PhD-FSTM-2021-039
Faculty of Science, Technology and Medicine

DISSERTATION

Defence held on 06/07/2021 in Esch-sur-Alzette
to obtain the degree of

DOCTEUR DE L'UNIVERSITE DU LUXEMBOURG
EN MATHEMATIQUES

by

Luca Notarnicola

Born on 12 October 1992 in Luxembourg (Luxembourg)

TOPICS IN COMPUTATIONAL NUMBER THEORY AND CRYPTANALYSIS

On Euclidean Lattices, Edwards Curves and Cryptographic Multilinear Maps

Dissertation defence committee:

Dr. Gabor Wiese, Dissertation supervisor
Professor, Université du Luxembourg

Dr. Jean-Sébastien Coron, Dissertation supervisor
Professor, Université du Luxembourg

Dr. Antonella Perucca, Chairman
Professor, Université du Luxembourg

Dr. Frederik Vercautern
Professor, KU Leuven

Dr. Phong Q. Nguyen
Professor, Inria and DIENS PSL Paris

Acknowledgments

Thank you to my supervisors, Jean-Sébastien Coron and Gabor Wiese. I am deeply grateful for their continuous support, for guiding me with patience and passion, and for everything they taught me over the last years. *Thank you*, Jean-Sébastien, for introducing me to the beautiful world of cryptography; your great positivity, intuition, and exemplary simplicity in tackling problems helped me a lot. *Thank you*, Gabor, for shaping my interest in number theory since my Bachelor studies and supervising all my theses; your rigour and extraordinary attention to detail made me improve every day.

Thank you to my collaborators, Jean-Sébastien Coron, Gabor Wiese, and Samuele Anni, for sharing many ideas and for all their time and efforts. It is a pleasure to work with you. *Thank you*, Samuele, for the opportunity to work together, for pointing me to many interesting problems, and for the great enthusiasm during our meetings.

Thank you to the non-supervising jury members, Antonella Perucca, Frederik Vercauteren, and Phong Nguyen. I am sincerely thankful for your interest in my work. *Thank you*, Frederik, for joining my thesis supervising committee, and the fruitful discussions.

Thank you to my colleagues and friends, for contributing to a great working environment. I am particularly thankful to the members close to the number theory group or the Applied Crypto group, with whom I had the chance to spend my everyday life: Gabor, Antonella, Alexander, Shaunak, Samuele, Andrea, Lassina, Alexandre, Jasper, Mariagiulia, Emiliano, Daniel, Pietro, Sebastiano, Bryan, Flavio, Fabio, and Jean-Sébastien, Benoît, Moon Sung, Rajeev, Razvan, François, Vitor, Najmeh, Agnese, Lorenzo, Jim, Raluca, Paul. *Thank you* to my office mates, Jasper, Daniel, and Fabio, for the nice atmosphere. *Thank you* to Gerard van der Geer for his beautiful lectures during his regular visits at the department. *Thank you* to who preciously helped me in organizing events and trips: Guenda and Gabor, for co-organizing the Mathematical Careers Day 2018 and 2019, Robert, for your outstanding help within the DPMA-PhD representative team, and Jim and Vincent Koenig, for your valuable support in coordinating events within the SP2 (Security and Privacy for System Protection) community. *Thank you* to many other friends from various departments, for all the fun activities, trips, football games, dinners, and mostly, your friendship. *Thank you* to the secretaries Marie Leblanc, Katharina Heil and Fabienne Schmitz. I could always count on your efficiency.

Thank you to all my family members, for always believing in me. Especially, my brother, Massimo, and my parents, Rosanna and Vito. A heartfelt thank you, for everything.

This work is supported by the Luxembourg National Research Fund through the grant PRIDE15/10621687/- SPsquared.

Luca Notarnicola

Bettembourg, May 2021

Abstract

The content of this thesis is situated between number theory and cryptology. It contributes in several directions of mathematical cryptanalysis and arithmetic geometry, by the use of explicit and computational methods in number theory. The thesis is divided into four main parts, describing four different works.

The first part of this thesis treats cryptographic multilinear maps, introduced in the early 2000s, as extension of bilinear pairings on abelian varieties following a survey by Boneh and Silverberg (Contemp. Math. 2003). The difficulties of constructing efficient cryptographic multilinear maps from algebraic geometry led cryptographers to obtain multilinear maps from the notion of graded encoding systems, following the formalism of Garg, Gentry and Halevi (Eurocrypt 2013), backgrounded from ideas on fully homomorphic encryption. To this date, there are three candidate constructions. Their security however is rather poorly understood and many attacks have flourished over the past years. In this thesis, we consider the CLT13 Scheme, proposed by Coron, Lepoint and Tibouchi (Crypto 2013), constructing multilinear maps over the integers. A vulnerability of this scheme was detected by Gentry, Lewko and Waters (Crypto 2014), by exploiting the composite-ring structure of the plaintext space, enabling a simple two-dimensional lattice attack to reveal secret information. At the same time, the authors suggested a simple countermeasure to prevent this line of attack. By relying on high-dimensional lattice reduction, we design new cryptanalysis against CLT13, overriding the countermeasure of Gentry, Lewko and Waters by several orders of magnitude. Combined with the Cheon et al. attack (Eurocrypt 2015), we reveal all secret parameters of CLT13. By applying our cryptanalysis to concrete instantiations of CLT13-based constructions, certain parameter ranges are found unsecure.

