant-ciphertext multiplications is very small in comparison to the time needed to perform all the automorphisms. The cost in noise is also about the same, essentially, one constant-ciphertext multiplication.

6.3 Baby-Step/Giant-Step Logic

We now present the logic for a new baby-step/giant-step multiplication algorithm. As discussed above in Sect. 4.1, this idea is very similar to the BlockMatMul1D implementation described in [8]. Set $g = \lceil \sqrt{D} \rceil$ and $h = \lceil D/g \rceil$. We have:

$$\begin{aligned}w &= \sum_{i \in [D]} \kappa(i)\rho^i(v) \\ &= \sum_{j \in [g]} \sum_{k \in [h]} \kappa(j + gk)\rho^{j+gk}(v) \\ &= \sum_{k \in [h]} \rho^{gk} \left[\sum_{j \in [g]} \kappa'(j + gk)\rho^j(v) \right],\end{aligned}$$

where $\kappa'(j + gk) = \rho^{-gk}(\kappa(j + gk))$.

Algorithm 1. In a good dimension, where $\rho = \theta$, we can implement the above logic using the following algorithm.

1. For each $j \in [g]$, compute $v_j = \theta^j(v)$.
2. For each $k \in [h]$, compute

$$w_k = \sum_{j \in [g]} \kappa'(j + gk)v_j.$$

3. Compute

$$w = \sum_{k \in [h]} \theta^{gk}(w_k).$$

Step 1 of the algorithm can be implemented by hoisting, or if we are using a minimal key-switching strategy, by iteration. Also, if we employ the minimal key-switching strategy, then Step 3 can be implemented using the Horner-rule idea discussed in Sect. 4.5 — for this, we just need a key-switching matrix for θ^g . Otherwise, if we have key switching matrices for all of the ρ^{gk} 's, it is somewhat faster to apply all of these automorphisms independently, which is also amenable to parallelization.

6.4 Revised Baby-Step/Giant-Step Logic for Bad Dimensions

Set $g = \lceil \sqrt{D} \rceil$ and $h = \lceil D/g \rceil$. Again, using Eq. (7), and the idea of algebraically decoupling the rotations and automorphisms in a bad dimension, we have:

$$\begin{aligned} w &= \sum_{i \in [D]} \kappa(i) \rho^i(v) \\ &= \sum_{i \in [D]} \kappa(i) \{ \mu(i) \theta^i(v) + \mu'(i) \theta^{i-D}(v) \} \\ &= \sum_{j \in [g]} \sum_{k \in [h]} \kappa(j + gk) \{ \mu(j + gk) \theta^{j+gk}(v) + \mu'(j + gk) \theta^{j+gk-D}(v) \} \\ &= \sum_{k \in [h]} \theta^{gk} \left[\sum_{j \in [g]} \{ \kappa'(j + gk) \theta^j(v) + \kappa''(j + gk) \theta^{j-D}(v) \} \right], \end{aligned}$$

where

$$\begin{aligned} \kappa'(j + gk) &= \theta^{-gk} \{ \mu(j + gk) \kappa(j + gk) \} \quad \text{and} \\ \kappa''(j + gk) &= \theta^{-gk} \{ \mu'(j + gk) \kappa(j + gk) \}. \end{aligned}$$

Based on this, we derive the following:

Algorithm 2.

1. Compute $v' = \theta^{-D}(v)$.
2. For each $j \in [g]$, compute $v_j = \theta^j(v)$ and $v'_j = \theta^j(v')$
3. For each $k \in [h]$, compute

$$w_k = \sum_{j \in [g]} \{ \kappa'(j + gk) v_j + \kappa''(j + gk) v'_j \}.$$

4. Compute

$$w = \sum_{k \in [h]} \theta^{gk}(w_k).$$

Step 2 of the algorithm can be implemented by hoisting, or if we are using a minimal key-switching strategy, by iteration. Also, if we employ the minimal key-switching strategy, then Step 4 can be implemented using Horner's rule.

As before, if we have key switching matrices for all of the ρ^{gk} 's, it is somewhat faster to apply all of these automorphisms independently, which is also amenable to parallelization.

Based on experimental data, we find that using the baby-step/giant-step multiplication algorithms are faster in dimensions for which we are using a baby-step/giant-step key-switching strategy. Moreover, even if we are using the full key-switching strategy, and we have all key-switching matrices for that dimensions available, the baby-step/giant-step multiplication algorithms are still faster in very large dimensions (say, on the order of several hundred).

