Vive Galois! Part 1: Optimal SIMD Packing and Packed Bootstrapping for FHE

Chris Peikert* Zachary Pepin†

September 25, 2025

Abstract

The vast majority of work on the efficiency of lattice-based cryptography, including fully homomorphic encryption (FHE), has relied on *cyclotomic* number fields and rings. This is because cyclotomics offer a wide variety of benefits, including good geometrical properties, fast ring arithmetic, and rich homomorphic operations like vectorized (SIMD) operations on "packed" plaintexts, automorphisms, and ring-switching. However, selecting a suitable cyclotomic that has the desired number and type of plaintext "slots," while also balancing security and efficiency, is a highly constrained problem that often lacks an ideal solution, resulting in wasted SIMD capacity and lost efficiency.

This work provides a suite of tools for instantiating ring-based lattice cryptography to work over *subfields* of cyclotomics, which provide more flexibility and better-fitting parameters for applications. A particular focus is on realizing FHE with *optimal plaintext packing* and homomorphic SIMD parallelism for *any* plaintext characteristic, along with efficient *packed bootstrapping* that fully exploits this parallelism.

Toward this end, this (two-part) work makes the following main technical contributions, all of which are catalyzed by Galois theory:

- For sampling and decoding errors in encryption and decryption (respectively), we construct *geometrically short, structured bases* for the number rings of arbitrary subfields of prime-power cyclotomics (and hence their composites as well).
- For fast ring arithmetic, we define and establish analogous structural properties for Chinese Remainder Theorem (CRT) bases in *abelian* number rings, and give *specialized fast transforms* that map between CRT bases and any similarly structured bases.
- For packed bootstrapping and homomorphic linear algebra, we give a general framework for *homomorphic evaluation of structured linear transforms* in abelian number rings, and show that CRT transforms can be evaluated using relatively few homomorphic operations.

1 Introduction

Since Gentry's seminal work [Gen09b, Gen09a] on fully homomorphic encryption (FHE), there has been enormous progress in its efficiency, security, and utility, both theoretically and in practice. See, e.g., [SV11, BV11a, BV11b, BGV12, GHS12a, GHS12b, GHS12c, Bra12, GSW13], for some of the key developments of the first few years.

*University of Michigan, cpeikert@umich.edu. This material is based upon work supported by Intel Corporation and by DARPA under Agreement No. HR00112020025. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of Intel, the United States Government, or DARPA.

†University of Michigan, zapepin@umich.edu.

The present work is mainly concerned with the efficiency and flexibility of so-called "second generation" *exact* FHE schemes, in the style of [BGV12, Bra12, FV12]. This is in contrast to more recent *approximate* schemes for complex or real numbers as introduced in [CKKS17], and "third-generation" exact FHE as introduced in [GSW13]. While approximate FHE has undergone major progress for applications like neural networks, and third-generation schemes have seen many improvements and powerful applications, second-generation exact FHE remains the leading approach for *bulk algebraic computations* on "large" plaintexts over finite fields or rings, particularly those with a good deal of inherent parallelism. These computations include important operations like *transciphering* [Gen09a]: this "upgrades" lightweight symmetric cryptography to have full homomorphism via homomorphic evaluation of pseudorandom functions, providing enormous time and bandwidth improvements for clients. (See, e.g., [GHS12c, ARS+15, CPS18, ADE+23, DJL+24] for various implementations of transciphering for certain symmetric-key primitives.)

1.1 Efficiency of FHE

One main efficiency technique for FHE, introduced in [SV11], is "SIMD packing." This encrypts a *vector* of values from a small plaintext space into a single ciphertext, so that homomorphic addition and multiplication on ciphertexts induce *component-wise* ("single instruction, multiple data") addition and multiplication on the plaintext vectors. In addition, *automorphisms* of the underlying ring induce certain permutations or other algebraic operations on the plaintext vectors, which can unlock major efficiency gains [BGV12, GHS12b]. However, the precise nature and number of plaintext "slots" that can be obtained is a subtle and delicate issue that we discuss in more detail in Section 1.1.1 below.

Another central FHE technique is Gentry's idea of *bootstrapping*, which to date has been necessary for obtaining FHE schemes that can evaluate *arbitrary* (unbounded) functions. In addition, "boostrapping as an optimization" [BGV12] is also useful for more efficiently evaluating functions of sufficient complexity. We recall that homomorphic operations increase the intrinsic "noise" in ciphertexts, which decryption removes by a certain kind of decoding; too much noise would result in an incorrect output. Bootstrapping effectively reduces the noise in a ciphertext, by *homomorphically evaluating the decryption function*—expressed, for the ciphertext in question, as a *function of the secret key*—on an encryption of that key. For appropriate parameters, this yields a lower-noise encryption of the same message, which supports further useful homomorphic computation. For the overall efficiency of FHE, it is therefore vital to express the decryption function as efficiently as possible in terms of the scheme's intrinsic homomorphic features (e.g., SIMD slots) and operations (addition, multiplication, automorphisms, etc.).

1.1.1 Cyclotomics and their Limitations

The vast majority of work on efficiency for FHE has used *cyclotomic* number fields and rings. There are several reasons for this focus. For security, cyclotomics were the first class of number fields to have a worst-case hardness theorem for the Ring-LWE problem [LPR10] (though later work gave such a theorem for *all* number fields [PRS17]), and they have good geometrical properties that yield favorable parameters. For functionality and efficiency, they have fast specialized algorithms for the ring operations used in cryptography [LMPR08, LPR13]; they support SIMD packing and have a full set of automorphisms, i.e., they are *Galois* extensions of the rationals; they support "ring/field-switching" [GHPS12] between related cyclotomics, etc.

On the other hand, when it comes to SIMD packing, cyclotomics can be cumbersome to use and wasteful. Typically, an application of FHE will desire plaintext "slots" that are isomorphic to a specific finite ring or field, e.g., $\mathbb{Z}_q$ or $\mathbb{F}_q$ for a particular integer q. One can search for a cyclotomic that has the desired slot

type, or at least an *extension* of it. The degree f of the extension times the number g of slots equals the dimension of the cyclotomic (ignoring ramification, which can only decrease fg), so degrees $f > 1$ represent a suboptimal "slot type" (for packing plaintexts) and "SIMD capacity" (for parallel homomorphic operations). Unfortunately, f is difficult to control, and $f > 1$ is *inherent* for prime characteristics smaller than the cyclotomic conductor, which is typically in the many thousands. (The effect of a large extension degree f was mitigated somewhat in [ALJ+22], but via complex homomorphic 'recoding' algorithms, and still with wasted slot space.) In addition, the degree of the cyclotomic should lie in a relatively narrow range, to provide the desired level of security with reasonable efficiency. Altogether, the designer faces a highly constrained optimization problem, whose solution may be far from ideal. For instance, for homomorphic evaluation of the AES function, the most natural plaintext slot type is $\mathbb{F}_{2^8}$, but the use of cyclotomics in [GHS12c] induced a slot type of $\mathbb{F}_{2^{24}}$, representing a threefold loss in SIMD capacity.

Beyond cyclotomics. An interesting approach to circumvent these kinds of difficulties was given in [AH17], which proposed working in *decomposition subfields* of cyclotomics. Essentially, the decomposition subfield for a prime is the largest subfield in which the prime "splits completely"; in term of SIMD slots, the subfield has exactly the desired slot type with no wasted space (extension degree $f = 1$), and the optimal number of such slots (g is the dimension of the subfield), making it ideal for homomorphic SIMD computations.[1] It is straightforward to find the decomposition subfield of any cyclotomic, for any particular prime.

However, many aspects of a complete solution for FHE, and for efficient ring-based cryptography more generally, based on decomposition subfields were left untreated in [AH17]. As smaller matters, it considered only *prime* cyclotomics (not prime-power or more general cyclotomics), and it did not give any specialized fast algorithms for arithmetic in these subfields (like we have for arbitrary cyclotomics), but instead relied on generic FFT convolution. More importantly, it did not consider *(packed) bootstrapping* in these decomposition subfields, which is a very important tool to go with the optimal SIMD packing provided by these fields.

1.2 Contributions

This two-part work provides a suite of tools for instantiating ring-based cryptography, and FHE in particular, over a very wide class of *subfields* of cyclotomics.[2] By the Kronecker–Weber theorem, any *abelian* number field—i.e., a finite-degree Galois extension of the rationals whose Galois group is abelian—is a subfield of some cyclotomic, so cyclotomic subfields are indeed a broad class of number fields. A primary achievement of our work is the realization of FHE and efficient *packed bootstrapping* algorithms, with *optimal* plaintext packing and SIMD parallelism for *any* plaintext characteristic—avoiding the wasted SIMD capacity of cyclotomics.

More broadly, our overarching contribution is a mathematical and algorithmic framework for this setting, with an emphasis on *generality*, optimization of *concrete geometric bounds*, and asymptotic *algorithmic efficiency*. We believe that these tools will be useful for many other applications in ring-based cryptography, beyond FHE. (Implementing and optimizing our techniques in practice, and evaluating them for specific applications of interest, will take substantial additional effort, which we leave to future work.) In summary, our work's main technical contributions are:

[1] In the setting of approximate FHE, an analogous special case of working in the *maximal totally real subfield* of a cyclotomic, to get a number of *real*-valued slots equaling the field degree, was proposed in [KS18].

[2] Specifically, our treatment covers all number fields that are *composites* of subfields of *prime-power* cyclotomics. This mild restriction ensures that all the number-ring extensions we consider are free modules over their base rings, which is needed for some of our goals. In general, this free-module structure is not present for arbitrary cyclotomic subfields (see Footnote 5), so some restriction is necessary.

1. constructions of *short, structured bases* for a wide class of abelian number rings and their duals;
2. definitions, constructions, and factorizations of *Chinese Remainder Theorem (CRT) bases* in arbitrary towers of abelian number rings, which yield *fast ring arithmetic* via CRT transforms;
3. a general framework for *sparse decompositions* of linear transforms in terms of automorphisms, which yields *fast "packed" bootstrapping* and related tools for homomorphic linear algebra.

In Sections 3 and 6 we show how all this fits together in the context of homomorphic encryption with packed bootstrapping, and in Appendix B we revisit homomorphic evaluation of the AES function [GHS12c] and obtain optimally packed parameterizations with other beneficial features.

A common theme of our technical contributions is their heavy use of Galois theory, and the related theory of prime splitting in Galois number-field extensions. Working with arbitrary (abelian) Galois extensions not only provides a high level of generality and flexibility, but also highlights the fundamental aspects of the tools, independent of implementation details.

Due to the total amount of material and the varying mathematical background needed for each specific contribution, we have split our work into two parts. The first (present) part is focused mainly on algorithmic and cryptographic aspects, and covers Items 2 and 3 above. The second part [PP25] covers Item 1 (along with some associated algorithms), and is technically much heavier, involving several mathematical tools that are not needed in this first part (and exceeding it in length). In the rest of this introduction, we give an overview of all these contributions and how they come together for cryptographic applications.

1.2.1 Non-Contributions

We set the context by first explaining what this work does *not* contribute, because the prior literature already provides it for our setting of cyclotomic subfields, without any modification. First, for Ring-LWE in *arbitrary* number fields there are *worst-case hardness theorems* for the search [LPR10] and decision [PRS17] variants, whose quantitative parameters do not depend on the choice of field, only its degree n.[3] So, we already have the same kinds of hardness guarantees for cyclotomic subfields as we have for cyclotomics themselves.[4]

Second, consistent with the theorems and recommended usage of Ring-LWE (see [LPR10, Section 3.3]), we work with error (i.e., ciphertext "noise") that is nearly spherical in the *canonical embedding* of the number field into $\mathbb{C}^n$. As noted in prior works, this makes it straightforward to analyze the error growth in ciphertexts, because both addition and multiplication are coordinate-wise, and we can obtain rather tight bounds using tools like subgaussianity (see, e.g., [LPR13]). Because the ambient space is merely $\mathbb{C}^n$, all this analysis works just as well in arbitrary numbers fields as it does in cyclotomics.

However, as explained in [LPR13], the size of the error in the canonical embedding is not the only relevant quantity for decryption. Under the recommended usage of Ring-LWE and its hardness results, what we really need is to be able to *decode*, under the accumulated error, a certain lattice $R^\vee$ that is dual to the ring of integers R of the number field. This task depends on the geometry of R in the canonical embedding, and hence on the choice of number field. This motivates the first main contribution of our work.

[3] In fact, even the original search-to-decision reduction for cyclotomics [LPR10] turns out to work verbatim for any Galois number field, also with no change in the parameters.

[4] We remark that certain parameterizations of Ring-LWE over certain specially crafted, non-Galois number fields were shown to be insecure [ELOS15]. However, subsequent work [CIV16, Pei16] showed that the *error distributions* of these weak parameters have a very different shape from those covered by the cited hardness theorems, and in fact they are so narrow that they reveal several *errorless* LWE equations, making them trivially easy to break. Moreover, it was also shown in [Pei16] that any parameterization conforming to the hardness theorems, over *any* number field, is provably immune to the class of attacks from [ELOS15].

1.2.2 Short, Structured Bases for Abelian Number Rings

A basic requirement for ring-based cryptography over a number ring R is to have a known, *relatively "short,"* and preferably *structured* $\mathbb{Z}$-basis of the ring, i.e., a set of ring elements for which every $x \in R$ can be written uniquely as an integer linear combination of these elements. Equivalently, viewing R as a lattice, one can see these elements simply as short vectors that form a basis of the lattice. Here "short" is typically measured in the canonical embedding of the ring, and "structured" is elaborated upon below.

Knowledge of *some* basis of R is needed for merely representing and operating on ring elements. Moreover, a *short* basis is needed for removing error in decryption, which recovers some "noisy" encoding of the plaintext, and decodes it using the short basis. Finally, for computational efficiency (see Section 1.2.3 below for details), it is advantageous to have a *structured* basis, i.e., one that is the Kronecker (or tensor) product of *relative* bases going down a tower of intermediate subrings of small relative degree at each step. For example, the "powerful" basis (so named in [LPR13]) of an arbitrary cyclotomic is excellent in all these respects: it consists of optimally short ring elements, and it is the tensor product of relative bases going down a tower of cyclotomic rings of minimal relative degrees.

Our contribution. In the second part of this work [PP25], we construct two kinds of short, structured, and efficiently computable bases for a wide class of abelian number rings, namely, *any subfield* of any *power-of-*p cyclotomic for prime p, or the *composite* of such subfields for distinct primes. The first kind of basis is highly structured as a tensor product going down a tower (like the powerful basis), and has length within a $\sqrt{n}$ factor of optimal in its number field, where n is the degree of the field. Moreover, its length is within a $\sqrt{dn}$ factor of the best possible in *any* degree-n number field, where $d \mid (p-1)$ measures "how far" the field is from cyclotomic (formally, d is the relative degree of the smallest extension field that is cyclotomic). The second kind of a basis has less structure, though still enough to support at least one kind of fast CRT transform, and both it and its dual are *optimally short* for their number field. Finally, by tensoring we immediately get short, structured bases for the *composites* of any number of prime-power cyclotomic subfields, for distinct primes.

Our constructions build on ideas from, and significantly generalize, the work of [Bre97], which gave bases (that happen to be short) for the number rings of cyclotomic subfields over the *integers* $\mathbb{Z}$, or more generally, relative bases over *cyclotomic* base rings. However, this is not sufficient for our purposes, because to obtain the desired structure we need relative bases over *non-cyclotomic* number rings. Indeed, within any non-cyclotomic subfield of a prime-power cyclotomic, *all* of its subfields are non-cyclotomic (except for $\mathbb{Q}$). Therefore, the main results from [Bre97] do not provide any nontrivially structured $\mathbb{Z}$-bases for our desired fields, just "monolithic" bases over $\mathbb{Z}$.

We point out that *some* limitation on the abelian number ring or underlying tower structure is necessary for constructing relative bases (whether short or not), because some abelian number-ring extensions *do not have a relative basis at all*, i.e., they are not *free* as modules over their base rings.[5] Our limitation is a mild one that supports a natural approach for choosing a suitable ring: select a suitable subfield of the power-of-p cyclotomics for various distinct primes p, then tensor the results together to get a large enough dimension and number of slots.[6]

Our constructions of short, structured bases are technically heavy, but for cryptographic applications, only

[5] A simple non-free example arises from certain subfields of the 15th cyclotomic $L = \mathbb{Q}(\zeta_{15})$. Then $\mathrm{Gal}(L/\mathbb{Q}) \cong \mathbb{Z}_{15}^* \cong \mathbb{Z}_3^* \times \mathbb{Z}_5^*$. Let $H := \langle (1,-1) \rangle \subseteq H' := \langle (-1,2) \rangle$ be the (multiplicative) cyclic subgroups generated by $(1,-1)$ and $(-1,2)$, respectively; note that $(1,-1) = (-1,2)^2$. Then letting $K' := L^{H'} \subseteq K := L^H$ respectively be the fixed fields of H' and H, the ring of integers $\mathcal{O}_K$ turns out to be *non-free* as a module over $\mathcal{O}_{K'}$, i.e., it has no $\mathcal{O}_{K'}$-basis.

[6] Indeed, to get slots that are isomorphic to a desired prime field $\mathbb{F}_r$, there is even a "best possible" subfield (yielding the most slots) of the power-of-p cyclotomics for each prime $p \neq r$; see Lemma A.1.

their geometric norms and Kronecker-product structure are relevant. Therefore, we have separated the details of this contribution into the second part of this work [PP25]. See Section 2.3 for the formal statements of the constructions, their relevant properties, and some further details.

1.2.3 Fast Ring Arithmetic via Structured CRT Bases and Transforms

In ring-based cryptography over cyclotomics, the *Chinese Remainder Theorem (CRT) representation* is an important and widely used feature enabling efficient ring arithmetic modulo certain integers. In this representation, both addition and multiplication of (quotient-)ring elements respectively correspond to coordinate-wise addition and multiplication of their CRT-coefficient vectors, which is very fast. In addition, there are fast algorithms that map between the CRT representation and other bases that are used for various purposes, like sampling errors and decryption. These *CRT transform* algorithms are closely related to the Number Theoretic Transform (NTT), which is a finite-field variant of the Fast Fourier Transform (FFT). Specialized fast CRT transforms were given for, e.g., power-of-two cyclotomics in [LMPR08], and for arbitrary cyclotomics in [LPR13].

A second important application of the CRT representation is for efficient *bootstrapping of "packed" ciphertexts*—i.e., those that encrypt a large amount of plaintext data—as initially proposed in [GHS12a]. In one of two main parts of packed bootstrapping, we need to *homomorphically evaluate the CRT transform* (and its inverse) efficiently, using the FHE scheme's "native" homomorphic operations.[7] This homomorphically moves the "noisy decryption coefficients" into the CRT slots for SIMD noise removal, and then back again (see Section 3.2 for further details). Efficient homomorphic CRT transforms were given in [GHS12a] (and concretely implemented in [HS15]) based on automorphisms, and in [AP13] (implemented in [CPS18]) based on ring/field-switching [GHPS12].

Our contribution. The mathematical theory underlying CRT representations in cyclotomics holds more generally for arbitrary *abelian* (Galois) number-field extensions, and in particular for cyclotomic subfields. In Section 5 we build on this theory for computational and cryptographic purposes. First, we define the *CRT basis* of any abelian extension of number rings (modulo a suitable ideal), and derive some of its key structural properties. Most importantly, any CRT basis factors as the tensor product of relative CRT bases going down any tower of intermediate number rings (see Lemma 5.11). In addition, any CRT basis can be "lifted" or "lowered" to a "parallel" abelian extension, according to the fundamental Galois correspondence (see Lemma 5.7).

We then use the factorization of CRT bases to give fast CRT-transform algorithms that map between the CRT basis and *any other* similarly structured basis (including the short ones described above)—both "in the clear" for basic ring arithmetic, and homomorphically for packed bootstrapping. The former algorithms work directly on coordinate vectors (relative to the source and target bases), and immediately yield fast addition and multiplication in general abelian number rings. But for homomorphic evaluation, native homomorphic operations on ciphertexts do not support direct manipulation of plaintexts' coordinate vectors, so a different approach is needed. Using the general framework described next in Section 1.2.4, we show that CRT transforms can also be expressed in terms of *relatively few automorphisms*, which allows them to be efficiently evaluated homomorphically.

Interestingly, although the tensor-product form of the CRT basis is essential to both kinds of CRT-transform algorithms, they work in different ways, and each one relies on a different extra feature of the CRT factors.

[7] The other main part is a nonlinear "rounding" operation that is applied to all of the SIMD slots in parallel. This has known solutions (e.g., [GHS12a, AP13, CH18, GIKV23]) that are independent of the linear part, so we do not consider it further in this work.

Most notably, the "in the clear" algorithm is best run "bottom up" (see Section 5.4), whereas a "top down" evaluation is needed when using automorphisms (see Sections 1.2.4 and 4).

As a related contribution, we give tools for finding cyclotomic subfields that have desired features. Specifically, for a given "slot type"—e.g., a certain prime-power finite field—these tools give cyclotomic subfields that have a desired number of CRT slots of *exactly* that type, with no "wasted capacity." See Appendix A for details and Figure 1 for examples.

1.2.4 Homomorphic Structured Transforms via Automorphisms, for Bootstrapping

As mentioned above in Section 1.2.3, [GHS12a] expresses CRT transforms on certain cyclotomics in terms of *relatively few* automorphisms. Because ring-based FHE schemes support automorphisms as a native homomorphic operation, this immediately yields efficient homomorphic evaluation of CRT transforms, which is one of two main steps in packed bootstrapping. (In the realm of approximate FHE, [CCS19] did similarly for CRT transforms over the *complex* numbers, in power-of-two cyclotomics.[8]) Subsequent work [HS14, HS15, HS18] improved and generalized these ideas to build a flexible toolkit for homomorphic evaluation of various linear transforms and related linear-algebraic algorithms, but limited to cyclotomics.

Our contribution. In Section 4 we give analogous tools for expressing "structured" linear transforms in terms of automorphisms, in *arbitrary* (finite) Galois extensions, via a simple and general framework. We build upon the standard fact that in any such extension L/K, any K-linear function can be expressed as an L-linear combination of the automorphisms. (See Lemma 4.1.) For efficient homomorphic evaluation, we want this linear combination to be "sparse," i.e., to use only a small number of automorphisms. (Each automorphism has a moderate cost to evaluate homomorphically, because it involves a key switch.)

We achieve this goal by a combination of two techniques. First, we focus on *structured* linear transforms that map between bases having tensor-product factorizations going down a tower, like our short and CRT bases. As with the "in the clear" CRT transforms described above, this leads to a corresponding sparse decomposition of the transform, which maps each factor of one basis to its counterpart in the other, in sequence (see Equation (4.3)). Sparsely mapping a factor "high" in the tower is immediate, because this corresponds to a linear function on a low-degree extension. But in general we cannot map "low" factors sparsely, without changing the higher factors as a side effect. Yet amazingly, when the high factors form a *CRT basis*, it turns out that we can preserve them while sparsely mapping the "low" factors! (See Lemmas 4.3 and 5.13.) So overall, we can map between the CRT basis and any other similarly structured basis, using few automorphisms.

This sparse-decomposition perspective is also useful more broadly, for homomorphic linear algebra and other algorithms (cf. [HS14, HS18]). Notably, the tensor-product form of the CRT basis enables flexible *data movement* among SIMD slots. More specifically, the slots can be seen as arranged in a multidimensional array (or tensor), whose shape matches the factorization of the Galois group into a product of subgroups; each subgroup then acts independently and transitively along its own dimension of the array. (See Remark 5.12.) So, beyond having optimal SIMD packing via a desired slot "type," one can also design a ring so that the slots are arranged in a desired "shape," to support the application's specific needs. As an illustration of this flexibility and its tradeoffs, in Appendix B we give two example parameterizations for homomorphic AES evaluation, whose array of slots has a few moderate dimensions in one case, and several small dimensions in the other.

[8] The second part of this work [PP25] gives a fine-grained tensor-product factorization of a CRT-like basis (over $\mathbb{R}$ or $\mathbb{C}$) for the number field's *canonical embedding*. Our sparse-transform framework can be slightly adapted to this setting to recover the homomorphic CRT transform of [CCS19], along with analogous ones for non-power-of-two cyclotomics and cyclotomic subfields.

1.3 Guide to the Rest of the Paper

For the reader's convenience, here we summarize the structure, contents, and dependencies for the remainder of the paper.