The second work presented in this thesis isolates the Hidden Lattice Problem as a more general problem appearing in several cryptographic scenarios, among which the Hidden Subset Sum Problem, studied by Nguyen and Stern (Crypto 1999). The Nguyen-Stern algorithm for the Hidden Subset Sum Problem builds on the use of orthogonal lattices, which have shown powerful in public-key cryptanalysis. After extending their algorithm to our more general setting, our main contribution is to present, based on duality in lattice theory, a new algorithm for the Hidden Lattice Problem, providing a competitive alternative to the celebrated Nguyen-Stern orthogonal lattice attack. For both algorithms, we provide practical parameters, based on heuristic and rigorous studies of the geometry of the underlying lattices. The generality of our definition for the Hidden Lattice Problem allows encompassing multiple number-theoretic problems of cryptographic interest. For these problems, our new algorithm leads to a valuable alternative to the state-of-the-art algorithms.

The third work presented in this thesis aims at extending the powerful cryptanalysis by Cheon et al. (Eurocrypt 2015) against the multiparty Diffie-Hellman protocol instantiated under the aforementioned CLT13 Scheme. To do so, we formulate, independently of the CLT-framework, a general problem on simultaneous diagonalization of special matrices of low-rank, which we call "incomplete". Based on techniques from linear algebra, our major contribution is the design of efficient algorithms for this problem, providing explicit and practical parameters. We recognize that our algorithms also apply to solve the Approximate Common Divisor Problem based on the Chinese Remainder Theorem (CRT-ACD Problem). In particular, our algorithms lead to quadratic improvements in the input length for the CRT-ACD Problem, and some cases of the Cheon et al. attack.

The last work included in this thesis studies Edwards curves, a normal form for elliptic curves, first introduced by Edwards (Bull. Amer. Math. Soc. 2007). This model for elliptic curves has been proposed for elliptic-curve cryptography by the work of Bernstein and Lange (Asiacrypt 2007), by exploiting the efficient Edwards point arithmetic. Their work includes a construction of the Edwards model from the Weierstrass model. While not directly oriented towards cryptographic targets, our study of the Edwards model is of arithmetic-geometric nature. Our main contributions include abstract constructions for the Edwards model, an explicit generalization of the Bernstein-Lange construction, as well as general properties of geometric nature. From an analytic perspective, we complement our study with extensive computer calculations related to the ranks of rational elliptic curves in a family related to Edwards curves. Part of our statistical data is based on our extension of an algorithm involving L-functions and relying on the well-known Birch and Swinnerton-Dyer Conjecture for elliptic curves.

Contents

CHAPTER 1

Introduction

1.1 Framework and objects of study

The content of this thesis is situated between number theory and cryptography. To give the reader an idea for the topics of interest in this thesis, let us briefly describe them and explain their relevance for modern number theory and cryptography.

Algorithms for lattice problems

Geometry of numbers is the area of number theory dealing with lattices and convex bodies, and originated in important work of many mathematicians, such as Fermat, Euler, Lagrange and Minkowski (see e.g. [Cas71] for an introduction). Lattices not only have a rich history, but occupy today a fundamental place in number theory and computer science, with a variety of applications.

One of the most famous computational problems on lattices is the *shortest vector problem*, or shortly, SVP. The problem asks to compute a non-zero vector of minimal Euclidean norm in a lattice. In high dimensions this is difficult, and exact algorithms are expensive (see e.g. [AKS01] for a randomized algorithm). This is why approximation algorithms are frequently employed. Such algorithms are based on lattice reduction, which issued from the revolutionary article by Lenstra, Lenstra and Lovasz [LLL82] from 1982, leading to the LLL-reduction algorithm. This algorithm allows to compute somewhat short and nearly orthogonal bases of lattices efficiently. This solves an approximate-version of SVP, in that the norm of the shortest basis vector computed by LLL is an approximation of the norm of an actual shortest non-zero vector in the lattice. Many variations of the LLL algorithm have appeared in the literature (see e.g. [GHGKN06a, CN11]). The theory of lattice reduction gains a lot of interest among cryptographers and number theorists and serves in many applications (such as factoring integer polynomials or simultaneous Diophantine approximation, as described in [LLL82]). Moreover, it is soon recognized as a fundamental ingredient in lattice-based cryptanalysis, by a range of works (see e.g. [NS01]). Also many constructions based on lattices have been proposed, making lattices even more interesting. The first cryptographic constructions based on lattices are proposed by Ajtaj in 1996, [Ajt96]. Subsequent works include the well-studied public-key encryption scheme NTRU [HPS98] by Hoffstein, Pipher and Silverman, and the corresponding signature schemes NTRUSign [HHGP$^+$03]. The security of these schemes is based on the hardness of lattice problems, such as SVP. In 2005,

Regev introduces another public-key encryption scheme whose security is proven under the hardness of the Learning with Errors Problem (LWE), [Reg05]. The LWE Problem is later used to build even more advanced cryptographic primitives, such as the fully homomorphic encryption scheme by Gentry [Gen09].

Today, post-quantum cryptography largely benefits from the theory of lattices and lattice-based cryptography is entirely devoted to constructions involving lattices. Some constructions based on lattices are important candidates for post-quantum cryptography, as they appear to be resistant against classical as well as quantum attacks, unlike discrete logarithm or factoring-based constructions. Moreover, their simple design makes them a strong candidate for the NIST supported by the NIST Post-Quantum Cryptography Standardization Competition.

In this thesis, we largely deal with lattice problems and the design of algorithms for them. These algorithms, based on lattice reduction, are shown powerful in the area of cryptanalysis against cryptographic multilinear maps, and more general problems of interest in algorithmic number theory. By studying and comparing the geometry of certain lattices, we derive new and competitive algorithms.