6.5 Alternative Revised Baby-Step/Giant-Step Logic for Bad dimensions

We considered, implemented, and tested an alternative algorithm, which was found to be slightly slower and was hence disabled. It proceeds as follows: Set $g = \lceil \sqrt{D} \rceil$ and $h = \lceil D/g \rceil$.

$$\begin{aligned}
 w &= \sum_{i \in [D]} \kappa(i) \rho^i(v) \\
 &= \sum_{i \in [D]} \kappa(i) \{ \mu(i) \theta^i(v) + \mu'(i) \theta^{i-D}(v) \} \\
 &= \sum_{j \in [g]} \sum_{k \in [h]} \kappa(i) \{ \mu(j + gk) \theta^{j+gk}(v) + \mu'(j + gk) \theta^{j+gk-D}(v) \} \\
 &= \sum_{k \in [h]} \theta^{gk} \left[\sum_{j \in [g]} \kappa'(j + gk) \theta^j(v) \right] + \theta^{-D} \left(\sum_{k \in [h]} \theta^{gk} \left[\kappa''(j + gk) \theta^j(v) \right] \right),
 \end{aligned}$$

where

$$\begin{aligned}
 \kappa'(j + gk) &= \theta^{-gk} \{ \mu(j + gk) \kappa(j + gk) \} \quad \text{and} \\
 \kappa''(j + gk) &= \theta^{D-gk} \{ \mu'(j + gk) \kappa(j + gk) \}.
 \end{aligned}$$

Based on this, we derive the following:

Algorithm 3.

1. For each $j \in [g]$, compute $v_j = \theta^j(v)$
2. For each $k \in [h]$, compute

$$u_k = \sum_{j \in [g]} \kappa'(j + gk) v_j \quad \text{and} \quad u'_k = \sum_{j \in [g]} \kappa''(j + gk) v'_j.$$

3. Compute

$$u = \sum_{k \in [h]} \theta^{gk}(u_k) \quad \text{and} \quad u' = \sum_{k \in [h]} \theta^{gk}(u'_k).$$

4. Compute

$$w = u + \theta^{-D}(u').$$

6.6 BlockMatMul1D Logic

Recall from Sect. 3.2 that for the BlockMatMul1D calculation, we need to compute

$$\begin{aligned} w &= \sum_{j \in [d]} \sum_{i \in [D]} \kappa(i, j) \sigma^j(\rho^i(v)) \\ &= \sum_{j \in [d]} \sigma^j \left[\sum_{i \in [D]} \kappa'(i, j) \rho^i(v) \right], \end{aligned}$$

where $\kappa'(i, j) = \sigma^{-j}(\kappa(i, j))$. Here, σ is the Frobenius automorphism. This strategy is very similar to the baby-step/giant-step strategy used for the MatMul1D computation.

Algorithm 4. In a good dimension, where $\rho = \theta$, we can implement the above logic using the following algorithm.

1. Initialize an accumulator $w_j = 0$ for each $j \in [d]$.
2. For each $i \in [D]$:
 - (a) compute $v_i = \theta^i(v)$;
 - (b) for each $j \in [d]$, add $\kappa'(i, j)v_i$ to w_j .
3. Compute

$$w = \sum_{j \in [d]} \sigma^j(w_j).$$

Step 2(a) of the algorithm can be implemented by hoisting, or if we are using a minimal key-switching strategy, by iteration. Also, if we employ the minimal key-switching strategy, then Step 3 can be implemented using Horner's rule, using just a key-switching matrix for σ . If we have key switching matrices for all of the σ^j 's, it is somewhat faster to apply all of these automorphisms independently, which is also amenable to parallelization.

Often, D is much larger than d . Assuming we are using the hoisting technique in Step 2(a), it is much faster to perform Step 2(a) on the dimension of larger size D , and to perform Step 3 on the dimension of smaller size d . Indeed, the amortized cost of computing each of the d automorphisms in Step 3 is much greater than the amortized cost of computing each of the D automorphisms (via hoisting) in Step 2(a). Note that in our actual implementation, if it turns out that D is in fact smaller than d , then we switch the roles of θ and σ .

Observe that we store d accumulators w_0, \dots, w_{d-1} , rather than store the intermediate values v_0, \dots, v_{D-1} . Either strategy would work, but assuming D is much larger than d , we save space with this strategy (even though it is slightly more challenging to parallelize).