- Section 2 gives the mathematical preliminaries, covering the necessary Galois theory in Section 2.1, the needed algebraic number theory in Section 2.2 (which by now is mostly standard in the lattice cryptography literature), and the results we need from [PP25] in Section 2.3.
- Section 3 abstracts out (from [BGV12, Bra12, FV12, LPR13, CKKS17]) a general template for ring-based homomorphic encryption that works over the ring of integers in *any* number field, highlighting the computational aspects that need to be addressed. This template can be understood with just the basics of algebraic number theory from Section 2.2.
- Section 4 lays out a framework for expressing linear functions on arbitrary (finite) Galois extensions—and Galois number fields in particular—as linear combinations of their automorphisms, which can be evaluated homomorphically. We also give sufficient conditions that yield *sparse* decompositions, in terms of relatively few automorphisms. This framework can be understood with just the background on Galois theory from Section 2.1.
- Section 5 defines the Chinese Remainder Theorem (CRT) basis of an arbitrary abelian extension of number rings, modulo a suitable ideal. We factor the CRT basis as the Kronecker product of CRT bases going down any tower of intermediate number rings, and use this to obtain two kinds of sparse decompositions of CRT transforms: one that works directly on coefficient vectors "in the clear," and (using the framework from Section 4) one in terms of automorphisms. The material in this section can be understood with the background on algebraic number theory from Section 2.2, especially Section 2.2.3.
- Section 6 uses our tools to instantiate the homomorphic encryption template from Section 3 computationally, with fast algorithms. This material relies on the details of the template, and the main results from Sections 4 and 5.
- Appendix A characterizes the number and type of finite-field slots that can be obtained in abelian number fields of prime-power conductor and their composites, and provides several numerical examples. Appendix B gives various choices of abelian number fields that support homomorphic AES evaluation with no wasted SIMD "capacity," and compares them to the cyclotomic field used in [GHS12c].

2 Preliminaries

In this work, all rings are implicitly commutative with identity. For a ring R, a function f from an R-module to an R-module is *R-linear* if $f(a+b) = f(a) + f(b)$ and $f(r \cdot a) = r \cdot f(a)$ for all $r \in R$ and all a, b.

Vectors and matrices. We denote column vectors by lower-case letters that either have an arrow, like $\vec{a}$, or sometimes are in boldface, like $\mathbf{a}$ (so $\mathbf{a}^t$ and $\vec{a}^t$ are row vectors). We use the former for general domains, and the latter only for vectors with real or complex entries (possibly modulo some integer). The entries a_i of a vector $\vec{a}$ are indexed by $i \in I$ for some specified finite *index set* I; similarly, the entries $A_{i,j}$ of a matrix A are indexed by $(i,j) \in I \times J$ for row and column index sets I and J, respectively. Often in this work, an index set is not of the form $\{1, \ldots, n\}$, but is some other finite structure. We often apply functions to vectors or matrices, which means element-wise application of the function.

For matrices (including vectors as a special case) A and B over a ring, and having respective index sets $I \times J$ and $I' \times J'$, their *Kronecker product* $A \otimes B$ is the matrix having index set $(I \times I') \times (J \times J')$ whose entries are $(A \otimes B)_{(i,i'),(j,j')} = a_{i,j} \cdot b_{i',j'}$. A central fact about the Kronecker product is the *mixed-product property*, which says that for matrices A, B as above and C, D having respective index sets $J \times K, J' \times K'$, we have that $(A \otimes B) \cdot (C \otimes D) = (AC) \otimes (BD)$, which has index set $(I \times I') \times (K \times K')$.

Group actions. A *group action* for a group G and a set S, called a *G-set*, is a function $\star\colon G \times S \to S$, typically used as an infix operator, that satisfies $e \star s = s$ where $e \in G$ is the identity, and $g \star (h \star s) = (gh) \star s$. It is *free* if $g \star s = s$ for some $s \in S$ implies that $g = e$. It is *transitive* if for any $s, s' \in S$, there exists some $g \in G$ such that $g \star s = s'$. Finally, it is *regular* if it is both free and transitive. When a particular (free and/or transitive) action is clear by context, we often say that *G acts (freely and/or transitively) on S*, or that S is *acted upon (freely and/or transitively)* by G. For brevity, when G acts freely on S, we say that S is *G-normal*, and if G also acts transitively on S, we say that S is *G-regular*.

2.1 Field and Galois Theory

A *field extension* L/K is a pair of fields $K \subseteq L$ where the ring operations on K coincide with those of L when restricted to K. The field L is a vector space over K, and the *degree* $\deg(L/K)$ is defined as the dimension of this space. *All extensions considered in this work are implicitly of finite degree.*

For subfields L_1, L_2 some common field M, their *composite* field L_1L_2 (also sometimes known as their *compositum*) is the subfield $L_1L_2 = \{\sum_{i=1}^{r} \alpha_i\beta_i : \alpha_i \in L_1, \beta_i \in L_2, \text{finite } r\} \subseteq M$.

2.1.1 Automorphisms and Galois Extensions

An *automorphism* of a field extension L/K is a ring isomorphism $\tau\colon L \to L$ that fixes K pointwise, i.e., $\tau(a) = a$ for all $a \in K$. The *Galois group* $\mathrm{Aut}(L/K)$ is the group of all such automorphisms, with function composition as the group operation. A *Galois extension* is one for which $|\mathrm{Aut}(L/K)| = \deg(L/K)$, and its Galois group is usually denoted $\mathrm{Gal}(L/K)$. For concision, a Galois extension whose Galois group is abelian (or cyclic, etc.), is simply said to be abelian (or cyclic, etc.). *In this work we typically work with abelian Galois groups* (even though this is not required for a few select results), so throughout this overview the reader may wish to focus on that case.

Any Galois group $G = \mathrm{Gal}(L/K)$ acts on L in the natural way, via $\tau \star x = \tau(x)$. Therefore, it also acts on any subset of L (or collection of such subsets) that is closed under G. This action is not necessarily free, since $\tau(x) = x$ for any $\tau \in G$ and $x \in K$, but for certain (collections of) subsets of L it may be.

The fundamental theorem of Galois theory says that for any (finite) Galois extension L/K, there is a bijective correspondence between its *intermediate* fields, also known as *subextensions*—i.e., those fields F satisfying $K \subseteq F \subseteq L$—and the subgroups of $\mathrm{Gal}(L/K)$. Specifically, for any intermediate field F of L/K, the corresponding subgroup is $\mathrm{Gal}(L/F)$, i.e., the automorphisms of L/K that fix F pointwise. In the reverse direction, for any subgroup $H \subseteq \mathrm{Gal}(L/K)$, the corresponding intermediate field is $L^H := \{a \in L : \tau(a) = a \ \forall \tau \in H\}$, the subfield of L that is fixed pointwise by every automorphism in H.

It is easy to see that the above correspondence is *inclusion reversing*, i.e., for any intermediate fields F_1, F_2 of L/K, we have that $\mathrm{Gal}(L/F_1) \subseteq \mathrm{Gal}(L/F_2)$ if and only if $F_1 \supseteq F_2$. From this it follows that the intersection and join of (i.e., subgroup generated by) their Galois groups are, respectively,

$$\begin{aligned}\mathrm{Gal}(L/F_1) \cap \mathrm{Gal}(L/F_2) &= \mathrm{Gal}(L/(F_1F_2)) \\ \langle \mathrm{Gal}(L/F_1), \mathrm{Gal}(L/F_2) \rangle &= \mathrm{Gal}(L/(F_1 \cap F_2)) \ .\end{aligned}$$

In particular, if $L = F_1F_2$ and both F_1, F_2 are Galois over $F_1 \cap F_2$, then the join is the (internal) direct product: $\mathrm{Gal}(L/(F_1 \cap F_2)) = \mathrm{Gal}(L/F_1) \times \mathrm{Gal}(L/F_2)$.

2.1.2 Trace and Duality

For a Galois extension L/K, the *trace* $\mathrm{Tr}_{L/K} \colon L \to K$ is merely the sum of the automorphisms:

$$\mathrm{Tr}_{L/K}(x) := \sum_{\tau \in \mathrm{Gal}(L/K)} \tau(x) \in K \ .$$

By definition, this is K-linear, and the output is in K because it is fixed by any element of $\mathrm{Gal}(L/K)$.

Let $\vec{b}$ be a vector over L of K-linearly independent entries, with index set I. Then a vector $\vec{b}^\vee$ over L, also with index set I, is *dual* to $\vec{b}$ if

$$\mathrm{Tr}_{L/K}(b_i^\vee \cdot b_{i'}) = \delta_{i,i'} := \begin{cases} 1 & \text{if } i = i' \\ 0 & \text{otherwise.} \end{cases}$$

Clearly this is symmetric, i.e., $\vec{b}$ is dual to $\vec{b}^\vee$ as well. Such a $\vec{b}^\vee$ always exists, and is *unique* if $\vec{b}$ is a K*-basis* of L, in which case we call $\vec{b}^\vee$ *the dual basis* of $\vec{b}$. When the extension L/K may not be clear from context (e.g., when working with towers of extensions), we may write $\vec{b}^{\vee_{L/K}}$ in place of $\vec{b}^\vee$, to emphasize that it is defined using the trace from L to K.

Fixing a K-basis $\vec{b}$ of L, any $x \in L$ can be written uniquely as $x = \langle \vec{b}, \vec{x} \rangle = \vec{b}^t \cdot \vec{x}$ for some coefficient vector $\vec{x}$ over K. The dual basis directly yields this vector, as $\vec{x} = \mathrm{Tr}_{L/K}(\vec{b}^\vee \cdot x)$. This is because by K-linearity of $\mathrm{Tr}_{L/K}$ and by definition of $\vec{b}^\vee$, for the coefficient vector $\vec{x}$ over K of $x = \vec{b}^t \cdot \vec{x} \in L$,

$$\mathrm{Tr}_{L/K}(\vec{b}^\vee \cdot x) = \mathrm{Tr}_{L/K}(\vec{b}^\vee \cdot \vec{b}^t \cdot \vec{x}) = \mathrm{Tr}_{L/K}(\vec{b}^\vee \cdot \vec{b}^t) \cdot \vec{x} = \vec{x} \ .$$

In particular, the dual basis lets us transform to basis $\vec{b}$ from any K-basis $\vec{a}$ of L: we have that $\vec{a}^t = \vec{b}^t \cdot T$ where $T = \mathrm{Tr}_{L/K}(\vec{b}^\vee \cdot \vec{a}^t)$ is the change-of-basis matrix from $\vec{a}$ to $\vec{b}$, so $\vec{a}^t \cdot \vec{x} = \vec{b}^t \cdot (T\vec{x})$ for any coefficient vector $\vec{x}$ over K.

2.1.3 Towers of Extensions

If M/L and L/K are field extensions, then we often write $M/L/K$ as a *tower*; recall that L is called an intermediate field (or subextension) of M/K. The trace map is *transitive* on any such tower: $\mathrm{Tr}_{M/K} = \mathrm{Tr}_{L/K} \circ \mathrm{Tr}_{M/L}$.

Suppose that the entries of a vector $\vec{b}_1$ over M are L-linearly independent, and the entries of a vector $\vec{b}_2$ over L are K-linearly independent. Then the entries of the Kronecker product $\vec{b}_1 \otimes \vec{b}_2$ (which is a vector over M) are K-linearly independent, and $(\vec{b}_1 \otimes \vec{b}_2)^{\vee_{M/K}} = \vec{b}_1^{\vee_{M/L}} \otimes \vec{b}_2^{\vee_{L/K}}$.

Towers of Galois extensions. If M/K is Galois, then M/L is Galois, but L/K is Galois if and only if the subgroup $\mathrm{Gal}(M/L) \subseteq \mathrm{Gal}(M/K)$ is normal. In particular, if M/K is abelian, then both M/L and L/K are abelian as well. We have the following standard correspondence.

Lemma 2.1. *Let $M/L/K$ be a tower where M/K and L/K are Galois. Then $\mathrm{Gal}(M/K)/\mathrm{Gal}(M/L) \cong \mathrm{Gal}(L/K)$ via restriction to L.*

If L_1 and L_2 are Galois extensions of some base field (all contained in some common field), then both their composite L_1L_2 and their intersection $L_1 \cap L_2$ are also Galois over that base field.

Lemma 2.2. *Let L_1, L_2 be Galois extensions of $K = L_1 \cap L_2$, and let $M = L_1L_2$. Then $\mathrm{Gal}(M/L_1) \cong \mathrm{Gal}(L_2/K)$ via restriction to L_2. In particular, Tr_{M/L_1} restricted to L_2 is $\mathrm{Tr}_{L_2/K}$.*

Proof. We have that

$$\mathrm{Gal}(M/L_1) \cong (\mathrm{Gal}(M/L_1) \times \mathrm{Gal}(M/L_2))/\mathrm{Gal}(M/L_2) = \mathrm{Gal}(M/K)/\mathrm{Gal}(M/L_2) \cong \mathrm{Gal}(L_2/K) ,$$

where the first isomorphism is the natural homomorphism, and the last one is via restriction to L_2, by Lemma 2.1. □

2.2 Algebraic Number Theory

2.2.1 Number Fields

A *number field* K is a finite-degree field extension of the rationals $\mathbb{Q}$. Concretely, it can always be represented as $K = \mathbb{Q}(\gamma) \cong \mathbb{Q}[x]/f(x)$, where $f(x)$ is the minimal polynomial over $\mathbb{Q}$ of γ, i.e., the unique monic polynomial over $\mathbb{Q}$ of least degree for which $f(\gamma) = 0$. As a field extension, K can be Galois (and abelian or even cyclic), have extensions or subextensions, etc., according to the above conditions.

We endow K with a geometry in the standard way via its *canonical embedding* $\boldsymbol{\sigma} \colon K \to \mathbb{C}^{E_K}$, which is the concatenation of its set E_K of $n = \deg(K/\mathbb{Q})$ ring embeddings $\sigma \colon K \to \mathbb{C}$. This makes K a complex inner-product space via the *canonical (Hermitian) inner product*[9]

$$\langle \alpha, \beta \rangle = \langle \alpha, \beta \rangle_K := \langle \boldsymbol{\sigma}(\alpha), \boldsymbol{\sigma}(\beta) \rangle = \sum_{\sigma \in E_K} \overline{\sigma(\alpha)} \cdot \sigma(\beta) ,$$

with the standard Euclidean norm $\|\alpha\| = \|\alpha\|_K := \sqrt{\langle \alpha, \alpha \rangle} = \|\boldsymbol{\sigma}(\alpha)\|$; we call this the *canonical norm* of K. We also extend this to $\|\vec{x}\| = \max_i \|x_i\|$ for any vector $\vec{x}$ over K.

2.2.2 Rings of Integers and Ideals

The *ring of integers* (or *number ring*) of K, denoted $\mathcal{O}_K$, is the ring of all algebraic integers (i.e., roots of monic polynomials with integer coefficients) in K. The ring of integers is a free $\mathbb{Z}$-module of rank $n = \deg(K/\mathbb{Q})$, and thus has a $\mathbb{Z}$-basis $\vec{b}$ consisting of n elements of $\mathcal{O}_K$. For a number field extension L/K, an element of L is in $\mathcal{O}_L$ if and only if its minimal polynomial over K has coefficients in $\mathcal{O}_K$ (see, e.g., [Mat89, Theorem 9.2]).

A (nonzero) *ideal* of $\mathcal{O}_K$ is a (nontrivial) additive subgroup $\mathfrak{a} \subseteq \mathcal{O}_K$ that is closed under multiplication by $\mathcal{O}_K$, i.e., $\mathcal{O}_K \cdot \mathfrak{a} \subseteq \mathfrak{a}$; indeed, this is an equality because $1 \in \mathcal{O}_K$. A *fractional* ideal (of $\mathcal{O}_K$) is a set $\mathfrak{a} \subseteq K$ such that $d\mathfrak{a}$ is an ideal (of $\mathcal{O}_K$) for some $d \in \mathcal{O}_K$. For convenience, throughout this work *we implicitly restrict all (fractional) ideals to be nonzero*, unless stated otherwise.

The product of (fractional) ideals $\mathfrak{a}, \mathfrak{b}$ is defined as the set of all finite *sums* of terms ab for $a \in \mathfrak{a}, b \in \mathfrak{b}$. The set of fractional ideals of $\mathcal{O}_K$ forms a group under multiplication with $\mathcal{O}_K$ as its identity element; the multiplicative inverse of $\mathfrak{a}$ is denoted $\mathfrak{a}^{-1}$. Ideals $\mathfrak{a}, \mathfrak{b}$ of $\mathcal{O}_K$ are *relatively prime*, also called *coprime*, if $\mathfrak{a} + \mathfrak{b} = \mathcal{O}_K$. Because $\mathcal{O}_K$ is commutative, the product of any (finite number of) pairwise coprime ideals is equal to their intersection.

[9] For convenience and consistency with later definitions, we arbitrarily define the inner product to be linear in its second argument and *conjugate* linear in its first argument.

Chinese Remainder Theorem. The *Chinese Remainder Theorem* says that for any (finite) collection of pairwise coprime ideals $\mathfrak{a}_i$ of $\mathcal{O}_K$ defining $\mathfrak{a} = \prod_i \mathfrak{a}_i$, the natural ring homomorphism from the quotient ring $\mathcal{O}_K/\mathfrak{a}$ to the product of quotient rings $\prod_i (\mathcal{O}_K/\mathfrak{a}_i)$ is in fact an *isomorphism*. So, there is a unique *CRT vector* $\vec{c}$ over $\mathcal{O}_K/\mathfrak{a}$, indexed by the i, for which $\vec{c} = \vec{e}_i \pmod{\mathfrak{a}_i}$, where $\vec{e}_i$ is the vector with 1 at index i and 0 elsewhere. By definition, the forward direction of the isomorphism maps $x \in \mathcal{O}_K/\mathfrak{a}$ to the tuple $\vec{x}$ whose ith entry is $x_i = x + \mathfrak{a}_i \in \mathcal{O}_K/\mathfrak{a}_i$, and the reverse direction maps such a tuple to $\langle \vec{c}, \vec{x} \rangle = \sum_i c_i x_i \in \mathcal{O}_K/\mathfrak{a}$. (Observe that because $c_i = 0 \pmod{\mathfrak{a}_{i'}}$ for all $i' \neq i$, and $x_i \in \mathcal{O}_K/\mathfrak{a}_i$, each product $c_i x_i \in \mathcal{O}_K/\mathfrak{a}$.)

2.2.3 Prime Ideals and Splitting

In $\mathcal{O}_K$, an ideal $\mathfrak{p}$ is *prime* if and only if it is *maximal*, i.e., $\mathfrak{p} \neq \mathcal{O}_K$ and there does not exist any ideal $\mathfrak{r}$ of $\mathcal{O}_K$ such that $\mathfrak{p} \subsetneq \mathfrak{r} \subsetneq \mathcal{O}_K$. In this case, the quotient ring $\mathcal{O}_K/\mathfrak{p}$ is isomorphic to a finite field, called the *residue field* of $\mathfrak{p}$. A fundamental theorem is that the ring $\mathcal{O}_K$ has unique factorization of ideals into prime ideals, i.e., it is a Dedekind domain. Therefore, $\mathfrak{a} \mid \mathfrak{b}$ (i.e., there exists an ideal $\mathfrak{c}$ such that $\mathfrak{a}\mathfrak{c} = \mathfrak{b}$) if and only if $\mathfrak{a} \supseteq \mathfrak{b}$ for ideals $\mathfrak{a}, \mathfrak{b}$. For an extension L/K of number fields, an ideal $\mathfrak{p}'$ of $\mathcal{O}_L$ is said to *lie over* the ideal $\mathfrak{p} = \mathfrak{p}' \cap \mathcal{O}_K$ of $\mathcal{O}_K$. If $\mathfrak{p}'$ is prime, then so is $\mathfrak{p}$ (but not necessarily vice versa).

Now let L/K be an abelian extension. Then any prime ideal $\mathfrak{p}$ of $\mathcal{O}_K$ "splits" in $\mathcal{O}_L$ into equal-exponent powers of the prime ideals $\mathfrak{p}_\ell$ lying over $\mathfrak{p}$ in $\mathcal{O}_L$, in the following way.[10] Let $G = \mathrm{Gal}(L/K)$ and

$$D = D_{L/K}(\mathfrak{p}_\ell) := \{\tau \in G : \tau(\mathfrak{p}_\ell) = \mathfrak{p}_\ell\}$$
$$I = I_{L/K}(\mathfrak{p}_\ell) := \{\tau \in D : \tau(a) = a \pmod{\mathfrak{p}_\ell} \; \forall \, a \in \mathcal{O}_L/\mathfrak{p}_\ell\}$$

respectively be the common *decomposition group* of all automorphisms that fix some arbitrary $\mathfrak{p}_\ell$, and the common *inertia group* of all those automorphisms that induce the identity map on $\mathcal{O}_L/\mathfrak{p}_\ell$.[11] Then the $\mathfrak{p}_\ell$ can be indexed by $\ell \in G/D$, so there are $|G/D|$ of them, and the Galois group G acts on them by $\tau(\mathfrak{p}_\ell) = \mathfrak{p}_{\tau \circ \ell}$. Therefore, this action is transitive, and is free if and only if D is the trivial group. The splitting of $\mathfrak{p}$ in $\mathcal{O}_L$ is given by the factorization

$$\mathfrak{p}\mathcal{O}_L = \prod_{\ell \in G/D} \mathfrak{p}_\ell^e \tag{2.1}$$

into the product of $g = |G/D|$ distinct factors, where the common exponent $e = |I|$ is called the *ramification index* of $\mathfrak{p}$ in L. Letting $\mathbb{F}_p \cong \mathcal{O}_K/\mathfrak{p}$ denote the residue field of $\mathfrak{p}$ for some prime-power p, each residue field $\mathcal{O}_L/\mathfrak{p}_\ell \cong \mathbb{F}_{p^f}$, where the *residue degree* $f = |D/I|$ and hence $efg = |G| = \deg(L/K)$.

SIMD slots from ideal splitting. In the context of FHE, the factorization from Equation (2.1) and the Chinese Remainder Theorem together form the foundation for plaintext "SIMD slots." Letting the plaintext ring be $\mathcal{O}_L/\mathfrak{p}\mathcal{O}_L$, it is isomorphic to the product of the g rings $\mathcal{O}_L/\mathfrak{p}_\ell^e$, in which both addition and multiplication work component-wise ("single instruction, multiple data"), so these component rings are called *slots*. Assuming for now that $e = 1$, all the slots are isomorphic to the residue field $\mathbb{F}_{p^f}$, which is called their *type*; recall from above that $fg = \deg(L/K)$, so larger f means smaller g, and vice versa. Often, we want slots that support arithmetic in $\mathbb{F}_p \cong \mathcal{O}_K/\mathfrak{p}$ (e.g., $\mathbb{F}_p \cong \mathbb{Z}_p$ for $K = \mathbb{Q}$ and prime integer p). This is achieved with no "waste" if $f = 1$ and thus $g = \deg(L/K)$; otherwise, $\mathbb{F}_p$ is a strict subfield of the slot type, so some of the extension's degree is "wasted" on a larger slot type, and fewer slots.

[10] The same holds if L/K is merely Galois (but not abelian), but with a somewhat more complicated formalization. It even holds if L/K is not Galois, except that the exponents of the $\mathfrak{p}_\ell$ may vary.

[11] The fact that all the $\mathfrak{p}_\ell$ have the same decomposition group and inertia group is implied by the hypothesis that L/K is abelian.

Ramification, inertness, and splitting. Returning to the mathematical theory, if $e = 1$ and hence $fg = \deg(L/K)$, then $\mathfrak{p}$ is said to be *unramified* in L. In this case, if D is trivial, and hence $f = 1$ and $g = |G| = \deg(L/K)$, then $\mathfrak{p}$ is said to *split completely* in L; at the opposite extreme, if $D = G$ and hence $f = \deg(L/K)$ and $g = 1$, then $\mathfrak{p}$ is said to be *inert* in L. If $e > 1$, then $\mathfrak{p}$ is said to *ramify* in L, and the $\mathfrak{p}_\ell$ are said to ramify over K. If $e = \deg(L/K)$, i.e., $I = D = G$ and hence $f = g = 1$, then $\mathfrak{p}$ is said to be *totally ramified* in L.

The decomposition and inertia groups allow for decomposing L/K as a tower of extensions $L/L^I/L^D/K$ in which (total) ramification, inertness, and (complete) splitting each can be isolated:

- At the top, L^I is the *smallest* intermediate field F of L/K in which the $\mathfrak{p}_\ell \cap F$ are totally ramified in L. So, $(\mathfrak{p}_\ell \cap L^I)\mathcal{O}_L = \mathfrak{p}_\ell^e$, where $e = |I| = \deg(L/L^I)$, and $\mathfrak{p}$ is unramified in L^I.
- In the middle intermediate extension L^I/L^D, the $\mathfrak{p}_\ell \cap L^D$ are inert: they split as $(\mathfrak{p}_\ell \cap L^D)\mathcal{O}_{L^I} = \mathfrak{p}_\ell \cap L^I$, with relative residue degree $f = |D/I| = \deg(L^I/L^D)$.
- At the bottom, L^D is the *largest* intermediate field of L/K in which $\mathfrak{p}$ splits completely: it splits into the product of the $g = |G/D| = \deg(L^D/K)$ prime ideals $\mathfrak{p}_\ell \cap L^D$ (which do not split any further in L^I or L, as already noted).