Cryptographic Multilinear Maps and their cryptanalysis

The goal of study for *cryptographic multilinear maps* is to generalize bilinear pairings to higher degrees of linearity. Bilinear pairings over elliptic curves, or more generally, abelian varieties, have been extensively studied (see e.g. [Gal05] for an exposition) and are at the heart of pairing-based cryptography. Moreover, they come with a range of powerful applications, such as, non-interactive key exchange among three users (see [Jou04]), generalizing the DH key exchange from [DH82], identity-based encryption [BF01], and short signatures [BLS01]. Thus, it is a natural question to see to what extent such applications hold in the multilinear framework. A survey put forward by Boneh and Silverberg [BS03] points out the difficulties of constructing multilinear maps of degree higher than 2 from natural maps in algebraic geometry.

Inspired by the work on fully homomorphic encryption [Gen09], the work of Gentry et al. [GGH13a] makes a step forward in this direction, by introducing the notion of *graded encoding system*. Such a system is somehow functionally equivalent to a cryptographic multilinear map, and targets very similar applications. The construction by Gentry et al. leads to the GGH13 multilinear map scheme. Shortly after, two alternative constructions of graded encoding schemes appeared in the literature; the CLT13 scheme by Coron, Lepoint and Tibouchi [CLT13] and the GGH15 scheme by Gentry, Gorbunov and Halevi [GGH15]. To this date, those are the only multilinear map schemes described.

The security of the current constructions of multilinear maps is still today rather poorly understood, and many powerful attacks have appeared over the last years. An important family of attacks against multilinear maps are so-called "zeroizing attacks" (see [CHL$^+$15, CGH$^+$15, HJ16, CLLT16]), which efficiently recover the secret parameters from encodings of zero and break the key exchange protocols. However, more complex constructions based on multilinear maps are not necessarily broken, giving hope for these constructions to be useful. For example, for indistinguishability obfuscation (iO) (see e.g. [GGH$^+$13b] for GGH13), low-level encodings of zero are generally not available, preventing zeroizing attacks. A range of attacks has appeared in the works [CGH$^+$15, MSZ16, CLLT17, CGH17, CVW18]. In gen-

eral, the above attacks only apply against branching programs with a simple structure, and breaking more complex constructions (such as dual-input branching programs) is currently infeasible.

In this thesis, we deal with the CLT13 multilinear map scheme and contribute to new cryptanalysis in several directions. We extend an attack by Gentry et al. [GLW14], as well as the powerful attack by Cheon et al. [CHL$^+$15].

Topics in Arithmetic and Geometry of Elliptic Curves

Elliptic curves are fundamental objects among algebraic geometers, number theorists, and cryptographers. Problems in Diophantine geometry, which deals with polynomial equations in algebraic geometry and algebraic number theory, are often related to elliptic curves. Still today, they form an active area of research with a number of open questions.

In number theory, one of the most famous examples using the theory of elliptic curves is the proof of Fermat's Last Theorem, by Taylor-Wiles (see [Wil95, TW95] and [Gou94] for an exposition). As shown by the proof (e.g. the modularity theorem, known as the Taniyama-Shimura-Weil conjecture), elliptic curves find ramifications within the theory of modular forms and Galois representations. Elliptic curves are also especially important in analytic number theory, with the study of their L-functions (or L-series). Most importantly, the conjecture by Birch and Swinnerton-Dyer on the rank of an elliptic curve is one of the most important open questions in number theory and mathematics.

With the works by Koblitz and Miller (see [Kob87, Mil86]), elliptic curves entered in use in cryptography, and widened 20 years later, to form *elliptic-curve cryptography* (ECC). Cryptosystems over elliptic curves share similarities with protocols over abelian groups, and rely on the hardness of the discrete logarithm problem, which has a natural analogue (ECDLP) on the group of points on elliptic curves defined over a finite field. Conjecturally, ECDLP is harder to solve than DLP in the multiplicative group of a finite field. Modern ECC includes a Diffie-Hellman key exchange over elliptic curves (ECDH) as well as a digital signature algorithm (ECDSA). While ECC is also under the threat of classical exponential attacks and quantum polynomial attacks (Schor's quantum algorithm, [Sho97]), it continues to be an active area of research among cryptographers. For security or efficiency reasons in cryptography, it is of great importance to select good representations of elliptic curves. Next to the Weierstrass model, which is probably the best-known model, many other models of elliptic curves exist. For example, Koblitz, Montgomery, Edwards, or Hessian curves. Such models often prevent side-channel attacks, which allow the attacker to learn private data from the algorithm (see e.g. [JQNP01, BL07] and more generally, [Lan05]). Also, protocols, such as digital signatures, are designed over different models than the Weierstrass model (see e.g. [BDL$^+$11]). Other well-known algorithms using elliptic curves are Schoof's point counting algorithm on elliptic curves over finite fields [Sch85] and Lenstra's algorithm for integer factorization [Len87].

In this thesis, we consider Edwards curves, introduced in the works of H. Edwards [Edw07] and proposed for cryptographic use by Bernstein and Lange, [BL07]. While our questions of interest do not directly target cryptographic applications, we study questions of arithmetic-geometric flavour on these curves. Also, we provide statistics for the ranks of these curves.

1.2 Main contributions and results of the thesis

The contributions described in this thesis build on four different works:

First, we present new contributions to cryptanalysis of cryptographic multilinear maps. This is joint work with Jean-Sébastien Coron, published at Asiacrypt 2019, [CN19a]. We provide a description in Section 1.2.1.

Second, we investigate the Hidden Lattice Problem and its ramifications to cryptographic contexts. This is joint work with Gabor Wiese and an article has been submitted. We provide a description in Section 1.2.2.

Third, we contribute to new algorithms in number-theoretic and cryptographic contexts, by studying the problem of simultaneously diagonalizing incomplete matrices. This is joint work with Jean-Sébastien Coron and Gabor Wiese, published at the Fourteenth Algorithmic Number Theory Symposium 2020, [CNW20a]. We provide a description in Section 1.2.3.