6.7 Revised BlockMatMul1D Logic for Bad Dimensions

Again, using Eq. (7) and the idea of algebraically decoupling rotations and automorphism, we have:

$$\begin{aligned}
 w &= \sum_{j \in [d]} \sum_{i \in [D]} \kappa(i, j) \sigma^j(\rho^i(v)) \\
 w &= \sum_{j \in [d]} \sum_{i \in [D]} \kappa(i, j) \sigma^j \{ \mu(i) \theta^i(v) + \mu'(i) \theta^{i-D}(v) \} \\
 &= \sum_{j \in [d]} \sigma^j \left[\sum_{i \in [D]} \kappa'(i, j) \theta^i(v) \right] + \theta^{-D} \left(\sum_{j \in [d]} \sigma^j \left[\sum_{i \in [D]} \kappa''(i, j) \theta^i(v) \right] \right),
 \end{aligned}$$

where

$$\begin{aligned}
 \kappa'(i, j) &= \sigma^{-j}(\kappa(i, j)) \mu(i) \quad \text{and} \\
 \kappa''(i, j) &= \theta^D \{ \sigma^{-j}(\kappa(i, j)) \mu'(i) \}.
 \end{aligned}$$

Based on this, we derive the following:

Algorithm 5.

1. Initialize accumulators $u_j = 0$ and $u'_j = 0$ for each $j \in [d]$.
2. For each $i \in [D]$:
 - (a) compute $v_i = \rho^i(v)$;
 - (b) for each $j \in [d]$, add $\kappa'(i, j)v_i$ to u_j and add $\kappa''(i, j)v_i$ to u'_j
3. Compute

$$u = \sum_{j \in [d]} \sigma^j(u_j) \quad \text{and} \quad u' = \sum_{j \in [d]} \sigma^j(u'_j).$$

4. Compute

$$w = u + \theta^{-D}(u').$$

As above, Step 2(a) of the algorithm can be implemented by hoisting, or if we are using a minimal key-switching strategy, by iteration. Also, if we employ the minimal key-switching strategy, then Step 3 can be implemented using Horner's rule, using just a key-switching matrix for σ . Again, if it turns out that D is in fact smaller than d , then we switch the roles of θ and σ .

7 Algorithms for Arbitrary Linear Transformations

So far, we have described algorithms for applying one-dimensional linear transformations to an encrypted vector, that is, E - or \mathbb{Z}_{p^r} -linear transformations that act independently on the hypercolumns in a single dimension (i.e., the MatMul1D and BlockMatMul1D operations introduced in Sect. 3). Many of the techniques

we have introduced can be adapted to arbitrary linear transformations. However, from a software design point of view, we adopted a strategy of designing a simple reduction from the general case to the one-dimensional case. For some parameter settings, this approach may not be optimal, but it is almost always much faster than the previous implementations of these operations in `HElib`.

We first consider the `MatMulFull` operation, which applies a general E -linear transformation to an encrypted vector. Here, an encrypted vector is a ciphertext whose corresponding plaintext is a vector with $\ell = \phi(m)/d$ slots. One can easily extend the `MatMulFull` operation to E -linear transformations on larger encrypted vectors that comprise several ciphertexts, although we have not yet implemented such an extension.

Recall from Sect. 2.3 that $\ell = D_1 \cdots D_n$, where for $s = 1, \dots, n$, the size of dimension s is D_s , and ρ_s is the rotation-by-1 map on dimension s . In [7], it was observed that we can apply the `MatMulFull` operation to a ciphertext v by using a generalization of the simple rotation strategy we presented above in Eq. (4). More specifically, if T is an E -linear transformation on R_{pr} , then for every $v \in R_{pr}$, we have

$$T(v) = \sum_{i_1 \in [D_1]} \cdots \sum_{i_n \in [D_n]} \kappa_{i_1, \dots, i_n} \cdot (\rho_n^{i_n} \cdots \rho_1^{i_1})(v), \quad (8)$$

where the κ_{i_1, \dots, i_n} 's are constants in R_{pr} determined by the linear transformation. For each (i_1, \dots, i_{n-1}) , there is a one-dimensional E -linear transformation $T'_{i_1, \dots, i_{n-1}}$ that acts on dimension n , such that for every $w \in R_{pr}$, we have

$$T'_{i_1, \dots, i_{n-1}}(w) = \sum_{i_n \in [D_n]} \kappa_{i_1, \dots, i_n} \cdot (\rho_n^{i_n} \cdots \rho_1^{i_1})(w).$$

Therefore, we can refactor Eq. (8) as follows:

$$T(v) = \sum_{i_1 \in [D_1]} \cdots \sum_{i_{n-1} \in [D_{n-1}]} T'_{i_1, \dots, i_{n-1}} \{(\rho_{n-1}^{i_{n-1}} \cdots \rho_1^{i_1})(v)\}. \quad (9)$$