2.2.4 Duality for Ideals and the (Co)different

Let L/K be an extension of number fields. For a fractional ideal $\mathfrak{b}$ of $\mathcal{O}_L$, its *dual* relative to K, which is also a fractional ideal of $\mathcal{O}_L$, is defined as

$$\mathfrak{b}^{\vee_{L/K}} := \{x \in L : \mathrm{Tr}_{L/K}(x \cdot \mathfrak{b}) \subseteq \mathcal{O}_K\} .$$

When the field extension L/K is clear from context, we often drop the subscript and simply write $\mathfrak{b}^\vee$. It is straightforward to verify from the definition that $(\mathfrak{b}^\vee)^\vee = \mathfrak{b}$, that $\mathfrak{b}^\vee$ is a fractional ideal of $\mathcal{O}_L$, and that if $\vec{b}$ is an $\mathcal{O}_K$-basis of $\mathfrak{b}$, then its dual $\vec{b}^\vee$ (as defined in Section 2.1.2) is an $\mathcal{O}_K$-basis of $\mathfrak{b}^\vee$.

(Co)different ideal. The dual ideal $\mathfrak{C}_{L/K} := \mathcal{O}_L^\vee$ of the ring of integers, called the *codifferent* of L/K, trivially contains $\mathcal{O}_L$. So, its inverse $\mathfrak{D}_{L/K} := \mathfrak{C}_{L/K}^{-1} \subseteq \mathcal{O}_L$ is an ideal of $\mathcal{O}_L$, which is called the *different* ideal of L/K. The codifferent relates the dual and inverse of any fractional ideal, as $\mathfrak{b}^\vee = \mathfrak{b}^{-1} \cdot \mathcal{O}_L^\vee$. When L/K is Galois, it is immediate that $\mathcal{O}_L^\vee$ is fixed by $\mathrm{Gal}(L/K)$, i.e., $\tau(\mathcal{O}_L^\vee) = \mathcal{O}_L^\vee$ for all $\tau \in \mathrm{Gal}(L/K)$, because $\mathcal{O}_L$ is. In addition, $\mathrm{Tr}_{L/K}(\mathcal{O}_L^\vee) = \mathcal{O}_K$.

Duality over quotients. Here we naturally extend the definition of duality (Section 2.1.2) to work over *quotients* of certain ideals. Letting $\mathfrak{q}$ be an ideal of $\mathcal{O}_K$ and $\mathfrak{b}$ be a fractional ideal of $\mathcal{O}_L$, and $\vec{b}$ be a vector over the quotient $\mathfrak{b}/\mathfrak{q}\mathfrak{b}$, we say that a vector $\vec{b}^\vee$ over $\mathfrak{b}^\vee/\mathfrak{q}\mathfrak{b}^\vee$ is dual to $\vec{b}$ (and symmetrically) if $\mathrm{Tr}_{L/K}(\vec{b}_i^\vee \cdot \vec{b}_{i'}) = \delta_{i,i'} \in \mathcal{O}_K/\mathfrak{q}$. Note that this congruence relation is valid because $\mathrm{Tr}_{L/K}(\mathfrak{b}^\vee \cdot \mathfrak{b}) = \mathrm{Tr}_{L/K}(\mathcal{O}_L^\vee) = \mathcal{O}_K$, and $\mathrm{Tr}_{L/K}$ is K-linear so the $\mathfrak{q}$ factor "passes through" it.

This extended notion of duality is consistent with, and inherits the properties of, the standard notion. If a vector $\vec{b}$ is over $\mathfrak{b}$ and $\vec{b}^\vee$ is over $\mathfrak{b}^\vee$ and is dual to $\vec{b}$, then $\vec{b} \bmod \mathfrak{q}\mathfrak{b}$ and $\vec{b}^\vee \bmod \mathfrak{q}\mathfrak{b}^\vee$ are duals in this new sense. Also, the material from Section 2.1.2 about using duals to extract coefficients and obtain change-of-basis matrices adapts straightforwardly to this setting. That is, if $\vec{b}$ is an $(\mathcal{O}_K/\mathfrak{q})$-basis of $\mathfrak{b}_\mathfrak{q} := \mathfrak{b}/\mathfrak{q}\mathfrak{b}$ and $\vec{b}^\vee$ is dual to $\vec{b}$, then any $c \in \mathfrak{b}_\mathfrak{q}$ can be written as $c = \vec{b}^t \cdot \mathrm{Tr}_{L/K}(\vec{b}^\vee \cdot x)$, so any vector $\vec{c}$ over $\mathfrak{b}_\mathfrak{q}$ can be written as $\vec{c}^t = \vec{b}^t \cdot T$ where $T = \mathrm{Tr}_{L/K}(\vec{b}^\vee \cdot \vec{c}^t)$.

2.3 Short Structured Bases

In the second part of this work [PP25], we construct (in two related ways) short and structured integral bases for arbitrary abelian number fields of prime-power conductor, and bound their canonical norms. Then, moving beyond prime-power conductors, we show that for any two abelian number fields having coprime conductors, the Kronecker product of any respective integral bases is an integral basis of the composite field, and the canonical norms of these basis elements are the products of the norms of their multiplicands. These high-level statements are sufficient for cryptographic applications, and are restated below.

For any positive integer m, the mth *cyclotomic field* is $\mathbb{Q}(\zeta_m)$ where ζ_m is a primitive mth root unity. Because the mth and $2m$th cyclotomic fields are isomorphic for odd m, we assume without loss of generality that $m \neq 2 \pmod 4$. We let $m^* = 2\,\mathrm{rad}(m)$ if $4 \mid m$, and $m^* = \mathrm{rad}(m)$ otherwise, where $\mathrm{rad}(m)$ is the product of all the primes that divide m (so $m^* \neq 2 \pmod 4$ as well). The Kronecker–Weber theorem states that a number field K is abelian if and only if it is a subfield of a cyclotomic field. In this case, its *conductor* is defined as the smallest positive integer m for which $K \subseteq \mathbb{Q}(\zeta_m)$.

Theorem 2.3 ([PP25]). *Let L be an abelian number field of prime-power conductor m, let $m = m_\ell > m_{\ell-1} > \cdots > m_0 = 1$ be such that $m^* \mid m_1$ and $m_{i-1} \mid m_i$ for all $1 \leq i \leq \ell$, and let $M_i = \mathbb{Q}(\zeta_{m_i})$ and $L_i = L \cap M_i$.*[12] *There is an efficiently computable $\mathbb{Z}$-basis $\vec{b} = \bigotimes_{i=1}^{\ell} \vec{b}_i$ of $\mathcal{O}_L$, where each $\vec{b}_i$ is an $\mathcal{O}_{L_{i-1}}$-basis of $\mathcal{O}_{L_i}$, and $\|\vec{b}\|^2 \leq \deg(M/L)^\ell \cdot \deg(L/\mathbb{Q}) \leq \deg(M/\mathbb{Q}) \cdot \deg(L/\mathbb{Q})$.*

The Kronecker-product structure of the bases from Theorem 2.3 directly yields "sparse decompositions" and corresponding fast algorithms for CRT transforms, both "in the clear" on coefficient vectors (see Section 5.4) and homomorphically, using automorphisms (see Section 4). The bound on $\|\vec{b}\|$ is within a $\sqrt{\deg(M/\mathbb{Q})}$ factor of the minimum distance of *any* number ring of the same degree, and within about a $\sqrt{\deg(L/\mathbb{Q})}$ factor of the largest successive minimum of $\mathcal{O}_L$. As compared with cyclotomics, this induces a (typically mild) cost in the noise tolerance in cryptographic applications, which affects the ultimate parameters (see Sections 3 and 6).

The next result shows that by using a different and slightly weaker structure, which is still sufficient for fast "in the clear" transforms, we can obtain *optimally short* integral bases. These make it possible to decode the rings of integers and their duals from larger error than using the bases from Theorem 2.3.

Theorem 2.4 ([PP25]). *Adopt the setup from Theorem 2.3. There is an efficiently computable $\mathbb{Z}$-basis $\vec{b} = \bigoplus_{i=1}^{\ell} \vec{b}_i$ of $\mathcal{O}_L$, where each $\vec{b}_i$ has some additional structure (see [PP25] for details) and $\bigoplus_{i=1}^{j} \vec{b}_i$ is a $\mathbb{Z}$-basis of $\mathcal{O}_{L_j}$. Moreover, if $m = m^*$, then $\|\vec{b}\|^2 = m^* - \deg(M/L)$; if $m > m^*$, then $\|\vec{b}\|^2 = \deg(M/\mathbb{Q})$; and (for $L \neq \mathbb{Q}$) in all cases $\|\vec{b}^\vee\|^2 = (\deg(M/L) + 1)/m$.*

In fact, the bases $\vec{b}$ from this construction are optimal in that they, and their duals $\vec{b}^\vee$, attain *all the successive minima* of the lattices they generate. Moreover, the norm of $\vec{b}$ is within a $\sqrt{\deg(M/L)} \leq \sqrt{\varphi(m^*)}$ factor of the minimum distance for *any* number ring of the same degree, and the norm of the dual basis is less than 1, which suffices for applications. Lastly, these bases have enough structure to support fast "in the clear" CRT transforms for ring arithmetic (building on Section 5.4). However, we do not yet know if they have sparse decompositions in terms of automorphisms, but this is relevant only for *homomorphic* linear transforms.

Finally, with our short structured integral bases for abelian number fields of any *prime-power* conductor in hand, we can use the Kronecker product to get such bases for their *composite* fields (which can have arbitrary conductors).

[12] Note that $M_\ell/M_{\ell-1}/\cdots/M_1/M_0$ is a tower of cyclotomics (with $M_0 = \mathbb{Q}$), hence we also have the tower $L_\ell/L_{\ell-1}/\cdots/L_1/L_0$.

Lemma 2.5 ([PP25]). *Let $L = L_1 L_2$ for abelian number fields L_1 and L_2 with coprime conductors. Then $\|x_1 \cdot x_2\|_L = \|x_1\|_{L_1} \cdot \|x_2\|_{L_2}$ for any $x_i \in L_i$, and if $\vec{b}_1$ and $\vec{b}_2$ are $\mathbb{Z}$-bases of $\mathcal{O}_{L_1}$ and $\mathcal{O}_{L_2}$, respectively, then $\vec{b}_1 \otimes \vec{b}_2$ is a $\mathbb{Z}$-basis of $\mathcal{O}_L$.*

3 Homomorphic Encryption Template

There are several homomorphic encryption schemes based on the Ring-LWE problem over rings in number fields, where the primary focus in the literature has been on *cyclotomic* rings. Adapting the presentation in [LPR13], here we abstract out a general template that works over the ring of integers in *any* number field, defining just the parts of the scheme that are relevant to this work. For our purposes, the prior schemes differ mainly in how they encode plaintexts and perform homomorphic multiplication, so the template applies equally well to the one of [BGV12], ones following the "scale invariant" methods of [Bra12, FV12], and the "approximate arithmetic" one of [CKKS17].

Throughout the template, we first describe clusters of related features purely mathematically, without regard to algorithmic implementation. We then remark how a fast instantiation of those features is enabled by the new tools given in this work; the details may be found in Section 6.

3.1 Homomorphic Encryption Scheme

A Ring-LWE encryption scheme is defined over the ring of integers $R = \mathcal{O}_K$ of a number field $K/\mathbb{Q}$, and is parameterized by a plaintext modulus p and a ciphertext modulus $q \gg p$, which are positive integers.[13] Recall that $\mathfrak{C} := R^\vee \supseteq R$ is the fractional dual (or "codifferent") ideal of R, relative to $\mathbb{Q}$. For any positive integer r (and specifically, $r = p$ and $r = q$), define the quotient ring $R_r := R/rR$, the quotient R-module (and R_r-module) $\mathfrak{C}_r := \mathfrak{C}/r\mathfrak{C}$, and more generally, $\mathfrak{C}_r^i := \mathfrak{C}^i/r\mathfrak{C}^i$ for any power $i \geq 0$.

3.1.1 Plaintext Encoding, Ciphertexts, and Decryption

The plaintext ring is R_p, a secret key is an element $s \in \mathfrak{C}$, and a ciphertext is a pair

$$c = (c_0, c_1) \in \mathfrak{C}_q \times R_q \quad \text{for which} \quad c(s) := c_0 + c_1 \cdot s \in \mathfrak{C}_q$$

is a "noisy encoding," modulo $q\mathfrak{C}$, of the plaintext.[14] Essentially, the ciphertext may be seen as an affine linear polynomial $c(S)$ in a variable S that represents the secret key, though its coefficients come from different modules.

For concreteness, in this template we use a "least-significant digit" noisy plaintext encoding, à la [BV11a, BV11b, BGV12], for which p and q must be coprime for security.[15] An important part of this encoding is an R-module isomorphism $\theta\colon \mathfrak{C}_p \to R_p$ for which $\theta^{-1}(\mu) = \mu \bmod p\mathfrak{C}$ (using the natural inclusion $R \subseteq \mathfrak{C}$),

[13] More generally, the plaintext "modulus" could be any ideal $\mathfrak{p}$ of R having a known short $\mathbb{Z}$-basis. This approach was used in [CLPX18, GV25] with *cyclotomic* rings to get smaller noise growth (under homomorphic operations) for characteristic-p plaintext rings, for large p of very special form. This idea works equally well in our setting of *general* or *abelian* number fields, and may even enlarge the class of characteristics p for which this technique can usefully apply. At minimum, our techniques focused on decomposition subrings allow for avoiding wasted "SIMD capacity" in this approach.

[14] The use of the dual $\mathfrak{C} = R^\vee$ here is important for both security and error tolerance: the known hardness results for Ring-LWE are obtained most directly and tightly for the form of the problem involving $\mathfrak{C}$ and spherically bounded error (see [LPR10, Section 3.3]), and having a short basis of R enables efficient decoding of $\mathfrak{C}$ under such error; see Section 6.2 for details.

[15] An analogous "most-significant digit" encoding can be given for "scale-invariant" schemes à la [Bra12, FV12], along with an approximate encoding à la [CKKS17].

and $\theta(z) = t \cdot z \bmod pR$ for some $t \in \mathfrak{D} = \mathfrak{C}^{-1}$ satisfying $t = 1 \pmod{pR}$. Such t exists (and can be computed efficiently) by the Chinese Remainder Theorem as long as $\mathfrak{D}$ and pR are coprime, or equivalently, if no prime divisor of p ramifies in K.[16] A noisy encoding of $\mu \in R_p$ is an error term $e \in \theta^{-1}(\mu) = \mu + p\mathfrak{C}$ that is "decodably small" relative to $q\mathfrak{C}$—i.e., from $e \bmod q\mathfrak{C}$ (an element of $\mathfrak{C}_q$) we can efficiently recover $e \in \mathfrak{C}$. Accordingly, decryption computes $c(s) \in \mathfrak{C}_q$, decodes it to $e \in \mathfrak{C}$, and outputs $\mu = \theta(e \bmod p\mathfrak{C}) \in R_p$.

For an instantiation, fast multiplication in R_q, and more generally across the modules $\mathfrak{C}_q^i$, is enabled by the use of a *CRT-basis* representation and *fast CRT transforms*, as given in Section 5 (see Remark 5.6 and Section 6.1 for details). And we can efficiently sample error terms, and get suitable noisy-encoding and decoding functions, using *short, structured* bases of $\mathfrak{C}$ and R (see Section 6.2 for details).

3.1.2 Homomorphic Operations

The scheme supports various homomorphic operations on encrypted plaintexts. We recall the main ones: addition, multiplication by a public value in R_p, multiplication of two encrypted values, and (as shown in [GHS12b]) applying an automorphism of K. In all of the following, let ciphertexts $c = (c_0, c_1), c' = (c'_0, c'_1)$ respectively encrypt plaintexts $\mu, \mu' \in R_p$ via noisy encodings $e, e' \in \mathfrak{C}$.

Linear operations. To homomorphically add the plaintexts encrypted by c, c', we simply compute $c_+ = (c_0 + c'_0, c_1 + c'_1)$. Observe that

$$c_+(s) = (c_0 + c'_0) + (c_1 + c'_1) \cdot s = c(s) + c'(s) = e + e' \pmod{q\mathfrak{C}} ,$$

so c_+ decrypts to $\mu + \mu' = \theta(e + e' \bmod p\mathfrak{C})$, as long as the combined error $e + e'$ is small enough.

To homomorphically multiply the plaintext encrypted by c by a public value $v \in R_p$, we simply output the ciphertext $\tilde{c} = \tilde{v} \cdot c = (\tilde{c}_0 = \tilde{v} \cdot c_0, \tilde{c}_1 = \tilde{v} \cdot c_1)$, where $\tilde{v} \in R$ is a "small" representative of v. Observe that

$$\tilde{c}(s) = \tilde{v} \cdot c(s) = \tilde{v}e \pmod{q\mathfrak{C}} ,$$

so $\tilde{c}$ decrypts to $v \cdot \mu = \theta(\tilde{v}e \bmod p\mathfrak{C})$, as long as the enlarged error $\tilde{v}e$ remains small enough.

Multiplication. To homomorphically multiply the two encrypted plaintexts, we first multiply their ciphertexts as formal polynomials. That is, we compute $c_\times(S) = (c_0 + c_1 S)(c'_0 + c'_1 S)$, which we represent as its *triple* of coefficients

$$c_\times = (c_0 c'_0 \,,\, c_0 c'_1 + c_1 c'_0 \,,\, c_1 c'_1) \in \mathfrak{C}_q^2 \times \mathfrak{C}_q \times R_q \,.$$

Observe that $c_\times(s) = c(s) \cdot c'(s) = e \cdot e' \pmod{\mathfrak{q}\mathfrak{C}^2}$ is congruent to the product of the two noisy encodings.[17]

However, note that the resulting ciphertext and noise product involve $\mathfrak{C}^2 \supseteq \mathfrak{C}$, which no longer matches the initial setup. To address this, we multiply the ciphertext by a known *small* value $d \in \mathfrak{D} := \mathfrak{C}^{-1} \subseteq R$ for which pR is coprime to $d \cdot \mathfrak{D}^{-1} \subseteq R$, and hence to dR as well (because pR and $\mathfrak{D}$ are coprime by assumption). Specifically, we let

$$\tilde{c}_\times = d \cdot c_\times \in \mathfrak{C}_q \times R_q \times \mathfrak{D}_q \,,$$

so that $\tilde{c}_\times(s) = d \cdot e \cdot e' \pmod{q\mathfrak{C}}$. Thus, $\tilde{c}$ decrypts to $d\mu\mu' = \theta(dee' \bmod p\mathfrak{C}) \in R_p$, as long as $dee' \in \mathfrak{C}$ is small enough relative to $q\mathfrak{C}$. The extra factor(s) of d can be tracked and removed upon decryption.

[16] This coprimality condition can be avoided by generalizing to $\theta^{-1}(\mu) = u \cdot \mu \bmod p\mathfrak{C}$ for some $u \in \mathfrak{C}$ such that the ideals $u\mathfrak{D} \subseteq R$ and pR are coprime, and using $t \in \mathfrak{D}$ such that $tu = 1 \pmod{pR}$; see [LPR10, Section 2.3.9] for further details.

[17] For a most-significant-digit encoding as used in [Bra12, FV12], a slightly different kind of "scale-invariant" ciphertext multiplication is used, but the outcome is substantially the same.

Effectively, d is a kind of "expansion factor" associated with maintaining the invariant that products of noisy encodings remain in $\mathfrak{C}$, and remain spherically bounded in their distributions. In any abelian number field of conductor m, just as in the cyclotomic case we can use $d = \hat{m} \in \mathfrak{D}$, where $\hat{m} = m/2$ if m is even and $\hat{m} = m$ otherwise; moreover, some fields have even smaller choices.

Finally, we convert $\tilde{c}_\times$ back to a *linear* polynomial in S by applying a "key-switching" operation to the quadratic coefficient $c_1 c_1' \in R_q$ of $c_\times$. This uses a "gadget decomposition" to express the coefficient in terms of short elements of R, and a "key-switching hint" consisting of a suitable encryption of $ds^2 \in \mathfrak{C}$ under s.[18] This just additively increases the error size by some fixed amount, and results in a ciphertext whose form matches the initial setup.

Automorphisms. To homomorphically apply an automorphism τ of $K/\mathbb{Q}$, we first compute the ciphertext $c_\tau = (\tau(c_0), \tau(c_1))$, which is in $\mathfrak{C}_q \times R_q$ because $\tau(q) = q, \tau(R) = R, \tau(\mathfrak{C}) = \mathfrak{C}$. Observe that

$$c_\tau(\tau(s)) = \tau(c_0) + \tau(c_1) \cdot \tau(s) = \tau(c_0 + c_1 \cdot s) = \tau(e) \pmod{q\mathfrak{C}} .$$

Moreover, since automorphisms preserve the size of the noise (in the canonical embedding) and $\tau(t) = 1 \pmod{pR}$, we see that c_τ is an encryption of $\theta(\tau(e) \bmod p\mathfrak{C}) = t \cdot \tau(e) \bmod p\mathfrak{C} = \tau(t \cdot e \bmod p\mathfrak{C}) = \tau(\mu)$, but under the "conjugate" secret key $\tau(s) \in \mathfrak{C}$. To make c_τ a proper encryption under s, we apply key-switching to the coefficient $\tau(c_1)$, using a suitable encryption of $\tau(s)$ under s.

Instantiation. For an instantiation, both lifting $v \in R_p$ to a small representative $\tilde{v} \in R$, and key-switching with little additional noise, are enabled by having a short basis of R (the details are standard, and exactly as in [LPR13]). Fast multiplication by $\tilde{v}$ and d, and application of automorphisms on R_q and $\mathfrak{C}_q$, are also enabled by the use of a CRT-basis representation and fast CRT transforms, as given in Section 5 (see Remark 5.6 and Section 6.1 for details).

3.2 Packed Bootstrapping Framework

Here we recall the relevant details of the efficient "packed" bootstrapping template of [GHS12a], which was further refined in [AP13]. The main idea is to efficiently express, via the FHE scheme's native operations, the decryption of a fixed ciphertext c (which is to be bootstrapped) as a function of the secret key s. The bootstrapping algorithm homomorphically evaluates this function on an *encryption* of s, yielding an encryption of the decryption of c, i.e., an encryption of the same underlying plaintext. For appropriate parameters, the resulting ciphertext will have significantly smaller noise than c has, allowing further homomorphic operations.

Recall from above that decryption of a ciphertext c works primarily by "decoding" $c(s) \in R_q^\vee$ to $R_p^\vee$. (In the bootstrapping context we ignore the final module isomorphism θ that maps back to R_p, because we want to continue operating homomorphically on the plaintext.) As detailed in Section 6.2, this is implemented coordinate-wise relative to a certain $\mathbb{Z}$-basis $\vec{d}$ of $R^\vee$, which is thus also a $\mathbb{Z}_r$-basis of $R_r^\vee$ for $r = p, q$. More specifically, writing $c(s) = \langle \vec{d}, \mathbf{z} \rangle \in R_q^\vee$ for some coefficient vector $\mathbf{z}$ over $\mathbb{Z}_q$, decryption computes $\langle \vec{d}, \mathsf{Decode}(\mathbf{z}) \rangle \in R_p^\vee$, where $\mathsf{Decode}\colon \mathbb{Z}_q \to \mathbb{Z}_p$ is a suitable non-linear decoding function, applied entry-wise to $\mathbf{z}$. For example, for the least-significant-digit noisy encoding used in our template, Decode lifts its argument to the smallest $\mathbb{Z}$-representative, then reduces modulo p.[19]

[18] This assumes the security of a key-dependent encryption; alternatively, one can instead key-switch using encryptions of ds^2 and s under an independent key s'.

[19] For the most-significant-digit encoding, Decode simply "scales down" and rounds, as $\mathsf{Decode}(z) = \lfloor \frac{p}{q} \cdot z \rceil \in \mathbb{Z}_p$.

The bootstrapping template of [GHS12a, AP13] homomorphically does the decoding in parallel across all the coordinates at once, using the scheme's intrinsic SIMD operations. To do this, it expresses the decryption function as three phases, and evaluates them homomorphically on the encrypted secret key:

1. The first phase moves the entries of $\mathbf{z}$ into the SIMD "slots." More precisely, this phase computes $c(s) = \langle \vec{d}, \mathbf{z} \rangle$ and maps it to $\langle \vec{c}, \mathbf{z} \rangle$, where $\vec{c}$ is the standard basis of the SIMD slots, i.e., each entry of $\vec{c}$ is 1 in a distinct slot and 0 in all the other slots. In other words, this map is the $\mathbb{Z}_q$-linear function that sends each entry of $\vec{d}$ to the corresponding entry of $\vec{c}$.

2. The second phase applies Decode in parallel across all the slots (i.e., the coordinates of $\mathbf{z}$), yielding $\langle \vec{c}, \mathsf{Decode}(\mathbf{z}) \rangle$. This can be expressed algebraically using additions and multiplications (see, e.g., [GHS12a, AP13] and several subsequent works), and is outside the scope of this paper.