Last, we contribute to several generalizations and new results on questions related to Edwards curves. We report on work in progress, based on joint work with Samuele Anni. We provide a description in Section 1.2.4.

1.2.1 Cryptanalysis of a Multilinear Map Scheme

1.2.1.1 CLT13 Multilinear maps and their cryptanalysis

The CLT13 multilinear map scheme is a graded encoding scheme constructed over the ring of integers $\mathbb{Z}$. We give a high-level description of this scheme here, and postpone more details to Chapter 3, Section 3.2. Let $n \geq 1$ be an integer regarded as a dimension ensuring correctness and security of the scheme. The instance generation of CLT13 generates n distinct secret "large" prime numbers $p_1, \ldots, p_n$ of bit size η, and publishes $x_0 = \prod_{i=1}^{n} p_i$; the values n and η are so that direct factorization of x_0 is intractable in classical polynomial time. Further, one generates n distinct secret "small" prime numbers $g_1, \ldots, g_n$ of bit size α. We do not make the words "large" and "small" precise here, but the reader should interpret a significant size difference between α and η. The plaintext ring is the composite ring $\mathbb{Z}/G\mathbb{Z} \simeq \bigoplus_{i=1}^{n} \mathbb{Z}/g_i\mathbb{Z}$ where $G = \prod_{i=1}^{n} g_i$. Therefore, plaintext elements are vectors $m = (m_1, \ldots, m_n)$ with components defined modulo $\{g_i\}_i$. Let $\kappa \in \mathbb{Z}_{\geq 1}$ be the multilinearity degree of the multilinear map constructed by the scheme. While the prime numbers $\{g_i\}_i$ define the plaintext ring, the prime numbers $\{p_i\}_i$ define the encoding ring. For $k \in \{1, \ldots, \kappa\}$, an encoding at *level* k of $(m_i)_i \in \bigoplus_i \mathbb{Z}/g_i\mathbb{Z}$ is an integer $c \in \mathbb{Z}$ satisfying congruences $c \equiv (r_i g_i + m_i) z^{-k} \pmod{p_i}$ for all $1 \leq i \leq n$, for "small" random integers $\{r_i\}_i$; we refer to $\{r_i\}_i$ as the "noise" in the encodings – this is why CLT13 (and more generally, graded encoding schemes) is regarded as a noisy multilinear map model. The integer z is random and invertible modulo x_0 (with inverse z^{-1}) and is kept secret. A main feature of graded encoding schemes is that the encoding space supports homomorphic operations; for example, two encodings at the same level can be added, and the underlying plaintexts get added in the plaintext ring. Similarly, the product of two encodings at level i and j gives an encoding of the product plaintexts at level

$i + j$, in general (i.e. up to a technical constraint on the size of the parameters). Encodings of level $k = \kappa$ are called *top-level encodings* and play an important role – one can publicly test whether they encode the zero message or not. By the homomorphic property, this functionality also allows testing equality of messages. Let us briefly outline how zero-testing for a top-level encoding c works. One defines and publishes a *zero-testing parameter*, denoted by p_{zt}. It is not important to state the exact formula for p_{zt} at this point; it is defined modulo x_0 and is a multiple of z^κ. Zero-testing of a top-level encoding c (recall, that this is an integer satisfying the above congruences, for $k = \kappa$) consists in computing $c \cdot p_{zt} \in \mathbb{Z}/x_0\mathbb{Z}$ (using the Chinese Remainder Theorem) and checking whether the result (i.e. the unique representative in $[0, x_0) \cap \mathbb{Z}$), commonly denoted by ω, is small enough compared to x_0. The expression for p_{zt} contains the term z^κ, so that it is canceled with the denominator in the encoding c. We denote by ν the number of bits that can be extracted from zero-testing; i.e. the ν most significant bits of ω only depend on the plaintext messages, and not on the noise in the encodings.

1.2.1.2 Our contribution

Analysis of the attack from [GLW14]

In [GLW14, Appendix B], the authors observe that instantiating CLT13 with independent slots leads to a simple lattice-based attack in dimension 2, which efficiently recovers the (secret) plaintext ring $\bigoplus_{i=1}^{n} \mathbb{Z}/g_i\mathbb{Z}$. Namely, in this case, the attacker can obtain encodings of messages of which all slots are zero, except one. In simplified notation, assuming that the underling message is non-zero modulo g_{i_0} and zero modulo all other primes $\{g_i\}_{i \neq i_0}$ (without loss of generality, we may assume $i_0 = 1$), then by correct zero-testing, the attacker can derive an equation of the form

$$g_{i_0}\omega \equiv u_0 \pmod{x_0} \tag{1.1}$$

for certain specific integers ω and u_0, and where u_0 is significantly smaller than x_0, and ω has approximately the size of x_0. Here ω and x_0 are public and g_{i_0} and u_0 are secret. Since the prime g_{i_0} is rather small, lattice reduction applied to a lattice of dimension two constructed from ω and x_0 then allows computing the secret prime g_{i_0}.