To implement Eq. (9), we compute all of the rotations $(\rho_{n-1}^{i_{n-1}} \cdots \rho_1^{i_1})(v)$ using a simple recursive algorithm. The main type of operation performed here is to compute all of the rotations $\rho_s^{i_s}(w)$ for a given w , a given dimension, and for all $i_s \in [D_s]$. In a good dimension, where $\rho_s = \theta_{g_s}$, we can use hoisting (see Sect. 5) to speed things up, provided the required key-switching matrices are available, or sequentially if not. For bad dimensions, we can use the decoupling idea discussed in Sect. 4.4. Specifically, using Eq. (7), if $\theta \stackrel{\text{def}}{=} \theta_{g_s}$, then

$$\rho_s^{i_s}(w) = \mu_{i_s} \theta^{i_s}(w) + (1 - \mu_{i_s}) \theta^{i_s - D_s}$$

for an appropriate mask μ_{i_s} . Then we can compute $w' = \theta^{-D_s}$, which requires a single key-switching using our new key-switching strategy (see Sect. 4.3). After this, we need to compute $\theta^{i_s}(w)$ and $\theta^{i_s}(w')$ for all $i_s \in [D_s]$, which again, can be done by hoisting or iteration, as appropriate.

The other main type of operation needed to implement Eq. (9) is the application of all of the one-dimensional transformations $T'_{i_1, \dots, i_{n-1}}$ in dimension n , for which we can use our improved implementation of `MatMul1D`.

The speedup over the previous implementation in HELib will be roughly equal to the speedup of our new implementation of `MatMul1D` in dimension n . So to get the best performance, our implementation orders the dimensions so that D_n is the largest dimension size. If dimension n is a bad dimension, we also save on noise as well (we save noise equal to that of one constant-ciphertext multiplication). In many applications, it is desirable to choose parameters so that there is one very large dimension, and zero, one, or two very small dimensions — indeed, by default, HELib will choose parameters in this way. In this typical setting, the speedup for `MatMulFull` will be very significant.

Finally, we mention that the above techniques carry over in an obvious way to general \mathbb{Z}_{p^r} -linear transformations on R_{p^r} . As above, there is a simple reduction from the general `BlockMatMulFull` operation to the one-dimensional `BlockMatMul1D` operation. The previous implementation of `BlockMatMulFull` was not particularly well optimized, and because of this, the speedup we get is roughly equal to n times the speedup of our implementation of `BlockMatMul1D`, where, again, n is the number of dimensions in the underlying hypercube.

8 Application to “thin” Bootstrapping

HELib implements a general bootstrapping algorithm, which will convert an arbitrary noisy ciphertext into an equivalent ciphertext with less noise. However, in some applications, ciphertexts are not completely arbitrary. Recall that plaintexts can be viewed as vectors of slots, where each slot contains an element of $E = \mathbb{Z}_{p^r}[\zeta]$, where ζ is a root of a polynomial over \mathbb{Z}_{p^r} of degree d . In some applications, one sometimes works with “thin” plaintexts, where the slots contain “constants”, i.e., elements of the subring \mathbb{Z}_{p^r} of E .

One could of course apply the HELib bootstrapping algorithm directly to such “thin” ciphertexts, but that would be quite wasteful. We can get more efficient implementation (in an amortized sense) by bootstrapping “batches” of d ciphertexts at a time: We can take d thin ciphertexts, pack them together to form a single ciphertext where each slot is fully packed, bootstrap this fully packed ciphertext, and then unpack it back to d thin ciphertexts. This approach, however, is only applicable when we have many ciphertexts to bootstrap, and it is not very convenient from a software engineering perspective. Moreover it also introduces some additional noise in the packing/unpacking steps.

Recently, Chen and Han devised an approach for more efficient and direct bootstrapping of thin ciphertexts [4], and we adapted their approach to HELib. We combined Chen and Han’s ideas with numerous optimizations for the linear algebra part of the bootstrapping from [8], reducing the bulk computation to a sequence of `MatMul1D` operations, where our improved algorithms for these operations yield great performance dividends. We implemented this new thin bootstrapping, and report on its performance below in Sect. 9.

Let us review the bootstrapping procedure of [8], which has been implemented in `HElib`, and then outline how to adapt it to incorporate Chen and Han’s technique.

A plaintext element $\alpha \in R_{p^r}$ can be viewed in a couple of different ways. It can be viewed as a vector of plaintext slots:

$$\alpha = \left(\sum_j a_{1j} \zeta^j, \dots, \sum_j a_{\ell j} \zeta^j \right),$$

where the a_{ij} ’s are scalars in \mathbb{Z}_{p^r} . Here, $\sum_j a_{ij} \zeta^j \in E$ is the content of the i th slot of α . For a thin plaintext, only the a_{i0} ’s are non-zero elements in E .