3. The third phase moves the entries of the slots back to the original basis $\vec{d}$, essentially inverting the linear function from the first phase. In other words, it evaluates the linear function that sends each entry of $\vec{c}$ to the corresponding entry of $\vec{d}$, resulting in $\langle \vec{d}, \mathsf{Decode}(\mathbf{z}) \rangle \in R_p^\vee$.

For an instantiation, the basis $\vec{c}$ is in fact the CRT basis of $R_q^\vee$ (see Section 5.1). Using the matching Kronecker-product structures of both $\vec{c}$ and $\vec{d}$, and other advantageous properties of $\vec{c}$, we can efficiently *homomorphically* evaluate the CRT transforms from the first and third phases via homomorphic automorphisms, following the framework in Section 4 (see Section 6.3 for details).

4 Sparse (Automorphism) Decompositions

In this section we lay out a general framework for expressing linear functions on an arbitrary Galois extension in terms of its automorphisms. Our ultimate goal is to obtain *sparse decompositions* for "structured" functions of interest, like Chinese Remainder Transforms (CRTs). Such a decomposition expresses a function as a linear combination of *relatively few* automorphisms, or more generally, as the (sequential) *composition* of a small number of such linear combinations. This allows us to efficiently evaluate the function *homomorphically*, since applying an automorphism is an efficient "native" operation in homomorphic encryption.[20]

To get a sparse decomposition for a structured function, we view it as mapping from one structured (Kronecker-product) vector of elements to another, and map each factor to its counterpart in sequence. The primary challenge is to ensure that each factor can be mapped sparsely, without affecting the other factors.

4.1 Arbitrary Linear Functions

It is well known (and straightforward to prove) that the automorphisms of any field are linearly independent over the field. So, for any finite Galois extension L/K, because the automorphisms $\tau \in \mathrm{Gal}(L/K)$ are K-linear and there are $\deg(L/K)$ of them, they form a *basis* for the space of K-linear functions from L to itself. In other words, any such function can be expressed as an L-linear combination of automorphisms. Lemma 4.1 makes this explicit, by giving the L-coefficients for a linear function that maps particular inputs to desired outputs, using duality (see Section 2.1.2). It also adapts to linear functions on certain *quotients* of fractional ideals in number fields.

[20] By contrast, native homomorphic operations cannot directly manipulate coefficient vectors in most bases of interest, so standard sparse decompositions and algorithms that operate on coefficient vectors (e.g., the Number-Theoretic Transform) do not translate well to homomorphic evaluation.

Lemma 4.1. *Let L/K be any finite Galois extension with $G = \mathrm{Gal}(L/K)$, and $\vec{b}, \vec{c}$ be over L with the same index set, such that $\vec{b}$ is K-linearly independent with dual $\vec{b}^\vee$. Then $f\colon L \to L$ defined as*

$$f(x) := \sum_{\tau \in G} \langle \vec{c}, \tau(\vec{b}^\vee) \rangle \cdot \tau(x) \tag{4.1}$$

is a K-linear function for which $f(\vec{b}) = \vec{c}$.

Alternatively, let L/K be a Galois extension of number fields, and let

- *$\mathfrak{r}$ be an ideal of $\mathcal{O}_K$ that is coprime with the different ideal $\mathfrak{D} = \mathfrak{D}_{L/K} \subseteq \mathcal{O}_L$, with $d \in \mathfrak{D}$ satisfying $d = 1 \pmod{\mathfrak{r}}$;*
- *$\mathfrak{b}$ be a fractional ideal of $\mathcal{O}_L$ that is fixed by every $\tau \in G$ (e.g., $\mathcal{O}_L$ or $\mathcal{O}_L^\vee$);*
- *$\vec{b}$ and $\vec{c}$ be over $\mathfrak{b}/\mathfrak{r}\mathfrak{b}$, and $\vec{b}^\vee$ be over $\mathfrak{b}^\vee/\mathfrak{r}\mathfrak{b}^\vee$ and dual to $\vec{b}$.*

Then $g\colon \mathfrak{b}/\mathfrak{r}\mathfrak{b} \to \mathfrak{b}/\mathfrak{r}\mathfrak{b}$ defined as $g(x) = d \cdot f(x)$ is a $(\mathcal{O}_K/\mathfrak{r})$-linear function for which $f(\vec{b}) = \vec{c}$.

Note that in the second claim, the coefficients $d \cdot \langle \vec{c}, \tau(\vec{b}^\vee) \rangle$ of the automorphisms are in $\mathcal{O}_L/\mathfrak{r}\mathcal{O}_L$, because $\mathfrak{D} \cdot \mathfrak{b} \cdot \mathfrak{b}^\vee = \mathcal{O}_L$ (see Section 2.2.4).

Proof. First, the definition in Equation (4.1) is K-linear because every $\tau \in G$ is, so it suffices to show that $f(\vec{b}^t) = \vec{c}^t$. Indeed, by multplicativity of automorphisms and the definitions of trace and duality,

$$f(\vec{b}^t) = \sum_{\tau \in G} (\vec{c}^t \cdot \tau(\vec{b}^\vee)) \cdot \tau(\vec{b}^t) = \vec{c}^t \cdot \sum_{\tau \in G} \tau(\vec{b}^\vee \cdot \vec{b}^t) = \vec{c}^t \cdot \mathrm{Tr}_{L/K}(\vec{b}^\vee \cdot \vec{b}^t) = \vec{c}^t \ . \tag{4.2}$$

For the second claim, Equation (4.2) holds over $\mathfrak{b} \cdot \mathfrak{b}^{-1} \cdot \mathfrak{b} = \mathfrak{D}^{-1}\mathfrak{b}$ modulo $\mathfrak{D}^{-1}\mathfrak{r}\mathfrak{b}$. So,

$$g(\vec{b}) = d \cdot f(\vec{b}) = d \cdot \vec{c} = \vec{c} \pmod{\mathfrak{r}\mathfrak{b}}$$

because $d \in \mathfrak{D}$, and because $d - 1 \in \mathfrak{r}$ and $\vec{c}$ is over $\vec{b}$. □

4.2 Sparse Decompositions of Structured Linear Functions

Now let $M/L/K$ be a tower of finite extensions, where M/K is abelian and hence so are M/L and L/K. Let $G_{M/L} := \mathrm{Gal}(M/L)$, which is a subgroup of $G_{M/K} := \mathrm{Gal}(M/K)$, and $G_{L/K} := \mathrm{Gal}(L/K)$.[21] In our setting of interest, $\deg(L/K)$ is typically small, but $\deg(M/L)$ is potentially large.

Let $\vec{b}_{M/L}$ and $\vec{c}_{M/L}$ (respectively, $\vec{b}_{L/K}$ and $\vec{c}_{L/K}$) be vectors over M (resp., L) having the same index set, where $\vec{b}_{M/L}$ (resp., $\vec{b}_{L/K}$) is linearly independent over L (resp., K), and let $\vec{b}_{M/K} = \vec{b}_{M/L} \otimes \vec{b}_{L/K}$ and $\vec{c}_{M/K} = \vec{c}_{M/L} \otimes \vec{c}_{L/K}$. Alternatively, when M, L, K are number fields, we can let these vectors be over suitable quotients as in the second part of Lemma 4.1, with no change to any of the following treatment; this is actually the setup we will use in applications.

Our goal is to use an M-linear (alternatively, $(\mathcal{O}_M/\mathfrak{r}\mathcal{O}_M)$-linear) combination of the automorphisms in $G_{M/K}$ to express a K-linear (alternatively, $(\mathcal{O}_K/\mathfrak{r})$-linear) function f that satisfies

$$f(\vec{b}_{M/K}) = \vec{c}_{M/K} \ .$$

As shown in Lemma 4.1, this can be done generically using up to $|G_{M/K}| = \dim(M/K)$ such automorphisms. However, we seek a *sparse* decomposition, i.e., one that uses relatively *few* automorphisms. Here we show that it is often possible to do so by exploiting the Kronecker-product structure of $\vec{b}_{M/K}$ and $\vec{c}_{M/K}$, and particular properties of the component factors of $\vec{c}_{M/K}$.

[21] More generally, for all the material in this section it suffices for M/K to be Galois, and for $G_{M/L}$ to be a *normal* subgroup of $G_{M/K}$, in which case L/K is also Galois.

Decomposing f. First, we express our desired map as the *composition* of two linear maps, as follows:

$$\vec{b}_{M/L} \otimes \vec{b}_{L/K} \xrightarrow{f_L} \vec{c}_{M/L} \otimes \vec{b}_{L/K} \xrightarrow{f_K} \vec{c}_{M/L} \otimes \vec{c}_{L/K} \ . \tag{4.3}$$

More precisely,

- f_L is any L-linear function for which $f_L(\vec{b}_{M/L}) = \vec{c}_{M/L}$, which by L-linearity implies that

$$f_L(\vec{b}_{M/L} \otimes \vec{b}_{L/K}) = f_L(\vec{b}_{M/L}) \otimes \vec{b}_{L/K} = \vec{c}_{M/L} \otimes \vec{b}_{L/K} \ ;$$

- f_K is any K-linear function for which $f_K(\vec{c}_{M/L} \otimes \vec{b}_{L/K}) = \vec{c}_{M/L} \otimes \vec{c}_{L/K}$.

Then it is clear that $f = f_K \circ f_L$ is K-linear and satisfies $f(\vec{b}_{M/K}) = \vec{c}_{M/K}$, as needed.

Expressing f_L. As shown in Lemma 4.1, any L-linear f_L for which $f_L(\vec{b}_{M/L}) = \vec{c}_{M/L}$ can be obtained as an M-linear combination of up to $\deg(M/L)$ automorphisms in $G_{M/L}$. But if $\deg(M/L)$ is large, as is often the case in our setting, this may not be as sparse as we would like.

Instead, we can proceed inductively, as long as $\vec{b}_{M/L}$ and $\vec{c}_{M/L}$ themselves factor as Kronecker products of appropriate vectors along a tower $M/L'/L$, where we typically want $\deg(L'/L)$ to be small. Then we can express f_L using a decomposition analogous to the one in Equation (4.3).

After inductively "unfolding" all the decompositions, this approach requires a tower $K_t/K_{t-1}/\cdots/K_1/K_0$ of suitable extensions, each typically of small relative degree, and decomposes the function f as

$$\vec{b}_t \otimes \vec{b}_{t-1} \otimes \cdots \otimes \vec{b}_1 \xrightarrow{f_t} \vec{c}_t \otimes \vec{b}_{t-1} \otimes \cdots \otimes \vec{b}_1 \xrightarrow{f_{t-1}}$$
$$\vec{c}_t \otimes \vec{c}_{t-1} \otimes \cdots \otimes \vec{b}_1 \xrightarrow{f_{t-2}} \cdots \xrightarrow{f_1} \vec{c}_t \otimes \vec{c}_{t-1} \otimes \cdots \otimes \vec{c}_1 \ ,$$

for appropriate vectors $\vec{b}_i, \vec{c}_i$. The following is then immediate.

Lemma 4.2. *The total number of automorphisms used in the expression of $f = f_1 \circ \cdots \circ f_t$ is the sum of the number of automorphisms used in the expression of each f_i.*

Expressing f_K. Handling f_K is more subtle. While a K-linear function *on* L that maps $\vec{b}_{L/K}$ to $\vec{c}_{L/K}$ can be obtained generically from the (typically few) automorphisms of L/K, our goal and set of available automorphisms are different: we wish to obtain a K-linear function f_K *on* M that maps $\vec{c}_{M/L} \otimes \vec{b}_{L/K}$ to $\vec{c}_{M/L} \otimes \vec{c}_{L/K}$ using few of the (typically many) automorphisms *of* M/K. We will achieve this goal *for specific target vectors $\vec{c}_{M/L}$ of interest*, by showing that relatively few of the automorphisms are needed, thanks to particular properties of $\vec{c}_{M/L}$.

Lemma 4.3. *In the expression of f_K from Lemma 4.1, the coefficient of $\tau \in G_{M/K}$ is*

$$\langle \vec{c}_{M/L} \, , \, \tau(\vec{c}^{\vee}_{M/L}) \rangle \cdot \langle \vec{c}_{L/K} \, , \, \tau|_L(\vec{b}^{\vee}_{L/K}) \rangle \ . \tag{4.4}$$

In particular, the coefficient of τ is zero if $\langle \vec{c}_{M/L}, \tau(\vec{c}^{\vee}_{M/L}) \rangle = 0$.

Proof. By the definition of f_K and Lemma 4.1, the coefficient of τ is

$$\langle \vec{c}_{M/L} \otimes \vec{c}_{L/K} \, , \, \tau(\vec{c}^{\vee}_{M/L} \otimes \vec{b}^{\vee}_{L/K}) \rangle \ .$$

The claim then follows by the multiplicativity of automorphisms, the mixed-product property, and the fact that $\vec{b}^{\vee}_{L/K}$ is over L. □

Note that the factors $\langle \vec{c}_{L/K}, \tau|_L(\vec{b}^\vee_{L/K})\rangle$ in Equation (4.4) are exactly the coefficients in the above-mentioned K-linear function on L that maps $\vec{b}_{L/K}$ to $\vec{c}_{L/K}$, because the restriction of $G_{M/K}$ to L is $G_{L/K}$. There are $\deg(L/K)$ such factors (each repeated $\deg(M/L)$ times as τ ranges over $G_{M/K}$), which may be arbitrary for general f_K.

In Section 5.3 and Lemma 5.13 we show that for the $\vec{c}_{M/L}$ component in the Kronecker-product factorization of a *CRT basis*, relatively few of the factors $\langle \vec{c}_{M/L}, \tau(\vec{c}^\vee_{M/L})\rangle$ are nonzero, hence the expression of f_K from Lemma 4.1 is indeed sparse. In summary, using the approach from this subsection we can map *to* a CRT basis *from* any similarly structured basis, using relatively few automorphisms overall.

4.3 Working "Bottom Up"

The decomposition used above in Section 4.2 works "top down," first replacing $\vec{b}_{M/L}$ with $\vec{c}_{M/L}$ via some f_L (which may involve replacing some smaller "top" components, inductively), then replacing $\vec{b}_{L/K}$ with $\vec{c}_{L/K}$ via some f_K. By Lemma 4.3, this yields a sparse decomposition if the "top" component $\vec{c}_{M/L}$ of the *target* vector has suitable properties. Alternatively, we can work "bottom up," which is advantageous if the top component of the map's *source* vector yields a sparse decomposition.

As a primary example, consider the *inverse* function f^{-1} that maps $\vec{c}$ to $\vec{b}$ (which exists assuming $\vec{c}$ is linearly independent over K, or alternatively, $\mathcal{O}_K/\mathfrak{r}$). We express $f^{-1} = f_L^{-1} \circ f_K^{-1}$ as follows:

$$\vec{c}_{M/L} \otimes \vec{c}_{L/K} \xrightarrow{f_K^{-1}} \vec{c}_{M/L} \otimes \vec{b}_{L/K} \xrightarrow{f_L^{-1}} \vec{b}_{M/L} \otimes \vec{b}_{L/K} \ .$$

As with f_L above, the function f_L^{-1} can be decomposed in an analogous bottom-up way. And analogously to Lemma 4.3, in the expression of f_K^{-1}, the coefficient of $\tau \in G_{M/K}$ has the same multiplicand $\langle \vec{c}_{M/L}, \tau(\vec{c}^\vee_{M/L})\rangle$ as it does in the expression for f_K (but the other multiplicand is typically different). So, if relatively few of these multiplicands are nonzero, we get a sparse decomposition for both f and f^{-1}.

5 Chinese Remainder Theorem Bases and Transforms

In Section 5.1 we define the *Chinese Remainder Theorem (CRT) basis* of an arbitrary *abelian* (Galois) extension of number rings modulo a suitable ideal, and show how it yields fast multiplication in the quotient ring (and related quotient modules).[22] In Section 5.2 we show that CRT bases admit a natural Kronecker-product factorization into smaller CRT bases, going down any tower of intermediate number rings. We then exploit this structure to give two kinds of "sparse decompositions" of *CRT transforms*, and associated fast algorithms, that map between CRT bases and any similarly structured bases:

- "in the clear" transforms and algorithms (Section 5.4) that work directly on coefficient vectors, and
- ones expressed in terms of relatively few automorphisms (Section 5.3) via the framework of Section 4, which yield efficient *homomorphic* evaluations of CRT transforms suitable for packed bootstrapping.

5.1 CRT Basis

Let L/K be a finite abelian (Galois) extension of number fields with $G_{L/K} := \mathrm{Gal}(L/K)$. Let $\mathfrak{r}$ be a prime ideal in $\mathcal{O}_K$, and assume without loss of generality that the common decomposition group

[22] The abelian assumption is mostly for convenience; we mainly use it to ensure that all conjugate prime ideals have the same decomposition group. Alternatively, we can use the weaker assumption that all the relevant decomposition groups are normal subgroups.

$D_{L/K} := D_{L/K}(\mathfrak{r}_\ell)$ of the prime—and hence pairwise coprime—ideals $\mathfrak{r}_\ell$ lying over $\mathfrak{r}$ in L is *trivial*.[23] Then recall from Section 2.2.3 that the $\mathfrak{r}_\ell$ are indexed by $\ell \in G_{L/K}$, which acts regularly (i.e., freely and transitively) on them, and that $\mathfrak{r}$ splits completely in $\mathcal{O}_L$, as $\mathfrak{r}\mathcal{O}_L = \prod_{\ell \in G_{L/K}} \mathfrak{r}_\ell$.

For our purposes it is convenient to generalize the above setup to possibly *non-prime* (proper) ideals, as follows.

Definition 5.1 (Generalized complete splitting). We say that a proper ideal $\mathfrak{r}$ of $\mathcal{O}_K$ *splits completely* in $\mathcal{O}_L$ if $\mathfrak{r}\mathcal{O}_L = \prod_{\ell \in G_{L/K}} \mathfrak{r}_\ell$ for some *pairwise coprime* ideals $\mathfrak{r}_\ell$ of $\mathcal{O}_L$ that are *conjugates*, i.e., $G_{L/K}$ acts transitively upon them. Without loss of generality, we index them so that $\ell'(\mathfrak{r}_\ell) = \mathfrak{r}_{\ell' \circ \ell}$ for all $\ell, \ell' \in G_{L/K}$.

Note that $G_{L/K}$ also acts *freely*, and hence *regularly*, on the $\mathfrak{r}_\ell$ by their pairwise coprimality, and hence distinctness (this is where we use the fact that $\mathfrak{r}$ is proper, hence so are the $\mathfrak{r}_\ell$). In addition, the indexing of *all* the $\mathfrak{r}_\ell$ is determined by the index of any *one* of them, by transitivity.

Remark 5.2. A necessary and sufficient condition for complete splitting is that each prime ideal factor of $\mathfrak{r}$ (in $\mathcal{O}_K$) splits completely in $\mathcal{O}_L$. In brief, sufficiency is simply by multiplying corresponding factors of the complete splittings, and necessity is because a nontrivial decomposition group for some prime ideal factor of $\mathfrak{r}$ implies a failure of pairwise coprimality. A consequence of these observations is that each $\mathfrak{r}_\ell$ is prime if and only if $\mathfrak{r}$ is prime.

Suppose that proper ideal $\mathfrak{r}$ of $\mathcal{O}_K$ splits completely in $\mathcal{O}_L$, with factorization as in Definition 5.1. By the Chinese Remainder Theorem (see Section 2.2.2), the natural homomorphism induces a ring isomorphism

$$\mathcal{O}_L/\mathfrak{r}\mathcal{O}_L \cong \prod_{\ell \in G_{L/K}} (\mathcal{O}_L/\mathfrak{r}_\ell) \,.$$

That is, each element $x \in \mathcal{O}_L/\mathfrak{r}\mathcal{O}_L$ can be uniquely represented as a tuple, indexed by $G_{L/K}$, whose ℓth entry is $x \bmod \mathfrak{r}_\ell$, i.e., the coset $x + \mathfrak{r}_\ell \in \mathcal{O}_L/\mathfrak{r}_\ell$. Addition and multiplication in $\mathcal{O}_L/\mathfrak{r}\mathcal{O}_L$ then correspond to component-wise addition and multiplication (respectively) of these tuples.

Recall from Section 2.2.2 that any instance of a CRT isomorphism (i.e., any collection of pairwise coprime ideals) yields a natural CRT vector of elements modulo the product of those ideals. In the case of complete splitting we call this a *CRT basis*, as justified by Lemma 5.5 below.

Definition 5.3 (CRT Basis). Let $\mathfrak{r}$ be a proper ideal of $\mathcal{O}_K$ that splits completely in $\mathcal{O}_L$, as $\mathfrak{r}\mathcal{O}_L = \prod_{\ell \in G_{L/K}} \mathfrak{r}_\ell$. Then the associated *mod-$\mathfrak{r}$ CRT basis* of $\mathcal{O}_L/\mathcal{O}_K$ is simply the CRT vector $\vec{c}$ over $\mathcal{O}_L/\mathfrak{r}\mathcal{O}_L$, indexed by $G_{L/K}$, for these $\mathfrak{r}_\ell$. Namely, $c_\ell = \delta_{\ell,\ell'} \pmod{\mathfrak{r}_{\ell'}}$ for all $\ell, \ell' \in G_{L/K}$, or more compactly, $\vec{c} = \vec{e}_\ell \pmod{\mathfrak{r}_\ell}$ (where $\vec{e}_\ell$ is the indicator vector whose ℓth entry is 1 and the rest are 0).

Notice that $\vec{c}$ is uniquely defined up to the indexing of the ideals $\mathfrak{r}_\ell$, which is determined by the indexing of any one of them. However, the choice of the ideals $\mathfrak{r}_\ell$ themselves is not unique unless $\mathfrak{r}$ is prime, so the CRT basis is associated with a particular splitting, which will always be clear from context. Also, since $G_{L/K}$ acts regularly on the $\mathfrak{r}_\ell$ via $\tau(\mathfrak{r}_\ell) = \mathfrak{r}_{\tau \circ \ell}$, the same goes for the c_ℓ.

Remark 5.4 (Self-duality of the CRT basis). Observe that $c_\ell^2 = c_\ell$ and $c_\ell \cdot c_{\ell'} = 0$ for $\ell \neq \ell'$ (where recall that both equalities are modulo $\mathfrak{r}\mathcal{O}_L$). So, the CRT basis $\vec{c}$ is essentially *self-dual* modulo $\mathfrak{r}$ (see Section 2.2.4).

[23] If $D_{L/K}$ is non-trivial, then L can be replaced by the decomposition subfield $\tilde{L} = L^{D_{L/K}}$ of the $\mathfrak{r}_\ell$. Then $\mathfrak{r}$ splits completely in $\mathcal{O}_{\tilde{L}}$, or equivalently, the decomposition group of the primes lying over $\mathfrak{r}$ in $\mathcal{O}_{\tilde{L}}$ is trivial.

More specifically, $\vec{c}^\vee = \vec{c} \bmod \mathfrak{r}\mathcal{O}_L^\vee$, which can be seen as a vector over $\mathcal{O}_L^\vee/\mathfrak{r}\mathcal{O}_L^\vee$ by the inclusion $\mathcal{O}_L \subseteq \mathcal{O}_L^\vee$. This is because $\mathrm{Tr}_{L/K}(c_\ell \cdot c_{\ell'}^\vee)$ for $\ell \neq \ell'$ is

$$\mathrm{Tr}_{L/K}(\mathfrak{r}\mathcal{O}_L^\vee) = \mathfrak{r}\,\mathrm{Tr}_{L/K}(\mathcal{O}_L^\vee) = \mathfrak{r}\mathcal{O}_K = 0 \pmod{\mathfrak{r}} ,$$

and for $\ell = \ell'$ is

$$\mathrm{Tr}_{L/K}(c_\ell + \mathfrak{r}\mathcal{O}_L^\vee) = \sum_{j \in G_{L/K}} c_j + \mathfrak{r}\,\mathrm{Tr}_{L/K}(\mathcal{O}_L^\vee) = 1 \pmod{\mathfrak{r}}$$

by the definition and K-linearity of $\mathrm{Tr}_{L/K}$, the Chinese Remainder Theorem, and because $G_{L/K}$ acts regularly on the c_ℓ.

The following justifies the name "CRT basis."