To prevent this line of attack, [GLW14] considers the following countermeasure: instead of working with a single slot defined modulo g_{i_0}, work modulo several primes $\{g_i : i \in I\}$ for some $I \subseteq \{1, \ldots, n\}$; in this way one does not have a message which is non-zero modulo a single isolated prime, i.e. either it is zero modulo g_{i_0} and all other prime g_i with $i \in I$, or it is non-zero modulo all of them. Let $\theta := \#I$; thus, every slot is defined modulo a product of θ primes $\{g_i\}_i$, which gives a total of $\lfloor n/\theta \rfloor$ plaintext slots, instead of n. It is easy to see that the above mentioned attack continues to hold for g_{i_0} replaced by $g = \prod_{i \in I} g_i$, that is, an attacker can again derive an equation of the type

$$g\omega \equiv u \pmod{x_0} \tag{1.2}$$

for certain specific integers ω and u, with u again somewhat smaller than x_0. However, for a sufficiently large set I (i.e. sufficiently large θ), g is too large to be recovered by lattice reduction and the attack is thwarted. Our first contribution is a rigorous study of the 2-dimensional attack indicated in [GLW14]. More precisely, let ν be the number of extracted bits from zero-testing. We explain that under the simplified condition

$$\alpha\theta < \nu/2 \tag{1.3}$$

one can reveal g and thus its prime factors $g_1, \ldots, g_\theta$, by a two-dimensional lattice attack. We refer to Proposition 3.3.1 and the subsequent analysis.

Breaking the countermeasure from [GLW14]

We extend the attack mentioned above, to break the countermeasure for even larger values of θ. Our new attack is based on lattice reduction in *higher* dimensions and is described in two settings. The first setting amounts to a linear improvement on Equation (1.2), while the second one gives a quadratic improvement on Equation (1.2). For simplicity, we here provide a more extensive summary for the first setting; the second one works similarly with appropriate generalizations. Let $\ell \geq 1$ be another integer, and consider ℓ encodings $\{c_j : 1 \leq j \leq \ell\}$ where the corresponding message vectors $\{m_j = (m_{ji})_i : 1 \leq j \leq \ell\}$ only have θ non-zero components modulo the primes $\{g_i\}_i$. Then, the previous attack corresponds to the case $\ell = 1$. For $1 \leq i \leq \theta$, we let $\hat{m}_i = (m_{ji})_j \in \mathbb{Z}^\ell$ be the vector corresponding to the ith non-zero component of $\{m_j\}_j$. Instead of deriving an equation similar to Equation (1.2), an attacker can now derive, after multiple zero-testing evaluations, a vector equation modulo x_0 in dimension ℓ, and we justify that it can be written under the form

$$\omega \equiv \sum_{i=1}^{\theta} \alpha_i \hat{m}_i + R \pmod{x_0} . \tag{1.4}$$

Here, $\omega \in \mathbb{Z}^\ell$ is a vector corresponding to the zero-tested values of the encodings $\{c_j\}_j$, $\{\alpha_i\}_i$ are certain integers (each satisfying a specific congruence relation modulo the prime factors of x_0) and $R \in \mathbb{Z}^\ell$ is a certain vector with rather small (compared to x_0) entries and which we regard as a *noise* term. Abstracting the specific CLT13-based framework, we justifiedly recognize Equation (1.4) as a variant of the hidden subset sum problem, studied by Nguyen and Stern in [NS99], for which they propose an algorithm relying on the powerful notion of the *orthogonal lattice*. Our Equation (1.4) while being close to a hidden subset sum-type equation, nevertheless presents crucial deviations from this problem (for instance, the presence of structured coefficients $\{\alpha_i\}_i$ and the presence of a non-zero noise vector R), which allow us to follow a different path of attack. For our attack, we mainly consider two lattices which we put in relation. First, we consider the (scaled variant of the) lattice $\mathcal{L} = \mathcal{L}(\omega, x_0)$ in dimension $\ell + 1$, constructed from the public data ω and x_0, consisting of elements $(u, v) \in \mathbb{Z}^\ell \times \mathbb{Z}$ orthogonal to $(\omega, 1) \in \mathbb{Z}^\ell \times \mathbb{Z}$ over $\mathbb{Z}/x_0\mathbb{Z}$, that is, $\langle (u, v), (\omega, 1) \rangle = \langle u, \omega \rangle + v \in x_0\mathbb{Z}$; the attacker can easily compute a basis of $\mathcal{L}$. Next, we consider the lattice $\Lambda := \{u \in \mathbb{Z}^\ell : \langle u, \hat{m}_i \rangle \equiv 0 \pmod{g_i},\ 1 \leq i \leq \theta\}$, which, if none of the message vectors $\hat{m}_i$ is a multiple of g_i, has volume $g = \prod_{i=1}^{\theta} g_i$ – a key fact for our cryptanalysis. This lattice is of course not known in advance (in the language of Section 1.2.2 of this chapter, we shall call it *hidden*), as the primes $\{g_i\}_i$ and the messages $\{\hat{m}_i\}_i$ are secret. Following the intuition of [NS99], we argue that sufficiently short vectors $(u, v) \in \mathcal{L}$ must have $u \in \Lambda$. In fact, as in the case of the hidden subset sum problem in [NS99], one would be tempted to expect, from short vectors (u, v) of $\mathcal{L}$, the stronger property, that $u \in \Lambda_0$ where Λ_0 is the lattice of vectors that are $\mathbb{Z}$-orthogonal to the vectors $\hat{m}_i$, that is u has inner product zero over $\mathbb{Z}$ with all the vectors $\{\hat{m}_i\}_i$. This requirement is however too strong in view of the extra congruence relations on the coefficients $\{\alpha_i\}_i$, which explains why we only obtain the weaker orthogonality over $\mathbb{Z}/g_i\mathbb{Z}$ for all i. We compute short vectors in $\mathcal{L}$ by lattice reduction and derive that heuristically, under the approximate

condition $(1 + 1/\ell)\,\alpha\theta < \nu$ one can compute a reduced basis $\{(u_j, v_j)\}_j$ of $\mathcal{L}$ having the first ℓ (projected) vectors $\{u_j\}_j$ in Λ. Since Λ has volume g, we reveal g as the determinant of the $\ell \times \ell$ matrix, say U, formed by the vectors $\{u_j\}_j$. Note that for moderately large values of ℓ, the mentioned bound gives the simpler condition

$$\alpha\theta < \nu\,, \tag{1.5}$$

which improves Equation (1.2) by a factor 2.