The above representation corresponds to some \mathbb{Z}_{p^r} -basis of R_{p^r} , namely $\alpha = \sum_{ij} a_{ij} \lambda_{ij}$ (with $\lambda_{ij} \in R_{p^r}$ being the element with ζ^j in the i th slot and zero elsewhere). But we can express the same α on an arbitrary \mathbb{Z}_{p^r} -basis $\{\beta_{ij}\}$ of R_{p^r} ,

$$\alpha = \sum_{ij} b_{ij} \beta_{ij} \quad (\text{where } b_{ij} \in \mathbb{Z}_{p^r}).$$

For example, for the *power basis*, the β_{ij} ’s are powers of X modulo $(p^r, \phi_m(X))$. As it turns out, for bootstrapping it is more convenient to use the *powerful basis*, introduced by Lyubashevsky et al. [11, 12] and developed further by Alperin-Sheriff and Peikert [1]. The bootstrapping algorithms in `HElib` make use of the powerful basis, as it allows us to decompose the required linear transformations into a sequence of one-dimensional linear transformations.

Here is a rough outline of `HElib`’s bootstrapping procedure for fully packed ciphertexts. We start out with a ciphertext encrypting a plaintext $\beta = \sum_{ij} b_{ij} \beta_{ij}$.

1. Perform a modulus switching and homomorphic inner product, obtaining a ciphertext with very little noise that encrypts some $\beta^* = \sum_{ij} b_{ij}^* \beta_{ij}$. The b_{ij}^* coefficients are actually in \mathbb{Z}_{p^s} for some $s > r$, and have the property that there is a (non-linear) “digit extraction” procedure that computes b_{ij} from b_{ij}^* . (In more detail, it computes $b_{ij} = \lfloor b_{ij}^* / p^{s-r} \rfloor$.)
2. Perform a linear “coefficient to slot” operation that transforms the ciphertext encrypting β^* to one encrypting $\alpha^* = (\sum_j b_{1j}^* \zeta^j, \dots, \sum_j b_{\ell j}^* \zeta^j)$.
3. Unpack the ciphertext encrypting α^* into d thin ciphertexts, where for $j = 0, \dots, d-1$, the j th unpacked ciphertext encrypts $(b_{1j}^*, \dots, b_{\ell j}^*)$.
4. Apply the above-mentioned “digit extraction” procedure to each unpacked thin ciphertext, obtaining d thin ciphertexts, where the j th ciphertext is an encryption of $(b_{1j}, \dots, b_{\ell j})$.
5. Repack the thin ciphertexts from the previous step, obtaining an encryption of $\alpha = (\sum_j b_{1j} \zeta^j, \dots, \sum_j b_{\ell j} \zeta^j)$.
6. Perform a linear “slot to coefficient” operation, which is the inverse of the “coefficient to slot” operation in Step 2, to transform the encryption of α in the previous step to an encryption of $\beta = \sum_{ij} b_{ij} \beta_{ij}$.

By careful usage of the powerful basis for $\{\beta_{ij}\}_{ij}$, each of the linear operations, “coefficient to slot” and “slot to coefficient”, can be implemented using

one `BlockMatMul1D` operation and a small number (typically one or two) `MatMul1D` operations. More specifically, the “slot to coefficient” transformation L can be decomposed as $L = L_t \cdots L_2 L_1$, where L_1 is a one-dimensional \mathbb{Z}_{p^r} -linear transformation (i.e., a `BlockMatMul1D` operation), and L_2, \dots, L_n are one-dimensional E -linear transformations (i.e., `MatMul1D` operations). The inverse “coefficient to slot” transformation can therefore also be decomposed as $L^{-1} = L_1^{-1} L_2^{-1} \cdots L_n^{-1}$. See [8] for details of the definitions of the maps L_1, \dots, L_n .

We now review Chen and Han’s technique from [4], adapted to HELib’s strategy to dealing with linear transformations. We start with a ciphertext encrypting a thin plaintext $\alpha = (a_{10}, \dots, a_{\ell 0})$.

1. First apply the “slot to coefficient” transformation, obtaining an encryption of $\beta = \sum_i a_{i0} \beta_{i0}$.
2. Perform the modulus switching and homomorphic inner product, obtaining a ciphertext with less noise that encrypts $\beta^* = \sum_{ij} a_{ij}^* \beta_{ij}$.
3. Apply the “coefficient to slot” transformation, which places $\sum_{ij} a_{ij}^* \zeta^j$ in the i th slot, followed by a slot-wise projection function π that maps each $\sum_{ij} a_{ij}^* \zeta^j$ to a_{i0}^* , obtaining a ciphertext that encrypts $\alpha^* = (a_{10}^*, \dots, a_{\ell 0}^*)$.
4. Apply the “digit extraction” procedure, obtaining a ciphertext encrypting $\alpha = (a_{10}, \dots, a_{\ell 0})$.