Lemma 5.5. *The mod-$\mathfrak{r}$ CRT basis $\vec{c}$ is an $(\mathcal{O}_K/\mathfrak{r})$-basis of $\mathcal{O}_L/\mathfrak{r}\mathcal{O}_L$. That is, $x \in \mathcal{O}_L/\mathfrak{r}\mathcal{O}_L$ if and only if it can be expressed uniquely as $x = \langle \vec{c}, \vec{x} \rangle = \vec{c}^t \cdot \vec{x}$ for some vector $\vec{x}$ indexed by $G_{L/K}$ over $\mathcal{O}_K/\mathfrak{r}$.*

Proof. The claim follows mainly from the (fairly standard) fact that for all $\ell \in G_{L/K}$, the natural ring homomorphism from $\mathcal{O}_K/\mathfrak{r}$ to $\mathcal{O}_L/\mathfrak{r}_\ell$ is an isomorphism, where the $\mathfrak{r}_\ell$ are as defined in Definition 5.1; recall that $G_{L/K}$ acts regularly on them. First, the homomorphism is injective: because any $z \in \mathcal{O}_K/\mathfrak{r}$ is fixed by $G_{L/K}$, its CRT representation must satisfy $z + \mathfrak{r}_{\ell'} = (\ell' \circ \ell^{-1})(z + \mathfrak{r}_\ell)$ for any $\ell, \ell' \in G_{L/K}$. So, if $z, z' \in \mathcal{O}_K/\mathfrak{r}$ are congruent modulo $\mathfrak{r}_\ell$, then they are congruent modulo *every* $\mathfrak{r}_{\ell'}$, and hence are equal by the CRT isomorphism. Similarly, the homomorphism is surjective: given any $z_\ell \in \mathcal{O}_L/\mathfrak{r}_\ell$, we construct the $z \in \mathcal{O}_L/\mathfrak{r}\mathcal{O}_L$ whose CRT representation has ℓ'th entry $z_{\ell'} = (\ell' \circ \ell^{-1})(z_\ell) \in \mathcal{O}_L/\mathfrak{r}_{\ell'}$. By construction, $z = z_\ell \pmod{\mathfrak{r}_\ell}$. Moreover, we have that $z \in \mathcal{O}_K/\mathfrak{r}$ because it is fixed by $G_{L/K}$: for any $\tau, \ell' \in G_{L/K}$, the CRT representation of $\tau(z)$ has ℓ'th entry

$$\tau(z) + \mathfrak{r}_{\ell'} = \tau(z + \mathfrak{r}_{\tau^{-1} \circ \ell'}) = \tau(z_{\tau^{-1} \circ \ell'}) = \tau\big((\tau^{-1} \circ \ell' \circ \ell^{-1})(z_\ell)\big) = (\ell' \circ \ell^{-1})(z_\ell) = z_{\ell'} ,$$

so (the CRT representations of) $\tau(z)$ and z are equal, as claimed.

Now recall that any element $x \in \mathcal{O}_L/\mathfrak{r}\mathcal{O}_L$ can be represented uniquely as $x = \sum_{\ell \in G_{L/K}} c_\ell \cdot x_\ell$, where $(x_\ell \in \mathcal{O}_L/\mathfrak{r}_\ell)_{\ell \in G_{L/K}}$ is the CRT representation of x. The final claim follows by applying the inverses of the above natural ring isomorphisms to the x_ℓ (respectively), to obtain the unique coefficient vector $\vec{x}$ over $\mathcal{O}_K/\mathfrak{r}$ for which $x = \langle \vec{c}, \vec{x} \rangle$. □

Remark 5.6 (Fast computation in the CRT basis). By Lemma 5.5, the CRT basis enables fast addition and multiplication in $\mathcal{O}_L/\mathfrak{r}\mathcal{O}_L$ via the corresponding operations in $\mathcal{O}_K/\mathfrak{r}$. Specifically, if $x = \langle \vec{c}, \vec{x} \rangle, y = \langle \vec{c}, \vec{y} \rangle$ for coefficient vectors $\vec{x}, \vec{y}$ over $\mathcal{O}_K/\mathfrak{r}$, then by the properties of the CRT basis we have that $x + y = \langle \vec{c}, \vec{x} + \vec{y} \rangle$ and $x \cdot y = \langle \vec{c}, \vec{x} \odot \vec{y} \rangle$, where $\odot$ denotes the component-wise (Hadamard) product. It also enables fast evaluation of automorphisms: because any $\tau \in G_{L/K}$ permutes $\vec{c}$ (since $G_{L/K}$ acts regularly on it) via $\tau(c_\ell) = c_{\tau \circ \ell}$, we have that $\tau(\vec{c}) = P_\tau \cdot \vec{c}$ for the permutation matrix P_τ whose entries $(\ell, \tau \circ \ell)$ are 1 for all $\ell \in G_{L/K}$ (and the rest are zero). Since τ fixes $\vec{x}$ (because it fixes K pointwise), we have that $\tau(x) = \langle \tau(\vec{c}), \vec{x} \rangle = \langle \vec{c}, P_\tau^t \cdot \vec{x} \rangle$.

Here we show that under the setup of Lemma 2.2, the CRT bases for corresponding extensions *coincide*, which is useful both conceptually and computationally. In other words, the CRT basis of an extension "lifts" to any corresponding "higher, parallel" extension (see also Remark 5.12 for additional consequences).

Lemma 5.7. *Let L_1, L_2 be abelian number fields with $M = L_1 L_2$ and $K = L_1 \cap L_2$, let $\mathfrak{r}$ be a proper ideal of $\mathcal{O}_K$ that splits completely in $\mathcal{O}_{L_1}$ as $\mathfrak{r}\mathcal{O}_{L_1} = \prod_{\ell \in \mathrm{Gal}(L_1/K)} \mathfrak{r}_\ell$, and let $\vec{c}$ be the corresponding mod-$\mathfrak{r}$ CRT basis of $\mathcal{O}_{L_1}/\mathcal{O}_K$. Then:*

- *$\mathfrak{r}\mathcal{O}_{L_2}$ splits completely in $\mathcal{O}_M$, as $\mathfrak{r}\mathcal{O}_M = \prod_{m \in \mathrm{Gal}(M/L_2)} \mathfrak{r}_m$ where $\mathfrak{r}_m = \mathfrak{r}_\ell \mathcal{O}_M$ for the restriction $\ell = m|_{L_1}$, and*
- *the corresponding mod-$\mathfrak{r}\mathcal{O}_{L_2}$ CRT basis of $\mathcal{O}_M/\mathcal{O}_{L_2}$ is $\vec{c} + \mathfrak{r}\mathcal{O}_M$, reindexed according to the restrict-to-L_1 isomorphism from $\mathrm{Gal}(M/L_2)$ to $\mathrm{Gal}(L_1/K)$ (see Lemma 2.2).*

Proof. Let $G_{L_1/K} := \mathrm{Gal}(L_1/K)$ and $G_{M/L_2} := \mathrm{Gal}(M/L_2)$. We first show the claimed complete splitting in $\mathcal{O}_M$. Indeed, the $\mathfrak{r}_m$ are pairwise coprime (in $\mathcal{O}_M$) because the $\mathfrak{r}_\ell$ are (in $\mathcal{O}_{L_1}$); we have that

$$\prod_{m \in G_{M/L_2}} \mathfrak{r}_m = \prod_{\ell \in G_{L_1/K}} \mathfrak{r}_\ell \mathcal{O}_M = \mathfrak{r}\mathcal{O}_{L_1}\mathcal{O}_M = \mathfrak{r}\mathcal{O}_{L_2}\mathcal{O}_M \ ;$$

and for any $m' \in G_{M/L_2}$ with restriction $\ell' \in G_{L_1/K}$,

$$m'(\mathfrak{r}_m) = m'(\mathfrak{r}_\ell \mathcal{O}_M) = \ell'(\mathfrak{r}_\ell) \cdot \mathcal{O}_M = \mathfrak{r}_{\ell' \circ \ell} \cdot \mathcal{O}_M = \mathfrak{r}_{m' \circ m} \ .$$

Finally, since $\vec{c} = \vec{e}_\ell \pmod{\mathfrak{r}_\ell}$ for all $\ell \in G_{L_1/K}$, we have that $\vec{c} + \mathfrak{r}\mathcal{O}_M = \vec{e}_m \pmod{\mathfrak{r}_m}$ for all $m \in G_{M/L_2}$ under the stated reindexing, as needed. □

5.2 Factorization of CRT Bases

In this section we show that for a tower of extensions, the CRT basis factors as the Kronecker product of CRT bases for each step of the tower.

For the rest of this section, let $M/L/K$ be a tower of number fields where M/K is abelian and hence so are M/L and L/K, and define $G_{M/K} := \mathrm{Gal}(M/K)$, $G_{M/L} := \mathrm{Gal}(M/L)$, and $G_{L/K} := \mathrm{Gal}(L/K)$. By Lemma 2.1, restricting $G_{M/K}$ to L induces an isomorphism $\rho \colon G_{M/K}/G_{M/L} \to G_{L/K}$. This yields the following bijective correspondence.

Definition 5.8. Fix a transversal $T \subseteq G_{M/K}$ of $G_{M/K}/G_{M/L}$, and define the following bijective mapping between $G_{M/K}$ and $G_{M/L} \times G_{L/K}$: any $(m, \ell) \in G_{M/L} \times G_{L/K}$ corresponds to $\phi(m, \ell) := m \circ t \in G_{M/K}$, where $t \in T \cap \rho^{-1}(\ell)$ is the (unique) representative element of T that restricts to ℓ.

Lemma 5.9. *We have that $m' \circ \phi(m, \ell) = \phi(m' \circ m, \ell)$ for any $m' \in G_{M/L}$ and $(m, \ell) \in G_{M/L} \times G_{L/K}$.*

Proof. This follows immediately from $m' \circ \phi(m, \ell) = (m' \circ m) \circ t$, where t is the representative of ℓ in T.□

We stress that in general, the correspondence from Definition 5.8 is *not* a group isomorphism between $G_{M/K}$ and $G_{M/L} \times G_{L/K}$ *as a product group*. Instead, one can verify that it is an isomorphism under the group law

$$(m, \ell) \diamond (m', \ell') = (c \circ m \circ m', \ell \circ \ell') \tag{5.1}$$

for the "carry" element $c = t \circ t' \circ \tilde{t}^{-1} \in G_{M/L}$, where $t, t', \tilde{t} \in T$ are the unique representatives for $\ell, \ell', (\ell \circ \ell') \in G_{L/K}$, respectively. Observe that c depends only on ℓ, ℓ' (not m or m'), and that it is an element of $G_{M/L}$ because $c|_L = \ell \circ \ell' \circ (\ell \circ \ell')^{-1}$ is identity.

Our next lemma shows that, analogously to the situation for completely splitting *prime* ideals, generalized complete splitting for M/K implies the same for both M/L and L/K.

Lemma 5.10. *Let $\mathfrak{r}$ be a proper ideal of $\mathcal{O}_K$ that splits completely in $\mathcal{O}_M$, as*

$$\mathfrak{r}\mathcal{O}_M = \prod_{\tau \in G_{M/K}} \mathfrak{r}_\tau = \prod_{\substack{m \in G_{M/L} \\ \ell \in G_{L/K}}} \mathfrak{r}_{m,\ell} \tag{5.2}$$

where $\mathfrak{r}_{m,\ell} := \mathfrak{r}_\tau$ for the corresponding $\tau = \phi(m,\ell)$ from Definition 5.8. Then:

- *$\mathfrak{r}\mathcal{O}_L$ splits completely in $\mathcal{O}_M$, as $\mathfrak{r}\mathcal{O}_M = \prod_{m \in G_{M/K}} \mathfrak{r}_m$ where $\mathfrak{r}_m := \prod_{\ell \in G_{L/K}} \mathfrak{r}_{m,\ell}$,*
- *$\mathfrak{r}$ splits completely in $\mathcal{O}_L$, as $\mathfrak{r}\mathcal{O}_L = \prod_{\ell \in G_{L/K}} \mathfrak{r}_\ell$ where $\mathfrak{r}_\ell := \mathcal{O}_L \cap \prod_{m \in G_{M/L}} \mathfrak{r}_{m,\ell}$, and*
- *each $\mathfrak{r}_\ell$ splits completely in $\mathcal{O}_M$, as $\mathfrak{r}_\ell \mathcal{O}_M = \prod_{m \in G_{M/L}} \mathfrak{r}_{m,\ell}$.*

Furthermore, $\mathfrak{r}_m + \mathfrak{r}_\ell \mathcal{O}_M = \mathfrak{r}_{m,\ell}$ for every $m \in G_{M/L}$ and $\ell \in G_{L/K}$.

Proof. First, we show the complete spliting of $\mathfrak{r}\mathcal{O}_L$ in $\mathcal{O}_M$. By hypothesis and definition of $\mathfrak{r}_m$, we have the factorization $\mathfrak{r}\mathcal{O}_M = \prod_{m \in G_{M/L}} \mathfrak{r}_m$, and the $\mathfrak{r}_m$ for $m \in G_{M/L}$ are pairwise coprime because the $\mathfrak{r}_{m,\ell}$ are. Lastly, the $\mathfrak{r}_m$ are conjugates under $G_{M/L}$ (with respect to the given indexing): for any $m, m' \in G_{M/L}$, by Lemma 5.9 and the fact that the $\mathfrak{r}_{m,\ell}$ are conjugates under $G_{M/K}$ (with respect to the given indexing),

$$m'(\mathfrak{r}_m) = m'\Big(\prod_{\ell \in G_{L/K}} \mathfrak{r}_{m,\ell}\Big) = \prod_{\ell \in G_{L/K}} \mathfrak{r}_{m' \circ m, \ell} = \mathfrak{r}_{m' \circ m} \ .$$

For the remaining claims it is helpful to use *norms* of ideals. Because M/L is Galois, its ideal-norm function can be defined as $\mathrm{N}_{M/L}(\mathfrak{a}) := \mathcal{O}_L \cap \prod_{m \in G_{M/L}} m(\mathfrak{a})$ for any ideal of $\mathcal{O}_M$, which satisfies $\mathrm{N}_{M/L}(\mathfrak{a}) \cdot \mathcal{O}_M = \prod_{m \in G_{M/L}} m(\mathfrak{a})$, and similarly for the other extensions. These functions are transitive down the tower: $\mathrm{N}_{M/K} = \mathrm{N}_{L/K} \circ \mathrm{N}_{M/L}$.

Now we show the complete splitting of each $\mathfrak{r}_\ell$ in $\mathcal{O}_M$. By hypothesis, the $\mathfrak{r}_{m,\ell}$ for $m \in G_{M/L}$ are pairwise coprime, and are conjugates under $G_{M/L}$ (with respect to the given indexing) by Lemma 5.9. So by definition, $\mathfrak{r}_\ell = \mathrm{N}_{M/L}(\mathfrak{r}_{m,\ell})$ for any $m \in G_{M/L}$, and hence $\mathfrak{r}_\ell \mathcal{O}_M = \prod_{m \in G_{M/L}} \mathfrak{r}_{m,\ell}$, as needed.

For the complete splitting of $\mathfrak{r}$ in $\mathcal{O}_L$, first observe that the $\mathfrak{r}_\ell \mathcal{O}_M$ for $\ell \in G_{L/K}$ are pairwise coprime because the $\mathfrak{r}_{m,\ell}$ are, hence the $\mathfrak{r}_\ell$ are as well. (Formally, $(\mathfrak{r}_\ell + \mathfrak{r}_{\ell'})\mathcal{O}_M = \mathfrak{r}_\ell \mathcal{O}_M + \mathfrak{r}_{\ell'} \mathcal{O}_M = \mathcal{O}_M$ for distinct ℓ, ℓ', and intersecting both sides with $\mathcal{O}_L$ yields the claim.) Next, the $\mathfrak{r}_\ell$ are conjugates under $G_{L/K}$ (with respect to the given indexing): for any $\ell, \ell' \in G_{L/K}$, taking any $\tau' \in G_{M/K}$ that restricts to ℓ', by the fact that the $\mathfrak{r}_{m,\ell}$ are conjugates under $G_{M/K}$ (with respect to the given indexing) and the group law in Equation (5.1),

$$\ell'(\mathfrak{r}_\ell) = \mathcal{O}_L \cap \tau'\Big(\prod_{m \in G_{M/L}} \mathfrak{r}_{m,\ell}\Big) = \mathcal{O}_L \cap \prod_{\tilde{m} \in G_{M/L}} \mathfrak{r}_{\tilde{m}, \ell' \circ \ell} = \mathfrak{r}_{\ell' \circ \ell} \ .$$

Finally, because $\mathfrak{r} = \mathrm{N}_{M/K}(\mathfrak{r}_{m,\ell})$ (seen by intersecting both sides of Equation (5.2) with $\mathcal{O}_K$) and $\mathfrak{r}_\ell = \mathrm{N}_{M/L}(\mathfrak{r}_{m,\ell})$ for any $m \in G_{M/L}$ and $\ell \in G_{L/K}$, by the transitivity of the ideal norm we have the needed factorization

$$\mathfrak{r}\mathcal{O}_L = \mathrm{N}_{M/K}(\mathfrak{r}_{m,\ell}) \cdot \mathcal{O}_L = \mathrm{N}_{L/K}(\mathfrak{r}_\ell) \cdot \mathcal{O}_L = \prod_{\ell \in G_{L/K}} \mathfrak{r}_\ell \ .$$

For the last claim, by the pairwise coprimality of the $\mathfrak{r}_{m,\ell}$ and the complete splitting of $\mathfrak{r}_\ell$ in $\mathcal{O}_M$,

$$\mathfrak{r}_m + \mathfrak{r}_\ell \mathcal{O}_M = \prod_{\ell' \in G_{L/K}} \mathfrak{r}_{m,\ell'} + \prod_{m' \in G_{M/L}} \mathfrak{r}_{m',\ell} = \mathfrak{r}_{m,\ell} \ .$$

□

The above results lead to a factorization of the mod-$\mathfrak{r}$ CRT basis of $\mathcal{O}_M/\mathcal{O}_K$ as the Kronecker product of the corresponding CRT bases of $\mathcal{O}_M/\mathcal{O}_L$ and $\mathcal{O}_L/\mathcal{O}_K$, as shown next in Lemma 5.11. We remark that [AP13] gave a similar factorization for the CRT bases of *cyclotomic* rings (but not more general abelian number rings). However, the factorization from [AP13] allows for a lot of arbitrariness in the definition of the "higher" component $\vec{c}_{M/L}$ (namely, the $\mathfrak{r}_{m,\ell}$ may be indexed arbitrarily and independently for each ℓ), whereas our factorization is *uniquely* defined by the choice of transversal T of $G_{M/K}/G_{M/L}$ (as in Definition 5.8). The present formulation turns out to be critical for expressing CRT transforms using relatively few automorphisms, as shown below in Section 5.3.

Lemma 5.11. *Adopt the setup and notation of Lemma 5.10. The mod-$\mathfrak{r}$ CRT basis $\vec{c}_{M/K}$ of $\mathcal{O}_M/\mathcal{O}_K$ factors, under the reindexing from Definition 5.8, as*

$$\vec{c}_{M/K} = \vec{c}_{M/L} \otimes \vec{c}_{L/K}\,,$$

where $\vec{c}_{L/K}$ is the mod-$\mathfrak{r}$ CRT basis of $\mathcal{O}_L/\mathcal{O}_K$, and $\vec{c}_{M/L}$ is the mod-$\mathfrak{r}\mathcal{O}_L$ CRT basis of $\mathcal{O}_M/\mathcal{O}_L$.

Proof. Let $\vec{c}' = \vec{c}_{M/L}$ and $\vec{c} = \vec{c}_{L/K}$. Note that by Definition 5.3, for all $m \in G_{M/L}$ and $\ell \in G_{L/K}$ we have that $\vec{c} = \vec{e}_\ell \pmod{\mathfrak{r}_{m,\ell}}$ because the same relation holds modulo $\mathfrak{r}_\ell \subseteq \prod_{m \in G_{M/L}} \mathfrak{r}_{m,\ell} \subseteq \mathfrak{r}_{m,\ell}$, and $\vec{c}' = \vec{e}_m \pmod{\mathfrak{r}_{m,\ell}}$ because the same relation holds modulo $\mathfrak{r}_m$. So, $\vec{c}' \otimes \vec{c} = \vec{e}_m \otimes \vec{e}_\ell = \vec{e}_{(m,\ell)} \pmod{\mathfrak{r}_{m,\ell}}$. Applying the reindexing from Definition 5.8, we have that $\vec{c}' \otimes \vec{c} = \vec{e}_\tau \pmod{\mathfrak{r}_\tau}$ for all $\tau \in G_{M/K}$, and the claim follows. □

Remark 5.12. Adopting the setup and notation of Lemma 5.7, the result of Lemma 5.11 also implies that

$$\vec{c}_{M/K} = \vec{c}_{M/L_1} \otimes \vec{c}_{M/L_2} = \vec{c}_{L_2/K} \otimes \vec{c}_{L_1/K} \pmod{\mathfrak{r}\mathcal{O}_M}$$

(with appropriate reindexing in each case). Each kind of factorization can be convenient for certain purposes, e.g., computing CRT bases more efficiently by working in smaller-dimensional fields, or analyzing the effect of automorphisms on the CRT basis.

Notably, the former factorization, together with the direct-product factorization of the Galois group $G_{M/K} = G_{M/L_1} \times G_{M/L_2}$, allows us to view the CRT "slots" as arranged in a two-dimensional array, where G_{M/L_i} acts regularly along the ith dimension (and has no effect on the other dimension). Naturally, even finer-grained factorizations, arising from composites of abelian number fields having a common pairwise intersection field, correspond to even higher-dimensional arrays, i.e., tensors. This allows us to design a field that induces a desired tensor shape, and to permute its contents in structured ways, e.g., for homomorphic linear algebra and other algorithms [HS14, HS18].

5.3 Sparsity of Automorphism Coefficients

For the mod-$\mathfrak{r}\mathcal{O}_L$ CRT basis $\vec{c} = \vec{c}_{M/L}$ of $\mathcal{O}_M/\mathcal{O}_L$, consider the linear function f_K from the second step of the "top-down" structured transform in Section 4.2. Recalling that $\vec{c}^\vee = \vec{c} \bmod \mathfrak{r}\mathcal{O}_M^\vee$, here we analyze the factor $\langle \vec{c}, \tau(\vec{c}^\vee)\rangle = \langle \vec{c}, \tau(\vec{c})\rangle \bmod \mathfrak{r}\mathcal{O}_M^\vee$ that appears in the coefficient of $\tau \in G_{M/K}$ in Equation (4.4) from Lemma 4.3, when expressing f_K as a linear combination of automorphisms. Specifically, we identify a necessary condition on τ for when this factor is nonzero, which implies that f_K can be expressed *sparsely* in terms of automorphisms. Note that it suffices to analyze the term $\langle \vec{c}, \tau(\vec{c})\rangle$ *without* the reduction modulo $\mathfrak{r}\mathcal{O}_M^\vee$, because the reduced term (which is what we ultimately care about) is nonzero only if the non-reduced term is nonzero, since $\mathfrak{r}\mathcal{O}_M \subseteq \mathfrak{r}\mathcal{O}_M^\vee$.

Lemma 5.13. *Adopt the reindexing of Definition 5.8 and the notation of Lemma 5.10, and let $\vec{c}$ be the mod-$\mathfrak{r}\mathcal{O}_L$ CRT basis of $\mathcal{O}_M/\mathcal{O}_L$. Then $\langle \vec{c}, \tau(\vec{c}) \rangle$ is nonzero (modulo $\mathfrak{r}\mathcal{O}_M$) only if $\tau = t' \circ t^{-1}$ for some $t, t' \in T$.*

Proof. Since $\langle \vec{c}, \tau(\vec{c}) \rangle \neq 0 \pmod{\mathfrak{r}\mathcal{O}_M}$ and the modulus factors as the product of the pairwise coprime $\mathfrak{r}_m$, it follows that $\langle \vec{c}, \tau(\vec{c}) \rangle \neq 0 \pmod{\mathfrak{r}_m}$ for some $m \in G_{M/L}$. Because $\vec{c} = \vec{e}_m \pmod{\mathfrak{r}_m}$ by definition, it must be that $\tau(c_m) \neq 0 \pmod{\mathfrak{r}_m}$. So, by definition of $\mathfrak{r}_m$, $\tau(c_m) \neq 0 \pmod{\mathfrak{r}_{t' \circ m}}$ for some $t' \in T$. Because $c_m \neq 0 \pmod{\mathfrak{r}_{t \circ m}}$—or equivalently, $\tau(c_m) \neq 0 \pmod{\mathfrak{r}_{\tau \circ t \circ m}}$—only if $t \in T$, it must hold that $\tau \circ t \circ m = t' \circ m$ for some $t \in T$. This implies that $\tau = t' \circ t^{-1}$, as claimed. □

Number of nonzero coefficients. We now use Lemma 5.13 to concretely bound the number of automorphisms that suffice for evaluating each linear function f_K from Section 4.2, in several settings of interest. By Lemma 4.2, the total number of automorphisms to homomorphically evaluate a complete CRT transform is just the sum of these over each step of the relevant tower. The following material shows that a worst-case bound on this total is $O(d^2 \log n)$, and $O(d \log n)$ is frequently achievable, where d is an upper bound on the degree of each step and n is the degree (over $\mathbb{Q}$) of the number field at the top of the tower.