The second setting of our attack aims at further improving Equation (1.5). Therefore, we introduce yet another integer parameter $d \geq 1$ and consider products of encodings of the form $c_j \cdot d_k$ for $1 \leq j \leq \ell$ and $1 \leq k \leq d$, where as previously, the underlying messages of the encodings $\{c_j\}_j$ have only θ non-zero components modulo the primes $\{g_i\}_i$. In that case, we justify that using a variant of the previous lattice attack (i.e. running lattice reduction on a generalization of $\mathcal{L}$), this time in dimension $\ell + d$, the bound asymptotically improves to:

$$\alpha\theta = O(\nu^2) \tag{1.6}$$

The above bound also applies when a vector of zero-testing elements is available, instead of a single zero-test parameter p_{zt}. We refer the reader to Section 3.3.6 for a precise description of the lattice attacks and Algorithm 3 for a description of the algorithm.

Our lattice attack not only applies for a much larger range of parameters than [GLW14], but is also more powerful. Namely, in addition to recovering the secret plaintext ring, we also show how to recover the plaintext messages $\{\hat{m}_i\}_i$ and thus $\{m_j\}_j$, up to a scaling factor. We propose two algorithms for this purpose, relying once on the factorization of g, feasible as its prime factors are not too large, and once on the Chinese Remainder Theorem without requiring factorization of g. For the precise content, we refer to Algorithms 4 and 5. Both algorithms make use of the matrix U derived from the vectors $\{u_j\}_j$ obtained by our lattice attack.

Our practical experiments (see Chapter 3, Section 3.3.8) confirm the theory. Concretely, for the original parameters from [CLT13], our attack takes a few seconds for $\theta = 40$, and a few hours for θ as large as 160, while the original attack from [GLW14] only works for $\theta = 1$. Finally, we propose a set of secure parameters for CLT13 multilinear maps that heuristically prevents our extended attack. For $\lambda = 80$ bits of security, we recommend to take $\theta \geq 1789$.

Recovering all the secret parameters of CLT13

For the range of parameters satisfying Equation (1.6), we next show how to combine our algorithm with the Cheon et al. attack from [CHL$^+$15], in order to reveal all secret parameters of CLT13, when instantiated with independent slots. More precisely, assuming that multiple intermediate-level encodings $\{c_j\}_j$ of partially zero messages are available, our approach consists in applying our lattice attack to generate intermediate-level encodings of zero. To do so, note that we crucially rely on computing relatively short vectors $\{u_j\}_j$ in Λ. By the linearity of the inner product, we easily see that for every j, the inner product of u_j with c_j is an encoding of zero (at the same level than c_j), with mildly larger noise than for the encoding c_j. Subsequently, we can apply the original Cheon et al. attack on these newly-created encodings of zero, which reveals the secret prime factors $\{p_i\}_i$ of x_0, and in particular, all the secret CLT13 parameters. This also contributes to an open problem described in [CFL$^+$16, Section 4], that

of cryptanalyzing CLT13 when no encodings of zero are available beforehand. Namely, the attack by Cheon et al. requires to have encodings of zero available to the attacker, which we create by our lattice-based attack in the first place.

We can summarize our complete cryptanalysis of CLT13 multilinear maps with independent slots as follows.

Proposition. *Consider the standard CLT13 notation with $G = \prod_{i=1}^{n} g_i$ for prime numbers $\{g_i\}_i$ defining the plaintext ring, and $x_0 = \prod_{i=1}^{n} p_i$ for prime numbers $\{p_i\}_i$ defining the encoding ring. Let $\ell \geq 1$ and $1 \leq \theta \leq n$ be integers. Let $\{c_j : 1 \leq j \leq \ell\}$ be encodings of elements $\{m_j : 1 \leq j \leq \ell\}$ such that $m_{ji} \not\equiv 0 \pmod{g_i}$ for $1 \leq i \leq \theta$ and $1 \leq j \leq \ell$. Let $\{\hat{m}_i : 1 \leq i \leq \theta\}$ be the vectors corresponding to the non-zero entries of $\{m_j\}_j$. Assume the asymptotic bound $\alpha\theta = O(\nu^2)$.*

(i) There exists an algorithm (Algorithm 3), which on input the encodings $\{c_j\}_j$, computes the secret factor $g = \prod_{i=1}^{\theta} g_i$ of G in heuristic polynomial time.

(ii) There exists an algorithm (Algorithm 4), which on input the encodings $\{c_j\}_j$, computes multiples $\{\lambda_i \hat{m}_i\}_i$ with $\lambda_i \not\equiv 0 \pmod{g_i}$ for $1 \leq i \leq \theta$, in sub-exponential time.

There exists an algorithm (Algorithm 5), which on input the encodings $\{c_j\}_j$, computes a vector $\lambda\hat{m}$ such that $\gcd(\lambda, g) = 1$ and $\hat{m} \equiv \hat{m}_i \pmod{g_i}$ for all $1 \leq i \leq \theta$, in polynomial time.

(iii) There exists an algorithm (Algorithm 6), which on input the encodings $\{c_j\}_j$ and a sufficiently large set of encodings, which is disjoint from the set $\{c_j\}_j$, computes the prime factors $\{p_i\}_i$ of x_0 in polynomial time.