Clearly, this procedure only performs a single digit extraction operation, versus the d digit extraction operations that are required for fully packed bootstrapping.

As another benefit, observe that in Step 1 we are applying the linear transformation $L = L_t \cdots L_2 L_1$ to a thin plaintext. It turns out, that the restriction of L_1 to the subspace of thin plaintexts is in fact an E -linear transformation (this is easily seen from the definition of L_1 in [8]). Therefore, we can implement L_1 as a `MatMul1D` operation, rather than as a more expensive `BlockMatMul1D` operation. (The other transformations L_2, \dots, L_n are already implemented as `MatMul1D` operations.)

Moreover, in Step 3, we are computing

$$\pi L^{-1} = (\pi L_1^{-1}) L_2^{-1} \cdots L_n^{-1}.$$

We can rewrite πL_1^{-1} as τK , where τ is the slot-wise trace map and K is a certain E -linear transformation derived from L_1^{-1} . The trace map τ_E on E sends $\eta \in E$ to $\sum_{j=0}^{d-1} \sigma_E^j(\eta)$, where σ_E is the Frobenius map on E . The decomposition of πL_1^{-1} as τK follows from the general fact that for every \mathbb{Z}_{p^r} -linear map M from E to \mathbb{Z}_{p^r} , there exists $\lambda_M \in E$ such that $M(\eta) = \tau_E(\lambda_M \eta)$ for all $\eta \in E$. Indeed, L_1^{-1} can be represented by a matrix whose entries are themselves \mathbb{Z}_{p^r} -linear maps on E , and so πL_1^{-1} can be represented by a matrix whose entries are \mathbb{Z}_{p^r} -linear maps from E to \mathbb{Z}_{p^r} . If we replace each such map M with the multiplication-by- λ_M map, we obtain the matrix for the E -linear map K , and we have $\pi L_1^{-1} = \tau K$.

Thus, we can implement πL_1^{-1} using one `MatMul1D` operation and one application of the slot-wise trace map τ . We can quickly compute the slot-wise trace using one of several strategies. If we have key switching matrices for σ^j , for all $j = 1, \dots, d-1$, where $\sigma \stackrel{\text{def}}{=} \theta_p$, we can compute the trace of a ciphertext v via hoisting by first computing $\sigma^j(v)$ for $j = 0, \dots, d-1$, and then adding these up. Alternatively, if $v^{(s)} \stackrel{\text{def}}{=} \sum_{j=0}^{s-1} \sigma^j(v)$, we can use the relation $v^{(s+t)} = \sigma^t(v^{(s)}) + v^{(t)}$. If we are using the baby-step/giant-step key switching strategy, then we can compute the trace of v using $O(\log d)$ key-switching operations via a “repeated doubling” computation strategy. If we are using the minimal key-switching strategy, we can use this same relation to compute the trace of v using $O(\sqrt{d})$ key-switching operations via a baby-step/giant-step computation strategy; for this to work, we just need key-switching matrices for σ and σ^g , where $g \approx \sqrt{d}$.

9 Timings

We now present some timing data that demonstrates the effectiveness of our new techniques. All of our testing was done on a machine with an Intel Xeon CPU, E5-2698 v3 @2.30 GHz (which is a Haswell processor), featuring 32 cores and 250 GB of main memory. The compiler was GCC version 4.8.5, and we used NTL version 10.5.0 and GMP version 6.0.

Table 1 shows the running time (in seconds) for the old default behavior (“old def”) and the new default behavior (“new def”) for `MatMul1D` computations (see Sect. 3.1). We do this for various values of m defining a cyclotomic polynomial of degree $\phi(m)$. The quantity d is the order of $p \bmod m$ (which represents the “size” of each slot), while the quantity D is the size of the dimension. We worked with plaintext spaces modulo $p^r = 2$ in all of these examples. A value of D marked with “*” denotes a “bad” dimension. Table 1 does not show the time taken to build the constants associated with a matrix or to convert them to `DoubleCRT` representation. One sees that for the large dimension of size 682 (which is a typical size for many applications), we get a speedup of 30 if it is a good dimension, and a speedup of 75 if it is bad. Speedups for smaller dimensions are less dramatic, but still quite significant.

Table 2 shows more detailed information on various implementation strategies, as well as the cost of precomputing matrix constants. The “build” column shows the time to build the constants associated with the matrix in a polynomial representation. The “conv” column shows that time required to convert these constants to `DoubleCRT` representation. The following columns show that time required to perform the matrix-vector multiplication, based on a variety of key switching and algorithmic strategies. The columns are labeled as “[MBF]/[BF][HN]”, where

MBF: **M** is for **Min** KS strategy, **B** is for **Baby-step/giant-step** key-switching strategy, **F** is for **Full** key-switching strategy,

BF: **B** is for **Baby-step/giant-step** multiplication strategy, **F** is for **Full** multiplication strategy,

HN: **H** is for **Hoisting**, **N** is for **No** hoisting.