Recall from Section 4.2 that in this context (and in contrast to Section 5.4 below), $\deg(L/K)$ is typically taken to be small, but $\deg(M/L)$ can be large (because we consider "top down" transforms to the CRT basis). Since $|T| = |G_{M/K}/G_{M/L}| = \deg(L/K)$, there are at most $|T|^2 = \deg(L/K)^2$ values of $\tau = t' \circ t^{-1}$ for which $\langle \vec{c}, \tau(\vec{c}) \rangle$ is nonzero. Moreover, in many cases of interest, the number of distinct $t' \circ t^{-1}$ can be significantly smaller than $|T|^2$.

1. The most favorable case is when we have a direct product $G_{M/K} = G_{M/L} \times G_{M/L'}$, which by the Galois correspondence holds if and only if $M = LL'$ and $L \cap L' = K$ (as in the setup of Lemma 2.2). We can then let T be the subgroup $G_{M/L'}$, so the number of distinct $t' \circ t^{-1}$ is only $|T| = \deg(M/L') = \deg(L/K)$.

   In particular, this case applies when the degrees of M/L and L/K are coprime (by the fundamental theorem of finite abelian groups), such as when M has odd prime-power conductor and the conductor of L is that prime.

2. Even if T is not a subgroup, there can be many duplicates among the $t' \circ t^{-1}$. As a common case, if $G_{M/K}/G_{M/L} \cong G_{L/K} \cong \mathbb{Z}/d\mathbb{Z}$ is *cyclic* of small order $d = \deg(L/K)$, we can let T correspond to $\{0, 1, \ldots, d-1\} \subset \mathbb{Z}$, in which case every $t' \circ t^{-1}$ corresponds to an element of $\{-d+1, \ldots, d-1\} \subset \mathbb{Z}$, which has cardinality $2|T| - 1$.

   In particular, this case applies if L has odd prime-power conductor, because $\mathrm{Gal}(L/\mathbb{Q})$ is cyclic; if L is a power-of-two cyclotomic with $K \supseteq \mathbb{Q}(\zeta_4)$ (or $L = \mathbb{Q}(\zeta_4), K = \mathbb{Q}$, though this even falls under the previous item); or if L is totally real with power-of-two conductor. These cases significantly generalize [CCS19], which obtained the same sparsity solely for towers of power-of-two cyclotomics, with a complex-CRT basis corresponding to the canonical embedding (not modulo an ideal), in the context of approximate FHE [CKKS17].

3. Finally, and analogously to Lemma 5.7, we can generically "lift" or "lower" a transversal of one quotient of Galois groups to corresponding "parallel" one. Specifically, let F/E be an extension of abelian number fields for which $M \cap KF = K$, and hence $L \cap KE = K$. Then by Lemma 2.2, we may lift the elements of any transversal T of $G_{M/K}/G_{M/L}$ from $G_{M/K}$ to $G_{MF/KF} \subseteq G_{MF/KE}$. This mapping is an isomorphism, so it preserves the number of distinct $t' \circ t^{-1}$. Furthermore,

it induces a homomorphism $G_{M/K} \mapsto G_{MF/KE}/G_{MF/LE}$ with kernel $G_{M/L}$, and thus also an injective homomorphism from $G_{M/K}/G_{M/L}$ to $G_{MF/KE}/G_{MF/LE}$. Now by Lemmas 2.1 and 2.2, $G_{MF/KE}/G_{MF/LE} \cong G_{LE/KE} \cong G_{L/K} \cong G_{M/K}/G_{M/L}$, so the previous injective homomorphism is actually an isomorphism, and therefore the original isomorphism $G_{M/K} \mapsto G_{MF/KF}$ maps any transversal of $G_{M/K}/G_{M/L}$ to one of $G_{MF/KE}/G_{MF/LE}$.

In particular, this case applies when the conductors of M and F are coprime. So, one can proceed modularly by separately considering various M whose conductors are powers of distinct primes, finding suitable transversals for the steps of their towers, and then finally "lifting" them to the Galois group of the composite field to get transversals for all the corresponding steps of the composite tower.

5.4 Fast "In the Clear" CRT Transforms

Here we give sparse decompositions of the linear CRT transforms between CRT bases and any similarly structured bases, which operate "in the clear"—i.e., directly on coordinate vectors (as opposed to homomorphically via automorphisms, in Section 5.3). These directly yield fast, highly parallel algorithms for converting between coordinate vectors relative to these bases. They generalize the prior Number Theoretic Transform (NTT) and Chinese Remainder Transform (CRT), which go between the CRT and power/"powerful" bases of *cyclotomics*, to fast transforms that go between CRT bases and *any* similarly structured basis of *any abelian* number field (equivalently, any subfield of a cyclotomic).

In the present context, a sparse decomposition for a change-of-basis transform is a factorization of its associated matrix into a product of a small number of sparse matrices. In our decompositions, each sparse matrix is of the form $I_l \otimes T \otimes I_r$, where T is some small-dimensional (but typically dense) square matrix, and I_l, I_r are identity matrices of certain dimensions (all of which vary from one sparse matrix to another). Multiplying such a sparse matrix with an input vector can be done efficiently (and with high parallelism) simply by multiplying T with each "strided block" of the input. So, the entire transform can be evaluated by multiplying by all the factors of the sparse decomposition in sequence.

Sparse decomposition of CRT transforms. In this context (and in contrast to Section 5.3) we typically want $\deg(M/L)$ to be small, but $\deg(L/K)$ can be large (because we give "bottom up" transforms to the CRT basis). Let $\mathfrak{r}$ be a proper ideal of $\mathcal{O}_K$ that splits completely in $\mathcal{O}_M$, let $\vec{c}_{M/K}$ be the corresponding mod-$\mathfrak{r}$ CRT basis of $\mathcal{O}_M/\mathcal{O}_K$, and let $\vec{c}_{M/K} = \vec{c}_{M/L} \otimes \vec{c}_{L/K}$ be its factorization from Lemma 5.11.

Similarly, let $\vec{b}_{M/K} = \vec{b}_{M/L} \otimes \vec{b}_{L/K}$ be any structured $(\mathcal{O}_K/\mathfrak{r})$-basis of $\mathcal{O}_M/\mathfrak{r}\mathcal{O}_M$, where $\vec{b}_{M/L}$ is an $(\mathcal{O}_L/\mathfrak{r}\mathcal{O}_L)$-basis of $\mathcal{O}_M/\mathfrak{r}\mathcal{O}_M$, and $\vec{b}_{L/K}$ is an $(\mathcal{O}_K/\mathfrak{r})$-basis of $\mathcal{O}_L/\mathfrak{r}\mathcal{O}_L$. In particular, this holds if $\vec{b}_{M/L}$ is an $\mathcal{O}_L$-basis of $\mathcal{O}_M$ that has been reduced modulo $\mathfrak{r}\mathcal{O}_L$, and similarly for $\vec{b}_{L/K}$. As a primary example, the short integral bases from Theorem 2.3 have this Kronecker-product structure.

We give a sparse decomposition for the $(\mathcal{O}_K/\mathfrak{r})$-linear transform that maps $\vec{b}_{M/K}$ to $\vec{c}_{M/K}$. (A sparse decomposition for the "dual" transform on $\mathcal{O}_M^\vee/\mathfrak{r}\mathcal{O}_M^\vee$ has an analogous decomposition.) In brief, it proceeds in a "bottom-up" fashion, as the composition of two stages of sparse transforms:

$$\vec{b}_{M/K} = \vec{b}_{M/L} \otimes \vec{b}_{L/K} \xrightarrow{I \otimes T_{L/K}} \vec{b}_{M/L} \otimes \vec{c}_{L/K} \xrightarrow{T'_{M/L}} \vec{c}_{M/L} \otimes \vec{c}_{L/K} = \vec{c}_{M/K} \,.$$

The formalization is as follows.

Theorem 5.14. *Let $T_{M/K}$ be the $\vec{b}_{M/K}$-to-$\vec{c}_{M/K}$ (change of basis) matrix over $\mathcal{O}_K/\mathfrak{r}$, satisfying $\vec{b}_{M/K}^t = \vec{c}_{M/K}^t \cdot T_{M/K}$. It has the "bottom-up" sparse decomposition*

$$T_{M/K} = T'_{M/L} \cdot (I_{M/L} \otimes T_{L/K}) \,,$$

where $I_{M/L}$ is the $\deg(M/L)$-dimensional identity matrix, $T_{L/K}$ is the $\vec{b}_{L/K}$-to-$\vec{c}_{L/K}$ matrix over $\mathcal{O}_K/\mathfrak{r}$, and

$$T'_{M/L} = \mathrm{Tr}_{L/K}(T_{M/L} \otimes \mathrm{diag}(\vec{c}^{\vee}_{L/K})) \,,$$

where $T_{M/L}$ is the $\vec{b}_{M/L}$-to-$\vec{c}_{M/L}$ matrix over $\mathcal{O}_L/\mathfrak{r}\mathcal{O}_L$.

We give the (fairly routine) proof below, after discussing the implications. Observe that $I_{M/L} \otimes T_{L/K}$ is block diagonal, and hence (at least somewhat) sparse: it simply applies $T_{L/K}$ to each of the $\deg(M/L)$ blocks of $\deg(L/K)$ coordinates in the input vector. However, recall that $T_{L/K}$ may have large dimension $\deg(L/K)$. Yet if L/K has an intermediate field and $\vec{b}_{L/K}$ factors correspondingly, then $T_{L/K}$ can be sparsely decomposed in the same way, inductively. In addition, $T'_{M/L}$ is also sparse, since it is a block matrix of (typically small) dimension $\deg(M/L)$, with diagonal blocks of dimension $\deg(L/K)$.

Overall, after sparsely decomposing $T_{L/K}$ inductively, this approach uses a tower $K_t/\cdots/K_1/K_0$ of abelian number fields, each typically of small relative degree, and decomposes the CRT transform as the following composition of sparse-transform stages:

$$\vec{b}_t \otimes \cdots \otimes \vec{b}_2 \otimes \vec{b}_1 \to \vec{b}_t \otimes \cdots \otimes \vec{b}_2 \otimes \vec{c}_1 \to \vec{b}_t \otimes \cdots \otimes \vec{c}_2 \otimes \vec{c}_1 \to \cdots \to \vec{c}_t \otimes \vec{c}_{t-1} \otimes \cdots \otimes \vec{c}_1 \,.$$

The inverse transform, which maps $\vec{c}_{M/K}$ to $\vec{b}_{M/K}$, works simply by inverting the stages in reverse ("top-down") order, and has essentially the same complexity.

Complexity. In the unfolded decomposition, the ith stage can be implemented via multiplication by a structured matrix over $\mathcal{O}_{K_0}/\mathfrak{r}$—specifically, the Kronecker product of an identity matrix and one with diagonal blocks—having at most $n \cdot \deg(K_i/K_{i-1})$ nonzero entries, where $n = \deg(K_t/K_0)$ is the dimension of the input vector. Therefore, the overall complexity of the full transform is $n \cdot \sum_{i=1}^{t} \deg(K_i/K_{i-1})$ multiplications and additions in $\mathcal{O}_{K_0}/\mathfrak{r}$. For example, if each $\deg(K_i/K_{i-1})$ is bounded by a constant, then $t = O(\log n)$ and the overall complexity is $O(n \log n)$. In addition, each stage of the transform is parallelizable in the natural way, due to the sparse structure of its matrix.

Proof of Theorem 5.14. For the first stage, by the mixed-product property, $I \otimes T_{L/K}$ is the change-of-basis matrix from $\vec{b}_{M/L} \otimes \vec{b}_{L/K}$ to $\vec{b}_{M/L} \otimes \vec{c}_{L/K}$. Concretely, by the material in Section 2.2.4,

$$T_{L/K} = \mathrm{Tr}_{L/K}(\vec{c}^{\vee}_{L/K} \cdot \vec{b}^{t}_{L/K}) \quad \text{(over } \mathcal{O}_K/\mathfrak{r}\text{)}.$$

For the second stage, $T'_{M/L}$ is the change-of-basis matrix over $\mathcal{O}_K/\mathfrak{r}$ from $\vec{b}_{M/L} \otimes \vec{c}_{L/K}$ to $\vec{c}_{M/L} \otimes \vec{c}_{L/K}$, which we derive as follows. By hypothesis and Section 2.2.4, $\vec{b}^{t}_{M/L} = \vec{c}^{t}_{M/L} \cdot T_{M/L}$ where

$$T_{M/L} = \mathrm{Tr}_{M/L}(\vec{c}^{\vee}_{M/L} \cdot \vec{b}^{t}_{M/L}) \quad \text{(over } \mathcal{O}_L/\mathfrak{r}\mathcal{O}_L\text{)}.$$

Since the entries of $T_{M/L}$ during this stage already are represented in the CRT basis $\vec{c}_{L/K}$, they each expand into a diagonal matrix over $\mathcal{O}_K/\mathfrak{r}$.[24] Formally, by Section 2.2.4, transitivity of $\mathrm{Tr}_{M/K} = \mathrm{Tr}_{L/K} \circ \mathrm{Tr}_{M/L}$, L-linearity of $\mathrm{Tr}_{M/L}$, and self-duality of the CRT basis (Remark 5.4), the full change-of-basis matrix is

$$T'_{M/L} = \mathrm{Tr}_{M/K}\big((\vec{c}^{\vee}_{M/L} \cdot \vec{b}^{t}_{M/L}) \otimes (\vec{c}^{\vee}_{L/K} \cdot \vec{c}^{t}_{L/K})\big) = \mathrm{Tr}_{L/K}\big(T_{M/L} \otimes \mathrm{diag}(\vec{c}^{\vee}_{L/K})\big) \,. \qquad \square$$

[24] This is why we decompose the transform in a *bottom-up* fashion, mapping the "lower" component $\vec{b}_{L/K}$ to the CRT-basis component $\vec{c}_{L/K}$ first, so that the $(\mathcal{O}_L/\mathfrak{r}\mathcal{O}_L)$-linear second stage can be implemented in a sparse way.

6 Putting Everything Together

Here we show how to combine our results in the service of Ring-LWE cryptography and homomorphic encryption. These solutions adapt and generalize existing approaches for Ring-LWE cryptography [LPR10, LPR13] and packed bootstrapping [GHS12a, AP13], originally designed for cyclotomics, to the abelian number fields (cyclotomic subfields) considered in this work.

Recall that Section 3 gave a generic mathematical template for Ring-LWE-based homomorphic encryption and packed bootstrapping over arbitrary number rings. Using the results of the preceding sections, here we fill in the details of how the template can be instantiated computationally, in any ring of integers $R = \mathcal{O}_K$ of an abelian number field K for which we can construct short, structured bases of R and $R^\vee$ (e.g., by the results from [PP25] recalled in Section 2.3):

- in Section 6.1 we describe fast operations on the ring and (powers of) its dual ideal, based on the CRT basis and fast transforms between it and other structured bases;
- in Section 6.2 we show how a short basis for the ring R is used for noisy encoding of plaintexts in encryption, decoding in decryption, and "gadget decomposition" in key-switching;
- in Section 6.3 we instantiate homomorphic CRT transforms via automorphisms, which are central to packed bootstrapping.

6.1 Fast Ring and Dual Operations

We first show how the basic operations in (and between) R_q and $R_q^\vee$ can be performed efficiently, assuming that q (formally, the ideal $q\mathbb{Z}$) splits completely in R, using the mod-q CRT basis $\vec{c}$ of $R/\mathbb{Z}$ (see Section 5). Recall from Remark 5.6 that the CRT basis yields fast addition and multiplication in R_q via the corresponding component-wise operations in $\mathbb{Z}_q$, along with fast evaluation of automorphisms on R_q because any $\tau \in \mathrm{Gal}(K/\mathbb{Q})$ simply permutes $\vec{c}$.

These same algorithms also work equally well across the fractional "codifferent" ideal $\mathfrak{C} := R^\vee \supseteq R$ and its powers, modulo their scalings by q. Define the quotient $\mathfrak{C}_q^k := \mathfrak{C}^k/q\mathfrak{C}^k$, and $\vec{c}^{(k)} := \vec{c} \bmod q\mathfrak{C}^k$ to be the vector over $\mathfrak{C}_q^k$ obtained by natural inclusion and reduction modulo $q\mathfrak{C}^k \supseteq qR$. We call this the CRT basis of $\mathfrak{C}_q^k$ (it is indeed a $\mathbb{Z}_q$-basis).

Suppose that $x = \langle \vec{c}^{(k)}, \mathbf{x} \rangle \in \mathfrak{C}_q^k$ and $y = \langle \vec{c}^{(k')}, \mathbf{y} \rangle \in \mathfrak{C}_q^{k'}$ for some non-negative integers k, k' and vectors $\mathbf{x}, \mathbf{y}$ over $\mathbb{Z}_q$. Then they can be efficiently added and multiplied as follows:

$$\begin{aligned} x + y &= \langle \vec{c}^{(k)}, \mathbf{x} + \mathbf{y} \rangle \in \mathfrak{C}_q^k \quad \text{if } k = k', \text{ and} \\ x \cdot y &= \langle \vec{c}^{(k+k')}, \mathbf{x} \odot \mathbf{y} \rangle \in \mathfrak{C}_q^{k+k'} \ , \end{aligned}$$

where $\odot$ denotes the component-wise (Hadamard) product. The latter follows from the fact that because $\vec{c}$ is over R_q and $R \subseteq \mathfrak{C}$,

$$c_i^{(k)} \cdot c_j^{(k')} = (c_i + q\mathfrak{C}^k) \cdot (c_j + q\mathfrak{C}^{k'}) \subseteq c_i \cdot c_j + q\mathfrak{C}^{k+k'} = \delta_{i,j} \cdot c_i^{(k+k')} \ ,$$

where $\delta_{i,j}$ is taken to be in R_q. Similarly, automorphisms on $\mathfrak{C}_q^k$ can be evaluated in the same way as for R_q in the CRT basis: any $\tau \in \mathrm{Gal}(K/\mathbb{Q})$ permutes $\vec{c}^{(k)}$ just as it does $\vec{c}$, by definition of $\vec{c}^{(k)}$.

Other operations like noisy encoding, key-switching, and decryption rely on representing elements with respect to other short bases (see Section 6.2 below). Using the fast linear transforms given in Section 5.4, we can efficiently convert between the CRT basis $\vec{c}$ (or more generally, $\vec{c}^{(k)}$) and any similarly structured basis

arising from the tower of subrings. In particular, this includes the short, structured integral basis $\vec{b}$ of R (or its dual basis $\vec{b}^\vee$ of $\mathfrak{C}$) given by Theorem 2.3 and Lemma 2.5 (reduced modulo qR and $q\mathfrak{C}$, respectively). Similar remarks apply for the optimally short, structured basis of $R^\vee$ constructed in Theorem 2.4, though in general that basis is not quite as structured—it is not simply a Kronecker product of relative bases going down the tower—so more work is needed to get fast transforms; we refer to [PP25] for the details.

6.2 Encoding, Decoding, and Decomposition

Here we give more details for the (least-significant-digit) noisy encoding and decoding as used in the homomorphic encryption template of Section 3.1, using a *short* basis of R.[25] We also briefly describe how such a basis is used for "gadget decomposition" in key-switching. We first recall the standard noisy encoding and decoding for *integers*, then extend them to $R^\vee$ using suitable $\mathbb{Z}$-bases.

A noisy encoding of $\mu \in \mathbb{Z}_p$ is a "small" $e \in \mu + p\mathbb{Z}$, i.e., an integer $e \in \mathbb{Z}$ such that $e = \mu \pmod p$ and $|e| < q/2$. Note that the magnitude of a typical encoding is proportional to p. In cryptographic applications, encodings are typically reduced modulo some $q > p$. To reverse this, the decoding function $\mathsf{Decode}\colon \mathbb{Z}_q \to \mathbb{Z}_p$ lifts the argument to its smallest (in magnitude) integer representative e and outputs $\mu = e \bmod p$.

We can extend the above to $R_p^\vee$ in a coordinate-wise fashion relative to a suitable basis of $R^\vee$, as follows. Recall from Section 3.1.1 that a noisy encoding of $\mu \in R_p^\vee$ is a short $e \in \mu + pR^\vee$, i.e., $e \in R^\vee$ and $e = \mu \pmod{pR^\vee}$. Letting $\vec{e}$ (for "encoding basis") be a suitably short $\mathbb{Z}$-basis of $R^\vee$ and expressing $\mu = \langle \vec{e}, \boldsymbol{\mu} \rangle$ for a coordinate vector $\boldsymbol{\mu}$ over $\mathbb{Z}_p$, we can generate such an encoding as $e = \langle \vec{e}, \mathbf{e} \rangle \in R^\vee$, where the vector $\mathbf{e}$ over $\mathbb{Z}$ is a coordinate-wise noisy encoding of $\boldsymbol{\mu}$. (See below for some alternative, more sophisticated methods.) In applications, encodings of $R_p^\vee$ are typically reduced modulo $qR^\vee$.

Similarly, let $\vec{b}$ be a suitably short $\mathbb{Z}$-basis of R; then its dual basis $\vec{d} = \vec{b}^\vee$ (for "decoding basis") is a $\mathbb{Z}$-basis of $R^\vee$, and hence is also a $\mathbb{Z}_r$-basis of $R_r^\vee$ for any positive integer r. (We stress that $\vec{d}$ and $\vec{e}$ need not be the same basis.) The decoding function $\mathsf{Decode}\colon R_q^\vee \to R_p^\vee$ is defined analogously, as coordinate-wise integer decoding relative to $\vec{d}$. That is, on input $z \in R_q^\vee$, we express $z = \langle \vec{d}, \mathbf{z} \rangle$ and output $\mathsf{Decode}(z) := \langle \vec{d}, \mathsf{Decode}(\mathbf{z}) \rangle \in R_p^\vee$.

Lemma 6.1. *Let $e \in R^\vee$ be a noisy encoding of $\mu \in R_p^\vee$. If $\|e\| < q/(2\|\vec{b}\|)$ where $\|\vec{b}\| = \max_i \|b_i\|$, then $\mathsf{Decode}(e \bmod qR^\vee) = \mu$. Alternatively, if e is subgaussian with parameter $r \le q/(2\|\vec{b}\|\sqrt{\ln(2n/\delta)})$ where $n = \deg(K/\mathbb{Q})$ and $\delta > 0$, then $\mathsf{Decode}(e \bmod qR^\vee) = \mu$ except with probability at most δ.*

Proof. By hypothesis and definition of Decode, it suffices to show that $|e_i| < q/2$ for every entry e_i of the coordinate vector $\mathbf{e}$ over $\mathbb{Z}$, where $e = \langle \vec{d}, \mathbf{e} \rangle$. Recall from Sections 2.1.2 and 2.2.1 that $e_i = \mathrm{Tr}(b_i \cdot e) = \langle \overline{b_i}, e \rangle$. So, for the first claim, $|e_i| \le \|b_i\| \cdot \|e\| < q/2$ by Cauchy-Schwarz, where $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ are respectively the standard inner product and Euclidean norm in the canonical embedding of K. For the second claim, each e_i is subgaussian with parameter $r\|b_i\|$, and hence $|e_i| < r\|b_i\|\sqrt{\ln(2n/\delta)}$ except with probability at most δ/n. The claim follows by the union bound over the n coordinates of $\mathbf{e}$. □

Finally, we briefly mention how a short $\mathbb{Z}$-basis of R can be used for "gadget decomposition" in key-switching (and other applications). In brief, this operation decomposes an element $c \in R_q$ as short "digits" $c_i \in R$ with respect to some small integer base $g \ge 2$, as $c = \sum_i c_i \cdot g^i \pmod{qR}$. This can be done

[25] Note that encoding and decoding *themselves* do not need the basis to have any Kronecker-product structure, but such structure is used to convert quickly to this basis from another structured one.

coefficient-wise with respect to the short basis of R, expressing each $\mathbb{Z}_q$-coefficient of c in base g using digits in $[-g/2, g/2)$, say. Because the basis of R is short, so are the resulting c_i. (All this adapts beyond powers of g to other kinds of "gadgets" as well, like CRT gadgets.)

Other noisy encoders/error samplers. Recall from above that to generate a noisy encoding for $\mu \in R_p^\vee$, we need to sample a short error $e \in \mu + pR^\vee$. For security, the distribution should have sufficiently large width in the canonical embedding of K, and to best control noise growth under homomorphic operations, it should have a nearly "spherical" (i.e., isotropic) shape. (And to conform with worst-case hardness theorems, the distribution should be Gaussian.) Unfortunately, the above coordinate-wise noisy encoding can produce a fairly "skewed" non-spherical distribution, depending on the geometry of the encoding basis $\vec{e}$.

There are at least two alternative distributions that satisfy the above criteria: a true discrete Gaussian, and a rounded-off Gaussian. The former can be efficiently sampled using the generic algorithm from [GPV08] with a short basis $\vec{e}$ of $R^\vee$. Moreover, we are optimistic that for bases $\vec{e}$ with a Kronecker-product structure, the more efficient techniques for cyclotomics from [DP16] should adapt to our more general setting of abelian number fields; we leave this to future work.