Application to CLT13-based constructions

Finally, we show how our attack impacts the parameters of several schemes based on CLT13 multilinear maps with independent slots. More precisely, we here consider the constructions from [GLW14, GLSW15, Zim15] and [FRS17]. First, we investigate the "multilinear subgroup elimination assumption", considered in [GLW14] and [GLSW15]. Following the line of our lattice-based attack, we describe a distinguishing attack for the range of parameters given by Equation (1.6). In [Zim15], Zimmerman describes a technique of program obfuscation based on composite-order multilinear maps and without relying on matrix branching programs. We explain that when instantiated with CLT13, the technique is subject to the constraint in Equation (1.6). Finally, we consider the construction from [FRS17] by Fernando et al. for preventing certain attacks against matrix branching programs, also described on CLT13 with independent slots. We argue how our lattice attack allows the recovery of the secret CLT13 plaintext ring for the range of parameters given by Equation (1.6). We note however that breaking the indistinguishability of the branching program remains an open problem. We describe these applications in Section 3.3.10.

In summary, our lattice-based attack breaks the countermeasure proposed in [GLW14] for a wide range of parameters, and comes with interesting cryptanalytic by-products: the recovery of the secret information about the plaintext messages, as well as the technique to generate encodings of zero, enabling the Cheon et. al attack to compute all the secret parameters. Further, concrete instantiations of CLT13 multilinear maps are found insecure in the range of parameters targeted by our attack, and therefore subjected to a new parameter selection.

1.2.2 The Hidden Lattice Problem

1.2.2.1 Motivation

The security of many cryptographic schemes based on lattices relies on the assumption of the intractability to efficiently solve certain computational lattice problems, often based on short vectors. Thus, it is natural to consider lattices possessing generating sets consisting of short vectors. The geometry of such lattices is usually impacted by the existence of such vectors. For example, for the hidden subset sum problem, one considers a public vector, say $v \in \mathbb{Z}^m$, such that every entry of v is a (binary) subset sum modulo a public integer N; in vector notation, this reads $v \equiv \sum_{i=1}^{n} \alpha_i v_i \pmod N$ for some integer $n < m$, a set of hidden binary vectors $\mathfrak{B} = \{v_i\}_i \subseteq \{0,1\}^m$ and hidden integers $\{\alpha_i\}_i$. The problem then asks to compute $\{v_i\}_i$ and $\{\alpha_i\}_i$, given v and N. It is likely that the binary vectors $\{v_i\}_i$ are linearly independent, thus generating a lattice $\mathcal{L} \subseteq \mathbb{Z}^m$ of rank n. Although the lattice $\mathcal{L}$ admits infinitely many bases (if $n > 1$) containing arbitrarily large vectors, the existence of the particularly "small" basis $\mathfrak{B}$ heavily impacts the geometry of $\mathcal{L}$. The problem of computing $\{v_i\}_i$ such that $v \equiv \sum_i \alpha_i v_i \pmod N$ can be weakened to the problem of computing any basis of $\mathcal{L}$ such that $v \in \mathcal{L} \pmod N$ and $\mathcal{L} = \bigoplus_i \mathbb{Z} v_i$.

1.2.2.2 Our contribution

We first propose a formal definition for the hidden lattice problem, following [NS99]. Our definition is however a bit more general and contains more parameters. Then we describe two algorithms for it, which we analyse theoretically, heuristically and practically. Finally, we give two variations of the problem and list a number of concrete applications.

To define the hidden lattice problem, we first define "small" lattices. We quantify this notion by a positive real parameter μ: consider that a lattice $\mathcal{L}$ is *μ-small* if it possesses a basis $\mathfrak{B}$ satisfying $\sigma(\mathfrak{B}) := ((\#\mathfrak{B})^{-1} \cdot \sum_{v \in \mathfrak{B}} \|v\|^2)^{1/2} \leq \mu$. Note that $\#\mathfrak{B}$ is the same for every basis and is equal to the rank of $\mathcal{L}$. Of course, there are many ways to define small lattices, and our choice of σ is not a canonical one. However, we justify that this "measure" suits our theoretical study nicely. Lattices admitting bases consisting of binary vectors (as in the hidden subset sum problem), or more generally, vectors with bounded entries, all fall into our family of small lattices for appropriate choices of μ. It is clear by this definition, that, for every lattice $\mathcal{L}$ there exists $\mu \in \mathbb{R}$ such that $\mathcal{L}$ is μ-small. Further, if $\mathcal{L}$ is μ-small then $\mathcal{L}$ is also μ'-small for every $\mu' \geq \mu$. As such, our definition of being μ-small might seem uninteresting at first sight; however, for the purpose of our study, we consider μ to be rather small, in which case, a lattice $\mathcal{L}$ being μ-small becomes a strong assumption on $\mathcal{L}$. As is customary in many computational problems we also work modulo $N \in \mathbb{Z}$ and write $v \in \mathcal{L} \pmod N$ if there exists $w \in \mathcal{L}$ such that $v - w \in (N\mathbb{Z})^m$. If $\mathcal{M} \subseteq \mathbb{Z}^m$, then $\mathcal{M} \subseteq \mathcal{L} \pmod N$ means $v \in \mathcal{L} \pmod N$ for all $v \in \mathcal{M}$. For the hidden lattice problem, we assume that one knows a basis of some lattice $\mathcal{M}$ such that $\mathcal{M} \subseteq \mathcal{L} \pmod N$ and $\mathcal{L}$ is *small*; this generalizes the setting of the hidden subset sum problem. Formally, we define the hidden lattice problem (HLP) as follows.