As one can see from the data, the cost of converting constant to DoubleCRT representation can easily exceed the cost of the remaining operations, so it is essential that these conversions are done as precomputations, if at all possible.

Consider the first line in Table 2. Column B/BH represents the default behavior: baby-step/giant-step key switching (since it is a large dimension of size 682), baby-step/step-step multiplication, and hoisting (only the baby steps are subject to hoisting). The next column (B/BN) is the same, except the baby steps are not hoisted, which is why it is slower. Column B/FH shows what happens if we do not use baby-step/giant-step multiplication, and rely exclusively on hoisting (as in Sect. 5.1). One can see that for such a large dimension, this is not an optimal strategy. Column M/B shows what happens when we use the minimal key switching strategy (with baby-step/step-step multiplication). Even though it needs only two key switching matrices (rather than about 50), it is less than twice as slow as the best strategy (although it does not parallelize very well). The algorithm represented by column B/FN corresponds directly to the algorithm originally implemented in HELib. The next line in the table represents a bad dimension. We note that for bad dimensions, the algorithm originally implemented in HELib is about twice as slow as the one represented by column B/FN (this is why the timing data in Table 1 for bad dimensions is not equal to the numbers in column B/FN of Table 2).

Table 3 shows corresponding timing data for BlockMatMul1D computations (see Sect. 3.2). For good dimensions, the previous implementation in HELib roughly corresponds to the non-hoisting strategy in our new implementation. So one can see that with hoisting we get a speedup of up to 4 times over the previous implementation for large dimensions (but only about 1.5 for small dimensions). For large, bad dimensions, in the previous implementation in HELib, the running time will be close to twice that of the non-hoisting strategy in our new implementation; therefore, the speedup in such dimensions is close to a factor of about 5.

Table 4 shows the effectiveness of parallelization using multiple cores. We show times for both MatMul1D and BlockMatMul1D, using 1, 4, and 16 threads. These times are for the default strategies, and do not show the time required to build the matrix constants or convert them to DoubleCRT representation. While the speedups do not quite scale linearly with the number of cores, they are clearly significant, with 16 cores yielding roughly an $8\times$ speedup in large dimensions and $4\times$ speedup in small ones.

We do not present detailed results for the running times of our new implementation of MatMulFull and BlockMatMulFull, discussed in Sect. 7. However, our experiments indicate that the speedups predicted in Sect. 7 closely align with practice: the speedup for MatMulFull is about the same as our speedup for MatMul1D in the largest dimension; the speedup for BlockMatMulFull is roughly our speedup for BlockMatMul1D in the largest dimension, times the number of dimensions in the hypercube.

Finally, we present some timing results to demonstrate the efficacy of our new algorithms in the context of bootstrapping, as discussed in Sect. 8. We chose large parameters that demonstrate well the potential saving with our new

implementation. Specifically, we used $m = 49981$ and $p^r = 2$, for which we have $\phi(m) = 49500$ and $d = 30$. The hypercube structure for $\mathbb{Z}_m^*/(p^r)$ has two dimensions, one of size 150 and one of size 11, for a total of 1650 slots. We note that most parameter choices in [8] attempted to balance the size of the different dimensions, specifically because the linear transformations would take too long otherwise. One of the benefits of our faster algorithms is thus to free us from having to consider that aspect; indeed, our timing shows that the linear transformations are now quite fast even for this “unbalanced” setting.

We ran our tests with ciphertexts with 55 “levels” (for an estimated security parameter of about 80). For these parameters, the bootstrapping procedure consumes about 10 levels, leaving about 45 levels for other computations. Table 5 shows the running time (in seconds) for both the thin bootstrapping and packed bootstrapping routines with both the old and new matrix multiplication algorithms. These results make it clear that for such large hypercubes, thin bootstrapping must be done using our new, faster matrix multiplication to be truly practical.