The latter kind of distribution can be sampled by drawing from a continuous spherical Gaussian in the canonical embedding, then rounding it off to the desired coset by representing it relative to a short basis $\vec{e}$ of $R^\vee$. To support this, the companion paper [PP25] (see Theorem 2.4) constructs an *optimally short*, structured basis $\vec{e}$ of $R^\vee$ and an associated fast CRT-like transform between $\vec{e}$ and a known structured *orthonormal* $\mathbb{R}$-basis of the canonical embedding of K.[26] To sample an error, we first sample from a spherical Gaussian in the canonical embedding using the orthonormal basis, then apply the fast transform to represent it (with correlated real coordinates) in the basis $\vec{e}$ of $R^\vee$. Finally, we round the coordinates (deterministically or randomly) to get an element in $\mu + pR^\vee$; because $\vec{e}$ is optimally short, this increases the norm (or covariance) of the sample by relatively very little.

6.3 Fast Homomorphic CRT Transforms

Here we consider the number field K to be the "top" of some tower $K^{(\ell)}/K^{(\ell-1)}/\cdots/K^{(0)}$ of abelian number fields, where $K = K^{(\ell)}$ and $K^{(0)} = \mathbb{Q}$. This tower induces a tower of ring extensions $R^{(\ell)}/R^{(\ell-1)}/\cdots/R^{(0)}$ where $R^{(i)} = \mathcal{O}_{K^{(i)}}$, and in particular $R = R^{(\ell)}$ and $\mathbb{Z} = R^{(0)}$. Suppose that $q\mathbb{Z}$ splits completely in R, and let $\vec{c} = \vec{c}_\ell \otimes \vec{c}_{\ell-1} \otimes \cdots \otimes \vec{c}_1$ be the Kronecker-product factorization of the mod-q CRT basis $\vec{c}$ going down the tower (see Lemma 5.11), i.e., $\vec{c}_i$ is the mod-$qR^{(i-1)}$ CRT basis of $R^{(i)}/R^{(i-1)}$. Similarly, suppose that $\vec{b} = \vec{b}_\ell \otimes \vec{b}_{\ell-1} \otimes \cdots \otimes \vec{b}_1$ is a factorization of a $\mathbb{Z}$-basis $\vec{b}$ of R going down the tower, i.e., $\vec{b}_i$ is an $R^{(i-1)}$-basis of $R^{(i)}$. In particular, recall from Section 2.3 that Theorem 2.3 and Lemma 2.5 gives such (short) factored bases for a broad family of towers of abelian number fields.

Section 4 shows how to express the linear CRT transforms that map between the structured bases $\vec{b}$ and $\vec{c}$ using linear combinations of the automorphisms of the extensions $K^{(i)}/K^{(i-1)}$. By expressing the transforms in this way, they can be computed *homomorphically* on the plaintext using the operations recalled in Section 3.1.2, namely, linear operations and automorphisms. Moreover, Section 5.3 (and in particular Lemma 5.13) shows that for the CRT basis $\vec{c}$, this representation of the structured linear transforms is *sparse*, i.e., it can be expressed in terms of *relatively few* automorphisms. Together, this directly yields efficient algorithms for the homomorphic evaluation of CRT transforms.

[26] This structured orthonormal $\mathbb{R}$-basis is analogous, with a similar Kronecker-product structure, to the mod-$\mathfrak{q}$ CRT basis from Section 5, and the fast transform works similarly to the one given in Section 5.4, but over $\mathbb{R}$.

Recall from Section 3.2 that homomorphic CRT transforms make up the first and third steps of the packed bootstrapping template. However, recall that the plaintext is best encoded in $R_q^\vee$ (not R_q), so we actually want to homomorphically evaluate the "dual" CRT transforms, which map between $\vec{b}^\vee = \vec{b}_\ell^\vee \otimes \vec{b}_{\ell-1}^\vee \otimes \cdots \otimes \vec{b}_1^\vee$ and $\vec{c}^\vee = \vec{c}_\ell^\vee \otimes \vec{c}_{\ell-1}^\vee \otimes \cdots \otimes \vec{c}_1^\vee$, both of which are $\mathbb{Z}_q$-bases of $R_q^\vee$. Fortunately, the framework of Section 4 works just as well in this setting, simply by swapping $\vec{b}$ and $\vec{c}$ with their respective duals $\vec{b}^\vee$ and $\vec{c}^\vee$. Note that this replaces the factors $\langle \vec{c}_{M/L}, \tau(\vec{c}_{M/L}^\vee)\rangle$ appearing in Equation (4.4) with $\langle \vec{c}_{M/L}^\vee, \tau(\vec{c}_{M/L})\rangle$. Fortunately, $\langle \vec{c}_{M/L}^\vee, \tau(\vec{c}_{M/L})\rangle = \langle \vec{c}_{M/L}, \tau(\vec{c}_{M/L})\rangle \bmod \mathfrak{q}\mathcal{O}_M^\vee$ by the same reasoning given at the start of Section 5.3, so Lemma 5.13 yields the same level of sparsity for this setting.

References

[ADE+23] E. Aharoni, N. Drucker, G. Ezov, E. Kushnir, H. Shaul, and O. Soceanu. E2E near-standard and practical authenticated transciphering. Cryptology ePrint Archive, Paper 2023/1040, 2023. Page 2.

[AH17] S. Arita and S. Handa. Subring homomorphic encryption. In *ICISC*, pages 112–136. 2017. Page 3.

[ALJ+22] K. M. M. Aung, E. Lim, S. J. Jie, B. H. M. Tan, H. Wang, and S. L. Yeo. Field instruction multiple data. In *EUROCRYPT 2022*, pages 611–641. 2022. Page 3.

[AP13] J. Alperin-Sheriff and C. Peikert. Practical bootstrapping in quasilinear time. In *CRYPTO*, pages 1–20. 2013. Pages 6, 17, 18, 26, and 30.

[ARS+15] M. R. Albrecht, C. Rechberger, T. Schneider, T. Tiessen, and M. Zohner. Ciphers for MPC and FHE. In *EUROCRYPT*, pages 430–454. 2015. Page 2.

[BGV12] Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. *TOCT*, 6(3):13, 2014. Preliminary version in ITCS 2012. Pages 1, 2, 8, and 15.

[Bra12] Z. Brakerski. Fully homomorphic encryption without modulus switching from classical GapSVP. In *CRYPTO*, pages 868–886. 2012. Pages 1, 2, 8, 15, and 16.

[Bre97] T. Breuer. Integral bases for subfields of cyclotomic fields. *Appl. Algebra Eng. Commun. Comput.*, 8(4):279–289, 1997. Page 5.

[BV11a] Z. Brakerski and V. Vaikuntanathan. Fully homomorphic encryption from Ring-LWE and security for key dependent messages. In *CRYPTO*, pages 505–524. 2011. Pages 1 and 15.

[BV11b] Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. *SIAM J. Comput.*, 43(2):831–871, 2014. Preliminary version in FOCS 2011. Pages 1 and 15.

[CCS19] H. Chen, I. Chillotti, and Y. Song. Improved bootstrapping for approximate homomorphic encryption. In *EUROCRYPT*, pages 34–54. 2019. Pages 7 and 27.

[CH18] H. Chen and K. Han. Homomorphic lower digits removal and improved FHE bootstrapping. In *EUROCRYPT*, pages 315–337. 2018. Page 6.

[CIV16] W. Castryck, I. Iliashenko, and F. Vercauteren. Provably weak instances of Ring-LWE revisited. In *EUROCRYPT*, pages 147–167. 2016. Page 4.

[CKKS17] J. H. Cheon, A. Kim, M. Kim, and Y. S. Song. Homomorphic encryption for arithmetic of approximate numbers. In *ASIACRYPT*, pages 409–437. 2017. Pages 2, 8, 15, and 27.

[CLPX18] H. Chen, K. Laine, R. Player, and Y. Xia. High-precision arithmetic in homomorphic encryption. In *CT-RSA*, pages 116–136. 2018. Page 15.

[Con] K. Conrad. Prime-power units and finite subgroups of $\mathrm{GL}_n(\mathbb{Q})$. Available at https://kconrad.math.uconn.edu/blurbs/gradnumthy/primepowerunitsandGLnQ.pdf, last accessed 5 April 2023. Page 36.

[CPS18] E. Crockett, C. Peikert, and C. Sharp. ALCHEMY: A language and compiler for homomorphic encryption made easY. In *ACM CCS*, pages 1020–1037. 2018. Pages 2 and 6.

[DJL+24] A. Deo, M. Joye, B. Libert, B. R. Curtis, and M. de Bellabre. Homomorphic evaluation of LWR-based PRFs and application to transciphering. Cryptology ePrint Archive, Paper 2024/665, 2024. Page 2.

[DP16] L. Ducas and T. Prest. Fast Fourier orthogonalization. In *ISSAC*, pages 191–198. 2016. Page 32.

[ELOS15] Y. Elias, K. E. Lauter, E. Ozman, and K. E. Stange. Provably weak instances of Ring-LWE. In *CRYPTO*, pages 63–92. 2015. Page 4.

[FV12] J. Fan and F. Vercauteren. Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144, 2012. https://eprint.iacr.org/2012/144. Pages 2, 8, 15, and 16.

[Gen09a] C. Gentry. *A fully homomorphic encryption scheme*. Ph.D. thesis, Stanford University, 2009. http://crypto.stanford.edu/craig. Pages 1 and 2.

[Gen09b] C. Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178. 2009. Page 1.

[GHPS12] C. Gentry, S. Halevi, C. Peikert, and N. P. Smart. Field switching in BGV-style homomorphic encryption. *Journal of Computer Security*, 21(5):663–684, 2013. Preliminary version in SCN 2012. Pages 2 and 6.

[GHS12a] C. Gentry, S. Halevi, and N. P. Smart. Better bootstrapping in fully homomorphic encryption. In *Public Key Cryptography*, pages 1–16. 2012. Pages 1, 6, 7, 17, 18, and 30.

[GHS12b] C. Gentry, S. Halevi, and N. P. Smart. Fully homomorphic encryption with polylog overhead. In *EUROCRYPT*, pages 465–482. 2012. Pages 1, 2, and 16.

[GHS12c] C. Gentry, S. Halevi, and N. P. Smart. Homomorphic evaluation of the AES circuit. In *CRYPTO*, pages 850–867. 2012. Updated version at https://eprint.iacr.org/2012/099. Pages 1, 2, 3, 4, 8, 39, 40, and 41.

[GIKV23] R. Geelen, I. Iliashenko, J. Kang, and F. Vercauteren. On polynomial functions modulo p^e and faster bootstrapping for homomorphic encryption. In *EUROCRYPT*, pages 257–286. 2023. Page 6.

[GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206. 2008. Page 32.

[GSW13] C. Gentry, A. Sahai, and B. Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *CRYPTO*, pages 75–92. 2013. Pages 1 and 2.

[GV25] R. Geelen and F. Vercauteren. Fully homomorphic encryption for cyclotomic prime moduli. In *EUROCRYPT*, pages 366–397. 2025. Page 15.

[HS14] S. Halevi and V. Shoup. Algorithms in HElib. In *CRYPTO*, pages 554–571. 2014. Pages 7 and 26.

[HS15] S. Halevi and V. Shoup. Bootstrapping for HElib. In *EUROCRYPT*, pages 641–670. 2015. Pages 6 and 7.

[HS18] S. Halevi and V. Shoup. Faster homomorphic linear transformations in HElib. In *CRYPTO*, pages 93–120. 2018. Pages 7 and 26.

[KS18] D. Kim and Y. Song. Approximate homomorphic encryption over the conjugate-invariant ring. In *ICISC 2018*, pages 85–102. 2018. Page 3.

[LMPR08] V. Lyubashevsky, D. Micciancio, C. Peikert, and A. Rosen. SWIFFT: A modest proposal for FFT hashing. In *FSE*, pages 54–72. 2008. Pages 2 and 6.

[LPR10] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *Journal of the ACM*, 60(6):43:1–43:35, November 2013. Preliminary version in Eurocrypt 2010. Pages 2, 4, 15, 16, and 30.

[LPR13] V. Lyubashevsky, C. Peikert, and O. Regev. A toolkit for ring-LWE cryptography. In *EUROCRYPT*, pages 35–54. 2013. Pages 2, 4, 5, 6, 8, 15, 17, and 30.

[Mat89] H. Matsumura. *Commutative ring theory*, volume 8 (2nd ed.) of *Cambridge Studied in Advanced Mathematics*. Cambridge University Press, 1989. Page 11.

[Pei16] C. Peikert. How (not) to instantiate Ring-LWE. In *SCN*, pages 411–430. 2016. Page 4.

[PP25] C. Peikert and Z. Pepin. Vive Galois! Part 2: Short structured bases and fast transforms, 2025. In preparation. Pages 4, 5, 6, 7, 8, 14, 15, 30, 31, and 32.

[PRS17] C. Peikert, O. Regev, and N. Stephens-Davidowitz. Pseudorandomness of Ring-LWE for any ring and modulus. In *STOC*, pages 461–473. 2017. Pages 2 and 4.

[SV11] N. P. Smart and F. Vercauteren. Fully homomorphic SIMD operations. *Designs, Codes and Cryptography*, 71(1):57–81, 2014. Preliminary version in ePrint Report 2011/133. Pages 1 and 2.

A Slot Type and Number in Abelian Number Fields

In this section we characterize how primes split in abelian number fields of prime-power conductor and their composites. These tools can be used to identify abelian number fields that have a desired type and number of "SIMD slots." Figure 1 provides several numerical examples.

A.1 Cyclotomic Fields

An important family of abelian number fields is the *cyclotomic* fields. For a positive integer m called the *conductor*, the mth cyclotomic field $M = \mathbb{Q}(\zeta_m)$ is obtained by adjoining a primitive mth root of unity ζ_m to the rationals. Its degree over $\mathbb{Q}$ is $\varphi(m)$, the totient of m. Its automorphisms τ_i are defined by $\tau_i(\zeta_m) = \zeta_m^i$ for each $i \in \mathbb{Z}_m^*$, the multiplicative group of integers modulo m. Therefore, its Galois group $\mathrm{Gal}(M/\mathbb{Q})$ is isomorphic to $\mathbb{Z}_m^*$, so it is an abelian extension. It is a standard fact (see, e.g., [Con, Theorem 2.3]) that $\mathbb{Z}_{p^k}^*$ is cyclic for odd prime p and $k \geq 1$, and $\mathbb{Z}_{2^k} \cong \{\pm 1\} \times \{i : i = 1 \pmod 5\} \cong \langle -1 \rangle \times \langle 5 \rangle$ for $k \geq 2$.

For a prime integer r that is coprime with m, the common decomposition group of the prime ideals $\mathfrak{r}$ lying over r in $\mathcal{O}_M$ is $\langle \tau_r \rangle \subseteq \mathrm{Gal}(M/\mathbb{Q})$, the cyclic subgroup generated by τ_r, which is isomorphic to $\langle r \rangle \subseteq \mathbb{Z}_m^*$. The inertia group is trivial, i.e., r is unramified in M.

A.2 Slot Structure

The following lemma shows that there is a maximum number of "slots" of a given prime characteristic r that can be obtained in power-of-p cyclotomics (and, by implication, their subfields), and tells us where that maximum is obtained. See Figure 1 for some numerical examples.

Lemma A.1. *Let p, r be distinct prime integers, $\tilde{p} = 4$ if $p = 2$ and $\tilde{p} = p$ otherwise, d be the multiplicative order of r modulo $\tilde{p}$, and k be the greatest integer such that p^k divides $r^d - 1$ (so $\tilde{p} \mid p^k$). In the p^ℓth cyclotomic ring for $p^\ell \geq \tilde{p}$, the prime r splits as the product of g prime ideals each having residue field $\mathbb{F}_{r^f}$, where $f = d \cdot p^{\max(0,\ell-k)}$ and $fg = \varphi(p^\ell)$. In particular, $g = \varphi(p^k)/d$ for all $\ell \geq k$.*

Proof. By prime splitting in abelian extensions (see Section 2.2.3), and because the decomposition and inertia groups of the prime ideals lying over r in the p^ℓth cyclotomic respectively correspond to $\langle r \rangle \subseteq \mathbb{Z}_m^*$ and the trivial group, it suffices to show that f is the multiplicative order of r modulo p^ℓ. We proceed by cases. For $\ell \leq k$, this order is d, by definition of d and k, and because $\tilde{p} \mid p^\ell$. For $\ell > k$, by definition of k, we have that $r^* = r^d \bmod p^\ell$ is an element of the order-$p^{\ell-k}$ subgroup $S = \{i : i = 1 \pmod{p^k}\} \subseteq \mathbb{Z}_{p^\ell}^*$, which is cyclic (because it is a subgroup of a cyclic group, either $\mathbb{Z}_{p^\ell}^*$ if p is odd or $\{i : i = 1 \pmod 5\} \cong \langle 5 \rangle$). Moreover, r^* is not an element of the maximal proper subgroup $\{i : i = 1 \pmod{p^{k+1}}\} \subseteq S$, so it has order $p^{\ell-k}$. □

For an abelian number-field extension L/K and an unramified prime ideal of $\mathcal{O}_K$, the next lemma relates its number of prime factors in $\mathcal{O}_L$ and their common residue degree (as described in Section 2.2.3) to those in any subextension of L/K. As corollary, by specializing to $K = \mathbb{Q}$ and combining with Lemma A.1, for a prime characteristic r we can obtain a residue field of $\mathbb{F}_{r^{f'}}$ from a subfield of a power-of-p cyclotomic for prime $p \neq r$ if and only if $f' \mid dp^i$ for some integer $i \geq 0$, where d is as defined in Lemma A.1.

Lemma A.2. *Let:*

- *L/K be an abelian extension of number fields;*
- *$\mathfrak{r}$ be a prime ideal of $\mathcal{O}_K$ that does not ramify in L, having residue field $\mathcal{O}_K/\mathfrak{r} \cong \mathbb{F}_r$ for some prime-power r;*
- *D be the order-f decomposition group of the $g = \deg(L/K)/f$ prime $\mathcal{O}_L$-ideals lying over $\mathfrak{r}$ (which each have residue field $\mathbb{F}_{r^f}$); and*
- *H be a subgroup of $\mathrm{Gal}(L/K)$.*

Then $\mathfrak{r}$ splits in L^H as the product of $g' = g \cdot |H \cap D|/|H|$ prime ideals each having residue field $\mathbb{F}_{r^{f'}}$, where $f' = f/|H \cap D|$. In particular, $H \subseteq D$ if and only if $f' = f/|H|$ and $g' = g$, and $H \cap D$ is trivial if and only if $f' = f$ and $g' = g/|H|$.

Proof. We have that

$$D/(H \cap D) \cong \mathrm{Gal}(L^{H \cap D}/L^D) \cong \mathrm{Gal}(L^H/L^{HD}) \,,$$

where the first isomorphism is by restriction to $L^{H \cap D}$, and the second is by further restriction to L^H and Lemma 2.2, because $L^{H \cap D} = L^H L^D$ and $L^{HD} = L^H \cap L^D$ by the Galois correspondence. Next, we claim that $\mathrm{Gal}(L^H/L^{HD}) = D'$, the decomposition subgroup of the prime ideals lying over $\mathfrak{r}$ in L^H. This is because by maximality, $L^{HD} = L^H \cap L^D$ is the *largest* subfield of L^H in which $\mathfrak{r}$ splits completely, hence L^{HD} is indeed the fixed field of D' in L^H. It therefore follows that $\mathfrak{r}$ splits in L^H as the product of g' prime ideals each having residue field $\mathbb{F}_{r^{f'}}$, where $f' = |D'| = |D/(H \cap D)| = f/|H \cap D|$, and $g' = \deg(L^H/K)/f' = fg/(f'|H|) = g \cdot |H \cap D|/|H|$. □

Remark A.3. We discuss several useful implications of Lemma A.2. First note that any intermediate field L^H of L/K (for a subgroup $H \subseteq \mathrm{Gal}(L/K)$) can be reached by applying the two particular cases from Lemma A.2 in sequence. First, go to $L^{H'}$ for $H' = H \cap D \subseteq D$, which reduces just residue degree from f to $f' = f/|H'|$. Then, go from $L^{H'}$ to L^H using the restriction of H to $L^{H'}$ (which is a subgroup of $\mathrm{Gal}(L^{H'}/K)$ isomorphic to H/H'), which preserves the residue degree and hence reduces just the number of prime ideal factors from g to $g' = g/|H/H'| = g \cdot |H'|/|H|$.

Moreover, for *any* $f' \mid f$ and for $g' = g$ there exists a subgroup $H \subseteq D$ yielding these parameters, namely, any order-(f/f') subgroup of D. Such a subgroup exists because D is finite and abelian, and thus is a "converse of Lagrange's Theorem" group. Moreover, if D is cyclic—in particular, when L is a cyclotomic and $K = \mathbb{Q}$—then there is a *unique* such H. Finally, if $\mathrm{Gal}(L/K)$ is cyclic—e.g., when L is an odd prime-power cyclotomic—then every distinct subgroup H yields a distinct product $f'g' = \deg(L/K)/|H|$ and hence distinct pair (f', g'), because distinct subgroups have distinct orders.

Unfortunately, for a desired $f' \mid f$ it is not always possible to obtain an arbitrary divisor g' of g. For example, letting $L/K = \mathbb{Q}(\zeta_{17})/\mathbb{Q}$ and $\mathfrak{r} = 2\mathbb{Z}$, we have that $f = 8$ and $g = 2$. However, the only subgroup of $\mathbb{Z}_{17}^*$ having trivial intersection with the decomposition group $D = \langle 2 \rangle$ is the trivial subgroup itself, and so there is no subfield in which $f' = f = 8$ and $g' = 1$.

Finally, the next lemma makes it simple to construct an abelian number field having many slots of the desired type: simply take the *composite* of abelian number fields having coprime conductors (e.g., powers of distinct primes) and slot types whose composite is the desired slot type. The composite of $\mathbb{F}_{r^k}$ and $\mathbb{F}_{r^\ell}$ is $\mathbb{F}_{r^{\mathrm{lcm}(k,\ell)}}$, so we want to use abelian number fields whose residue degrees for the primes lying over r have the desired least common multiple. Also recall that Theorem 2.3 gives us short, structured bases for abelian number fields constructed in this way.

For subrings S_1, S_2 of some ring R, their *composite* ring is defined (just like the definition of composite field) as the subring $S_1S_2 = \{\sum_{i=1}^r \alpha_i\beta_i : \alpha_i \in S_1, \beta_i \in S_2, \text{finite } r\} \subseteq R$. Note that by Lemma 2.5, the hypothesis $\mathcal{O}_{L_1L_2} = \mathcal{O}_{L_1}\mathcal{O}_{L_2}$ in the following lemma is satisfied when L_1 and L_2 are abelian number fields with coprime conductors, which is the primary way in which we use the result.

Lemma A.4. *Let L_1, L_2 be number fields with composite field $M = L_1L_2$ such that $\mathcal{O}_M = \mathcal{O}_{L_1}\mathcal{O}_{L_2}$, let $\mathfrak{r}_M$ be a prime ideal in $\mathcal{O}_M$, and let $\mathfrak{r}_i := \mathcal{O}_{L_i} \cap \mathfrak{r}_M$ for $i \in \{1, 2\}$. Then $\mathcal{O}_M/\mathfrak{r}_M = \phi_1(\mathcal{O}_{L_1}/\mathfrak{r}_{L_1})\phi_2(\mathcal{O}_{L_2}/\mathfrak{r}_{L_2})$, where the natural map $\phi_i \colon \mathcal{O}_{L_i}/\mathfrak{r}_{L_i} \to \mathcal{O}_M/\mathfrak{r}_M$ is a field embedding.*

Proof. First, the natural map ϕ_i is clearly a field homomorphism. Furthermore, it is injective, since $\phi_i(x + \mathfrak{r}_{L_i}) = \mathfrak{r}_M$ for some $x_i \in \mathcal{O}_{L_i}$ if and only if $x \in \mathcal{O}_{L_i} \cap \mathfrak{r}_M = \mathfrak{r}_{L_i}$. So, $\mathcal{O}_M/\mathfrak{r}_M$ contains each $\phi_i(\mathcal{O}_{L_i}/\mathfrak{r}_{L_i})$, and thus contains their composite as well. Finally, for the reverse inclusion, since $\mathcal{O}_M = \mathcal{O}_{L_1}\mathcal{O}_{L_2}$, any $x \in \mathcal{O}_M$ can be written as a finite sum $x = \sum_i \alpha_i\beta_i$ for some $\alpha_i \in \mathcal{O}_{L_1}$, $\beta_i \in \mathcal{O}_{L_2}$, so

$$x + \mathfrak{r}_M = \sum_i \alpha_i\beta_i + \mathfrak{r}_M = \sum_i \phi_1(\alpha_i + \mathfrak{r}_{L_1})\phi_2(\beta_i + \mathfrak{r}_{L_2}) + \mathfrak{r}_M \ ,$$

hence $\mathcal{O}_M/\mathfrak{r}_M$ is contained in the composite of the fields $\phi_i(\mathcal{O}_{L_i}/\mathfrak{r}_{L_i})$, as needed. □

Summary. The above lemmas directly yield a procedure for constructing abelian number fields with any desired characteristic-r (finite field) slot type. First, Lemma A.1 characterizes what slot types $\mathbb{F}_{r^f}$ and number of slots can be obtained in power-of-p cyclotomics, for any prime $p \neq r$. However, the obtainable residue degrees f are often larger than desired. Next, Lemma A.2 and Remark A.3 show how to obtain any slot type of residue degree $f' \mid f$, i.e., any subfield $\mathbb{F}_{r^{f'}}$ of the cyclotomic's residue field $\mathbb{F}_{r^f}$, by using a suitable cyclotomic subfield (namely, the fixed subfield of a subgroup of appropriate order of the decomposition group). Finally, Lemma A.4 shows how to obtain more slots by compositing such power-of-p cyclotomic subfields for distinct primes p. The resulting slot type is the composite of the slot types for the component cyclotomic subfields, so these should each have a slot type that is a subfield of the ultimate desired slot type. Another consequence of this is that the number of slots obtained in this way is super-multiplicative.

A.3 Numerical Examples

Figure 1 gives some selected examples of parameters that can be obtained from Lemma A.1, which can be reduced and refined using Lemma A.2.

For the class of example where $r = p^k \cdot c + 1$, notice that given r and p, we can confirm the given values of d and k using the definitions in Lemma A.1. Indeed, $d = 1$ since $r = 1 \pmod{\tilde{p}}$, and p^k is the largest power of p that divides $r^1 - 1 = p^k \cdot m$ since $p \nmid m$. Similarly for $r = p^k \cdot m - 1$, we have that $r^1 = -1 \pmod{\tilde{p}}$, so $r^2 = 1 \pmod{\tilde{p}}$, and p^k is the largest power of p that divides $r^2 - 1 = p^{2k} \cdot m^2 - 2p^k \cdot m = p^k \cdot m \cdot (p^k \cdot m - 2)$, since $p \nmid m$ and $p \nmid (p^k \cdot m - 2)$.

Now consider the examples where $r = 2$. The case $p = 17$ represents the least p where Lemmas A.1 and A.2 yield an abelian number field where the prime-ideal factors of r have residue field $\mathbb{F}_{2^8}$; there are two such factors. The values $p = 31$, $p = 73$, and $p = 127$ are the least p where g is at least 4, 8, and 16, respectively. So, r splits into at least 8, 16, and 32 prime-ideal factors having common residue field $\mathbb{F}_{2^8}$ in a suitable subfield of the mth cyclotomic for $m = 17 \cdot 31$, $m = 17 \cdot 73$, and $m = 17 \cdot 127$, respectively. By contrast, the values $p = 241$, $p = 257$, and $p = 5153$ are the least prime p where g is at least 8, 16, and 32 (respectively) for a residue field of $\mathbb{F}_{2^8}$.

Also for $r = 2$, it is also worth comparing what can be obtained in cyclotomic *subfields*, versus in cyclotomics only. In any cyclotomic in which 2 does not ramify, every prime ideal lying over 2 has residue degree $f > 1$, i.e., it is *not possible* to obtain $\mathbb{F}_2$ as a residue field, only proper extensions of it. By contrast, in cyclotomic subfields—specifically, (subfields of) the decomposition subfield of the primes lying above 2—we can obtain a number of $\mathbb{F}_2$-slots matching the degree of the subfield. For example, we get two $\mathbb{F}_2$-slots in the decomposition subfield of $\mathbb{Q}(\zeta_{17})$; six in the decomposition subfield of $\mathbb{Q}(\zeta_{31})$; eight in the decomposition subfield of $\mathbb{Q}(\zeta_{73})$, etc.

The examples for $r = 3$, $r = 5$, and $r = 7$ give the $p \leq 256$ that yield the largest g for the given r. The examples where $r = 263, r = 443$ (respectively, $r = 79193$) represent the $r \in [2^8, 2^9]$ (respectively, $r \in [2^{16}, 2^{17}]$) that yield the largest g for the given p.

r	p	d	k	$g =$ # slots
$r = p^k \cdot c + 1$	any	1	$\geq \log_p(\tilde{p})$	$= \varphi(p^k)$
$r = p^k \cdot c - 1$	$\neq 2$	2	any	$= \varphi(p^k)/2$
2	17	8	1	$2 = 16 \cdot 17^0/8$
2	31	5	1	$6 = 30 \cdot 31^0/5$
2	73	9	1	$8 = 72 \cdot 73^0$
2	127	7	1	$18 = 126 \cdot 127^0/7$
2	241	24	1	$10 = 240 \cdot 241^0/24$
2	257	16	1	$16 = 256 \cdot 257^0/16$
2	5153	112	1	$46 = 5152 \cdot 5153^0/112$
3	11	5	2	$22 = 10 \cdot 11^1/5$
5	71	5	1	$14 = 70 \cdot 71^0/5$
7	191	10	1	$19 = 190 \cdot 191^0/10$
263	7	3	2	$14 = 6 \cdot 7^1/3$
443	5	4	4	$125 = 4 \cdot 5^3/4$
79193	5	4	6	$3125 = 4 \cdot 5^5/4$

Figure 1: Some examples of the maximum number of slots that can be obtained for a prime modulus $r \neq p$ in a power-of-p cyclotomic (sub)field. The values d, k, g are as in Lemma A.1, i.e., modulus r in the p^kth cyclotomic yields g slots of type $\mathbb{F}_r^d$, and higher powers of p do not yield any additional slots. In the first two generic examples, p and k may be arbitrary (subject to the minor listed restrictions), and r has the given form (for integer c) where $p \nmid r$.

B Example Instantiations for Homomorphic AES Evaluation

In this section, we expand upon the example application of homomorphic AES evaluation, and describe some new instantiations using the tools from this work. For this application, the most natural plaintext "SIMD slot" type is the finite field $\mathbb{F}_{2^8}$, because the AES function works with vectors over this field. However, because [GHS12c] used *cyclotomic* rings, it was induced to use $\mathbb{F}_{2^{24}}$ as its slot type, which is "wasteful" by a factor of three. More specifically, [GHS12c] uses the cyclotomic field $\mathbb{Q}(\zeta_{28679}) = \mathbb{Q}(\zeta_7)\mathbb{Q}(\zeta_{17})\mathbb{Q}(\zeta_{241})$, which has dimension $\varphi(28679) = 6 \cdot 16 \cdot 240 = 23040$, but only $23040/24 = 960$ $\mathbb{F}_{2^{24}}$-slots, enough for 60 AES blocks.

B.1 Example Instantiation 1

Using the tools from this paper, we can see that one good choice of field is

$$L = C_{17}^{(1)} C_{127}^{(7)} C_{241}^{(3)} ,$$

where $C_p^{(d)}$ is the unique subfield of $\mathbb{Q}(\zeta_p)$ such that $\deg(\mathbb{Q}(\zeta_p)/C_p^{(d)}) = d$, for odd prime p and $d \mid (p-1)$. By Lemmas A.1 and A.2,

- $C_{17}^{(1)}$ has 2 $\mathbb{F}_{2^8}$-slots,
- $C_{127}^{(7)}$ has 18 $\mathbb{F}_2$-slots, and
- $C_{241}^{(3)}$ has 10 $\mathbb{F}_{2^8}$-slots.

In summary, L has:

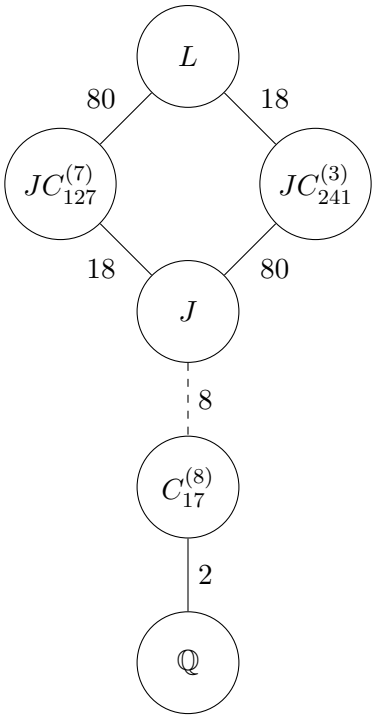

Figure 2: A diagram representing the structure of the various abelian number fields in our first example instantiation, where $J = C_{17}^{(1)}$. Each edge represents a field extension of the given dimension, with solid and dashed edges respectively representing complete splitting and total inertness for the prime ideals lying over 2. For any pair of depicted distinct field extensions L_1/K and L_2/K over a common base field K, we have that $L_1 \cap L_2 = K$. So by Remark 5.12, for any prime ideal $\mathfrak{r}$ of $\mathcal{O}_K$ lying over 2, the mod-$\mathfrak{r}$ CRT basis of $\mathcal{O}_{L_1 L_2}/\mathcal{O}_K$ is the Kronecker product of the mod-$\mathfrak{r}$ CRT bases of $\mathcal{O}_{L_1}/\mathcal{O}_K$ and $\mathcal{O}_{L_2}/\mathcal{O}_K$.

- dimension $16 \cdot 18 \cdot 80 = 23040$ over $\mathbb{Q}$ (cf. 23040 in [GHS12c]), and
- $23040/8 = 2880$ $\mathbb{F}_{2^8}$-slots (cf. 960 $\mathbb{F}_{2^{24}}$-slots) by Lemma A.4, enough for 180 AES blocks (cf. 60).

B.1.1 Tensor of CRT Slots

First, note that $J = C_{17}^{(1)}$ has two $\mathbb{F}_{2^8}$-slots. Specifically, by the facts recalled at the start of Appendix A, the prime $2 \in \mathbb{Z}$ splits completely over the quadratic extension $C_{17}^{(8)}/\mathbb{Q}$, and the two prime ideals lying over 2 are totally inert in the extension $J/C_{17}^{(8)}$. Furthermore, by Lemma A.4, J, $JC_{127}^{(7)}$, and $JC_{241}^{(3)}$ all have $\mathbb{F}_{2^8}$-slots, and so the two prime ideals lying over 2 in J split completely in the extensions $JC_{127}^{(7)}/J$ and $JC_{241}^{(3)}/J$. See Figure 2 for a diagram of these extensions and how they relate.

Observe that the composite of $JC_{127}^{(7)}$ and $JC_{241}^{(3)}$ is L, and their intersections is J. Therefore, by Remark 5.12, letting $\mathfrak{r} = 2\mathcal{O}_J$, the mod-$\mathfrak{r}$ CRT basis of $\mathcal{O}_L/\mathcal{O}_J$ is the Kronecker product of the mod-$\mathfrak{r}$ CRT bases of $\mathcal{O}_{JC_{127}^{(7)}}/\mathcal{O}_J$ and $\mathcal{O}_{JC_{241}^{(3)}}/\mathcal{O}_J$, so the mod-$\mathfrak{r}$ CRT basis of $\mathcal{O}_L/\mathcal{O}_J$ has an order-two tensor structure over $\mathcal{O}_J/2\mathcal{O}_J$, where the automorphisms of these extensions permute along the corresponding dimension.

Next, we obtain a $\mathbb{Z}_2$-basis of $\mathcal{O}_J/2\mathcal{O}_J$ as the Kronecker product of the mod-2 CRT basis of $\mathcal{O}_{C_{17}^{(8)}}/\mathbb{Z}$ and any $\mathcal{O}_{C_{17}^{(8)}}$-basis of $\mathcal{O}_J$, such as $\vec{d}_{17} = (\zeta_{17}^i)_{i=0}^7$. Finally, we obtain a $\mathbb{Z}_2$-basis of $\mathcal{O}_L/2\mathcal{O}_L$ as the Kronecker

product of the mod-$\mathfrak{r}$ CRT basis of $\mathcal{O}_L/\mathcal{O}_J$ with our $\mathbb{Z}_2$-basis of $\mathcal{O}_J/2\mathcal{O}_J$. Such a Kronecker product basis may be used to represent an element of $\mathcal{O}_L/2\mathcal{O}_L$ as an order-3 tensor over $\mathbb{Z}_2$, where the $\mathcal{O}_{C_{17}^{(8)}}$-basis of $\mathcal{O}_J$ corresponds to a $\mathbb{Z}_2$-basis of the $\mathbb{F}_{2^8}$-slot.

A key observation of this tensor view is that the automorphisms of J, $C_{127}^{(7)}$, and $C_{241}^{(3)}$ act independently on their respective dimensions of the tensor. Indeed, for the latter two components they simply *permute* the tensor, acting regularly. Furthermore, for both of the prime ideals $\mathfrak{r}$ lying over 2 in $\mathcal{O}_{C_{17}^{(8)}}$, the automorphisms of $\mathrm{Gal}(J/C_{17}^{(8)})$ induce (Frobenius) automorphisms of $\mathcal{O}_J/\mathfrak{r}\mathcal{O}_J$. Therefore, we have efficient homomorphic evaluation of arbitrary automorphisms of $\mathbb{F}_{2^8}$ in a SIMD fashion across all the slots via, the lifting of the corresponding automorphisms of $\mathrm{Gal}(J/C_{17}^{(8)})$ to $\mathrm{Gal}(L/C_{17}^{(8)}) \subseteq \mathrm{Gal}(L/\mathbb{Q})$.

B.1.2 Short, Structured Basis

Using Theorem 2.3, we can efficiently compute short, structured integral bases of $C_{17}^{(8)}$, $C_{127}^{(7)}$, and $C_{241}^{(3)}$. Let $\vec{b}_{m,d}$ be the integral basis of $C_m^{(d)}$ obtained from this theorem. Because these m are prime, upper bounds on the canonical norms of these bases (in their respective number fields) are $\|\vec{b}_{m,d}\|^2 \le d \cdot \deg(C_m^{(d)}/\mathbb{Q})$.

Additionally, $\vec{p}_{17} = (\zeta_{17}^i)_{i=0}^7$ is a power $\mathcal{O}_{C_{17}^{(8)}}$-basis of $\mathcal{O}_J$, because $\mathcal{O}_J = \mathbb{Z}[\zeta_{17}]$ and $\mathbb{Z} \subseteq \mathcal{O}_{C_{17}^{(8)}} \subseteq \mathcal{O}_J$. Thus, $\vec{p}_{17} \otimes \vec{b}_{17,8}$ is a structured $\mathbb{Z}$-basis of $\mathcal{O}_J$ with canonical norm bounded by $\|\vec{p}_{17} \otimes \vec{b}_{17,8}\|^2 = \deg(J/C_{17}^{(8)}) \cdot \|\vec{b}_{17,8}\|^2 \le 8 \cdot \deg(J/\mathbb{Q})$, where the equality holds because $\vec{p}_{17}$ consists of roots of unity. Then by Lemma 2.5,

$$(\vec{p}_{17} \otimes \vec{b}_{17,8}) \otimes \vec{b}_{127,7} \otimes \vec{b}_{241,3}$$

is a $\mathbb{Z}$-basis of $\mathcal{O}_L$ with norm upper bounded by

$$\sqrt{8 \cdot 7 \cdot 3} \cdot \sqrt{\deg(L/\mathbb{Q})} \approx 12.96148 \cdot \sqrt{\deg(L/\mathbb{Q})}\ .$$

For comparison, recall that *any* nonzero element of $\mathcal{O}_K$ for *any* number field K has norm at least $\sqrt{\deg(K/\mathbb{Q})}$, so this integral basis of L has norm within a factor of 13 of optimal for any number field of the same degree.

B.2 Example Instantiation 2

We can also use our tools to find an abelian number field of similar dimension to that of the previous example, while keeping all the prime divisors of the conductor small. This yields a much finer-grained (higher-order) tensor of CRT slots, which supports cheaper and richer homomorphic linear algebra. One example of such a field is

$$L = C_3^{(1)} C_5^{(1)} C_7^{(3)} C_{11}^{(5)} C_{13}^{(3)} C_{17}^{(1)} C_{19}^{(9)} C_{31}^{(5)}\ ,$$

where $C_p^{(d)}$ is the unique subfield of $\mathbb{Q}(\zeta_p)$ such that $\deg(\mathbb{Q}(\zeta_p)/C_p^{(d)}) = d$, for odd prime p and $d \mid (p-1)$. By Lemmas A.1 and A.2,

- $C_3^{(1)}$ has 1 $\mathbb{F}_{2^2}$-slot,
- $C_5^{(1)}$ has 1 $\mathbb{F}_{2^4}$-slot,
- $C_7^{(3)}$ has 2 $\mathbb{F}_2$-slots,
- $C_{11}^{(5)}$ has 1 $\mathbb{F}_{2^2}$-slot,
- $C_{13}^{(3)}$ has 1 $\mathbb{F}_{2^4}$-slot,
- $C_{17}^{(1)}$ has 2 $\mathbb{F}_{2^8}$-slots,
- $C_{19}^{(9)}$ has 1 $\mathbb{F}_{2^2}$-slot, and
- $C_{31}^{(5)}$ has 6 $\mathbb{F}_2$-slots.

In summary, L has:

- dimension $2 \cdot 4 \cdot 2 \cdot 2 \cdot 4 \cdot 16 \cdot 2 \cdot 6 = 24576$ over $\mathbb{Q}$ (cf. 23040 in [GHS12c]), and
- $24576/8 = 3072$ $\mathbb{F}_{2^8}$-slots (cf. 960 $\mathbb{F}_{2^{24}}$-slots) by Lemma A.4, enough for 192 AES blocks (cf. 60).

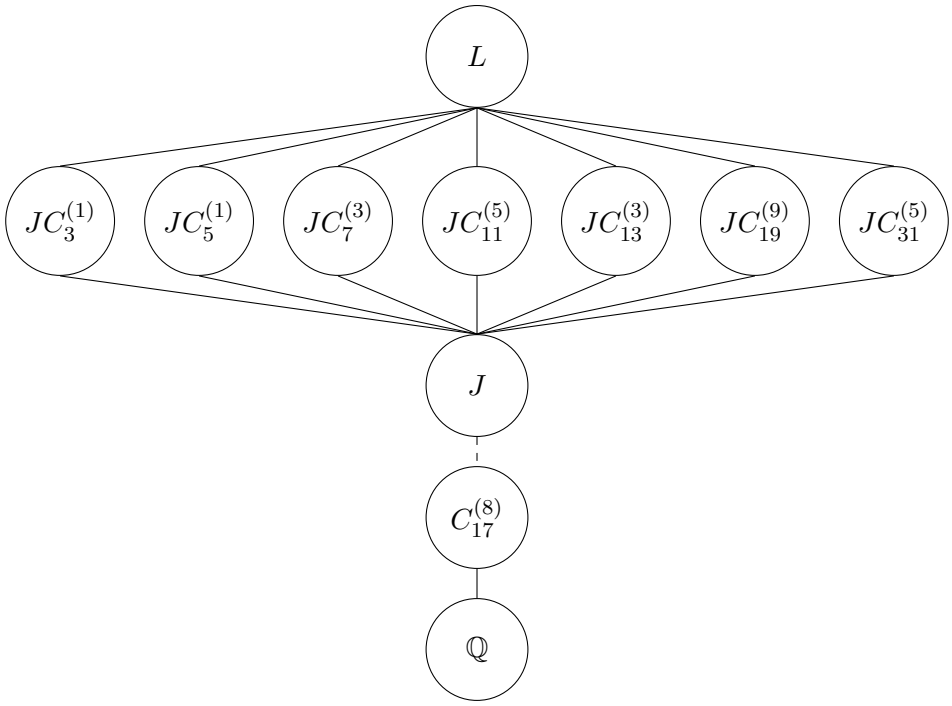

Figure 3: A diagram representing the structure of the various abelian number fields in our example construction, where $J = C_{17}^{(1)}$. Each edge represents a field extension, with solid and dashed edges respectively representing complete splitting and total inertness for the prime ideals lying over 2. For any pair of depicted distinct field extensions L_1/K and L_2/K over a common base field K, we have that $L_1 \cap L_2 = K$. So by Remark 5.12, for any prime ideal $\mathfrak{r}$ of $\mathcal{O}_K$ lying over 2, the mod-$\mathfrak{r}$ CRT basis of $\mathcal{O}_{L_1 L_2}/\mathcal{O}_K$ is the Kronecker product of the mod-$\mathfrak{r}$ CRT bases of $\mathcal{O}_{L_1}/\mathcal{O}_K$ and $\mathcal{O}_{L_2}/\mathcal{O}_K$.

B.2.1 Tensor of CRT Slots

We construct a $\mathbb{Z}_2$-basis of $\mathcal{O}_L/2\mathcal{O}_L$ in a similar fashion as in Appendix B.1.1, but with a richer tensor structure. Just as before, recall that $J = C_{17}^{(1)}$ has two $\mathbb{F}_{2^8}$-slots, and we get a $\mathbb{Z}_2$-basis of $\mathcal{O}_J/2\mathcal{O}_J$ as the Kronecker product of $\vec{d}_{17}$ and the mod-2 CRT basis of $\mathcal{O}_{C_{17}^{(8)}}/\mathbb{Z}$. Then by Lemma A.4, J and JK for

$$K \in \{C_3^{(1)}, C_5^{(1)}, C_7^{(3)}, C_{11}^{(5)}, C_{13}^{(3)}, C_{19}^{(9)}, C_{31}^{(5)}\}$$

all have $\mathbb{F}_{2^8}$-slots, so the two prime ideals lying over 2 in J split completely in the extension JK/J. See Figure 3 for a diagram of these extensions and how they relate.

Next observe that the composite of all the fields JK is L, and their pairwise intersections all are J. So by Remark 5.12, letting $\mathfrak{r} = 2\mathcal{O}_J$, the mod-$\mathfrak{r}$ CRT basis of $\mathcal{O}_L/\mathcal{O}_J$ is the Kronecker product of the mod-$\mathfrak{r}$ CRT bases of $\mathcal{O}_{JK}/\mathcal{O}_J$ for each K. This results in an order-7 tensor structure over $\mathcal{O}_J/2\mathcal{O}_J$, where the automorphisms of each extension JK/J permutes along the corresponding dimension of the tensor.

Next, we construct a $\mathbb{Z}_2$-basis of $\mathcal{O}_L/2\mathcal{O}_L$ as the Kronecker product of the mod-$\mathfrak{r}$ CRT basis of $\mathcal{O}_L/\mathcal{O}_J$ with the previously described $\mathbb{Z}_2$-basis of $\mathcal{O}_J$. Such a Kronecker product basis may be used to represent an element of $\mathcal{O}_L/2\mathcal{O}_L$ as an order-8 tensor over $\mathbb{Z}_2$, where the $\mathcal{O}_{C_{17}^{(8)}}$-basis of $\mathcal{O}_J$ corresponds to a $\mathbb{Z}_2$-basis of the F_{2^8}-slot.

Similarly to Appendix B.1.1, the automorphisms of $C_3^{(1)}$, $C_5^{(1)}$, $C_7^{(3)}$, $C_{11}^{(5)}$, J, $C_{19}^{(9)}$, and $C_{31}^{(5)}$ act independently on their respective dimensions of the tensor, and we have efficient homomorphic evaluation of arbitrary (Frobenius) automorphisms of $\mathbb{F}_{2^8}$ in a SIMD fashion.

B.2.2 Short, Structured Basis

Using Theorem 2.3, we can efficiently compute short, structured integral bases of $C_3^{(1)}$, $C_5^{(1)}$, $C_7^{(3)}$, $C_{11}^{(5)}$, $C_{13}^{(3)}$, $C_{17}^{(8)}$, $C_{19}^{(9)}$, and $C_{31}^{(5)}$. Let $\vec{b}_{m,d}$ be the integral basis of $C_m^{(d)}$ obtained from this theorem. Because these m are prime, upper bounds on the canonical norms of these bases (in their respective number fields) are $\|\vec{b}_{m,d}\|^2 \leq d \cdot \deg(C_m^{(d)}/\mathbb{Q})$.

Just as in Appendix B.1.2, we get that $\vec{p}_{17} \otimes \vec{b}_{17,8}$ is a structured $\mathbb{Z}$-basis of $\mathcal{O}_J$ with norm bounded by $\|\vec{p}_{17} \otimes \vec{b}_{17,8}\|^2 \leq 8 \cdot \deg(J/\mathbb{Q})$. So, by Lemma 2.5,

$$\vec{b}_{3,1} \otimes \vec{b}_{5,1} \otimes \vec{b}_{7,3} \otimes \vec{b}_{11,5} \otimes \vec{b}_{13,3} \otimes (\vec{p} \otimes \vec{b}_{17,8}) \otimes \vec{b}_{19,9} \otimes \vec{b}_{31,5}$$

is a $\mathbb{Z}$-basis of $\mathcal{O}_L$ with canonical norm upper bounded by

$$\sqrt{1 \cdot 1 \cdot 3 \cdot 5 \cdot 3 \cdot 8 \cdot 9 \cdot 5} \cdot \sqrt{\deg(L/\mathbb{Q})} \approx 127.27922 \cdot \sqrt{\deg(L/\mathbb{Q})}\,.$$

So, this integral basis of L has norm within a factor of 128 of optimal for any number field of the same degree.