Table 1. MatMul1D: summary of old vs new, time in seconds

m	$\phi(m)$	d	D	old def	new def	speedup
15709	15004	22	682	69.28	2.22	31.20
15709	15004	22	682*	138.20	3.14	75.86
18631	18000	25	120	20.27	1.38	14.69
18631	18000	25	120*	39.97	1.69	23.65
24295	18816	28	42	3.18	0.51	6.24
24295	18816	28	42*	6.20	0.55	11.27

Table 2. Different strategies for MatMul1D, time in seconds

m	$\phi(m)$	d	D	build	conv	M/B	M/F	B/BH	B/BN	B/FH	B/FN	F/FH	F/FN
15709	15004	22	682	0.47	5.54	3.80	44.81	2.22	3.19	6.46	69.28	5.30	28.30
15709	15004	22	682*	0.56	11.07	5.93	44.86	3.14	5.03	7.33	69.70	5.94	29.16
18631	18000	25	120	0.08	1.96	2.43	13.81	1.38	2.04	2.36	20.27	1.29	8.70
18631	18000	25	120*	0.10	3.91	3.68	13.95	1.69	2.89	2.45	20.27	1.29	8.78
24295	18816	28	42	0.03	0.70	1.39	5.09	0.82	1.17	1.11	6.87	0.51	3.18
24295	18816	28	42*	0.04	1.39	2.17	5.09	0.95	1.64	1.20	6.94	0.55	3.20

Table 3. Different strategies for BlockMatMul1D, time in seconds

m	$\phi(m)$	d	D	build	conv	M/	B/H	B/N	F/H	F/N
15709	15004	22	682	15.47	122.62	54.73	21.03	84.42	18.15	42.67
15709	15004	22	682*	17.31	246.89	64.98	36.81	99.84	32.41	57.07
18631	18000	25	120	2.44	49.59	18.83	9.84	27.90	6.88	14.66
18631	18000	25	120*	2.96	98.79	23.83	17.62	35.80	12.73	20.58
24295	18816	28	42	0.95	19.73	9.25	7.84	13.64	5.01	7.70
24295	18816	28	42*	1.15	39.72	13.49	14.73	20.45	9.65	12.47

Table 4. Multithreading for MatMul1D/BlockMatMul1D, time in seconds

m	$\phi(m)$	d	D	MatMul1D			BlockMatMul1D		
				$nt = 1$	$nt = 4$	$nt = 16$	$nt = 1$	$nt = 4$	$nt = 16$
15709	15004	22	682	2.18	0.67	0.29	20.21	7.60	2.47
15709	15004	22	682*	3.14	0.97	0.42	35.50	12.17	4.70
18631	18000	25	120	1.35	0.49	0.20	7.97	2.49	1.03
18631	18000	25	120*	1.65	0.58	0.29	13.89	4.30	1.67
24295	18816	28	42	0.47	0.23	0.15	4.98	1.37	0.61
24295	18816	28	42*	0.51	0.22	0.14	9.51	2.67	1.08

Table 5. Bootstrapping, time in seconds

	old		new	
	total	linear	total	linear
thin bootstrap	474.18	428.76	80.31	36.17
packed bootstrap	2120.05	804.30	1413.02	102.65

References

1. Alperin-Sheriff, J., Peikert, C.: Practical bootstrapping in quasilinear time. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8042, pp. 1–20. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40041-4_1
2. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: Fully homomorphic encryption without bootstrapping. In: Innovations in Theoretical Computer Science (ITCS 2012) (2012). <http://eprint.iacr.org/2011/277>
3. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) fully homomorphic encryption without bootstrapping. ACM Trans. Comput. Theory **6**(3), 13 (2014)
4. Chen, H., Han, K.: Homomorphic lower digits removal and improved FHE bootstrapping. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018. LNCS, vol. 10820, pp. 315–337. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78381-9_12

5. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: Proceedings of the 41st ACM Symposium on Theory of Computing - STOC 2009, pp. 169–178. ACM (2009)
6. Gentry, C., Halevi, S., Smart, N.P.: Fully homomorphic encryption with polylog overhead. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 465–482. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_28
7. Halevi, S., Shoup, V.: Algorithms in HELib. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 554–571. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44371-2_31
8. Halevi, S., Shoup, V.: Bootstrapping for HELib. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part I. LNCS, vol. 9056, pp. 641–670. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46800-5_25
9. Halevi, S., Shoup, V.: HELib - an implementation of homomorphic encryption, September 2014. <https://github.com/shaih/HELib/>
10. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 1–23. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5_1
11. Lyubashevsky, V., Peikert, C., Regev, O.: A toolkit for ring-LWE cryptography. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 35–54. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38348-9_3
12. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. *J. ACM* **60**(6), 43 (2013). Early version in EUROCRYPT 2010
13. Rivest, R., Adleman, L., Dertouzos, M.: On data banks and privacy homomorphisms. In: Foundations of Secure Computation, pp. 169–177. Academic Press (1978)
14. Roman, S.: Field Theory, 2nd edn. Springer, New York (2006). <https://doi.org/10.1007/0-387-27678-5>
15. Smart, N.P., Vercauteren, F.: Fully homomorphic SIMD operations. *Des. Codes Cryptogr.* **71**(1), 57–81 (2014). Early verion at <http://eprint.iacr.org/2011/133>