Definition (Hidden Lattice Problem, see Definition 4.1.2)**.** *Let $\mu \in \mathbb{R}_{\geq 1}$, integers $1 \leq r \leq n \leq m$ and $N \in \mathbb{Z}$. Let $\mathcal{L} \subseteq \mathbb{Z}^m$ be a μ-small lattice of rank n. Further, let $\mathcal{M} \subseteq \mathbb{Z}^m$ be a lattice of rank r such that $\mathcal{M} \subseteq \mathcal{L} \pmod N$.*

The Hidden Lattice Problem (HLP) *is the task to compute from the knowledge of n, N and a basis of $\mathcal{M}$, a basis of the completion of any μ-small lattice Λ of rank n such that $\mathcal{M} \subseteq \Lambda \pmod N$.*

Here we used the following piece of notation: let $\mathcal{L}_{\mathbb{Q}}$ (resp. $\mathcal{L}_{\mathbb{R}}$) be the $\mathbb{Q}$-span (resp. $\mathbb{R}$-span) of $\mathcal{L}$ in $\mathbb{R}^m$; the *completion* $\overline{\mathcal{L}}$ of $\mathcal{L}$ is $\mathcal{L}_{\mathbb{Q}} \cap \mathbb{Z}^m = \mathcal{L}_{\mathbb{R}} \cap \mathbb{Z}^m$ and we say that $\mathcal{L}$ is *complete* if $\overline{\mathcal{L}} = \mathcal{L}$. We will see that $\overline{\mathcal{L}}$ is very often equal to $\mathcal{L}$: it is the *hidden lattice* to be revealed. We analyse for which values of μ a generic HLP can be expected to be solvable.

Algorithms for the HLP

We describe two algorithms for the HLP, which we denote by Algorithm I and II, respectively. Our algorithms rely on the public lattices $\mathcal{M}^{\perp_N}$ of vectors orthogonal to $\mathcal{M}$ modulo N (we call this the *N-orthogonal lattice of* $\mathcal{M}$), and $\mathcal{M}_N$ consisting of vectors lying in $\mathcal{M}$ modulo N (we call this the *N-congruence lattice* of $\mathcal{M}$), respectively. For both algorithms, we employ in first place a lattice reduction algorithm (e.g. LLL) on bases of $\mathcal{M}^{\perp_N}$ and $\mathcal{M}_N$ to compute certain lattices $\mathcal{N}_{\mathrm{I}}$ and $\mathcal{N}_{\mathrm{II}}$, from which the completion of the hidden lattice $\mathcal{L}$ is revealed in a second step.

Our first algorithm (Algorithm I, or Algorithm 7) is a direct adaptation of the orthogonal lattice algorithm proposed by Nguyen and Stern [NS99] in the context of the hidden subset sum problem. In fact, we notice that this algorithm extends to our more general setting. It is based on the orthogonal lattice, which was formally introduced in [NS97] as a strong tool for lattice-based cryptanalysis. More precisely, $\mathcal{M}^{\perp_N}$ naturally contains $\mathcal{L}^{\perp}$, a relatively small lattice of rank $m-n$, of which we identify a sublattice $\mathcal{N}_{\mathrm{I}}$ of the same rank by lattice reduction, provided that the parameters satisfy appropriate conditions. From $\mathcal{N}_{\mathrm{I}}^{\perp}$, the completion of $\mathcal{L}$ is then revealed by computing $(\mathcal{N}_{\mathrm{I}}^{\perp})^{\perp} = \overline{\mathcal{L}}$, which solves the HLP.

Our main contribution is to propose an alternative algorithm (Algorithm II, or Algorithm 8), based on the (public) lattice $\mathcal{M}_N$, which lies in $\mathcal{L}_N$, by assumption. In this case, we explain how to recognize and compute n linearly independent short vectors, and thus an entire sublattice $\mathcal{N}_{\mathrm{II}}$ of rank n generated by those vectors, lying in $\overline{\mathcal{L}}$ directly, if the parameters are suitable. We again compute these vectors (thus $\mathcal{N}_{\mathrm{II}}$) by lattice reduction on $\mathcal{M}_N$. More precisely, we derive an explicit lower bound on the norm of the vectors lying in $\mathcal{M}_N$ but outside $\mathcal{L}_{\mathbb{Q}}$, which gives us a criterion for establishing an explicit parameter selection. To obtain the completion $\overline{\mathcal{L}}$ from $\mathcal{N}_{\mathrm{II}}$, it is enough to complete $\mathcal{N}_{\mathrm{II}}$. We observe experimentally that this can be carried out locally at N, speeding up our algorithm. Our new algorithm is therefore a competitive alternative to the celebrated orthogonal lattice algorithm.

We finally show that both algorithms are related by duality theory. Namely, the lattices $\mathcal{M}^{\perp_N}$ (considered for Algorithm I) and $\mathcal{M}_N$ (considered for Algorithm II) are dual up to scalar multiplication by N, that is, $\mathcal{M}^{\perp_N} = N(\mathcal{M}_N)^{\vee}$, where the symbol $(\cdot)^{\vee}$ stands for the dual operator. Using celebrated transference results for the successive minima of dual lattices, we explain how to bridge both algorithms theoretically.

Analysis of our algorithms

We provide a heuristic analysis of our algorithms based on the Gaussian Heuristic for "random lattices". Namely, for random instances of the HLP, we can consider all the lattices to behave like random lattices. For Algorithm I, we follow the intuition of [NS99]: short enough vectors $u \in \mathcal{M}^{\perp_N}$ (which we compute by lattice reduction) must lie in $\mathcal{L}^{\perp}$. Since $\mathcal{L}^{\perp}$ has rank $m-n$, we expect to find $m-n$ such vectors. For Algorithm II, we derive an explicit lower bound on the norm of the vectors lying in $\mathcal{M}_N$ but outside $\mathcal{L}_{\mathbb{Q}}$, which gives us a criterion for establishing an explicit parameter selection. In both cases, it turns out that the HLP is solv-

able when the size difference between N and μ is sufficiently large, which we read by explicit lower bounds for N in terms of μ, resp. upper bounds for μ in terms of N. For example, both algorithms detect hidden lattices of size $\mu = O(N^{\frac{r(m-n)}{nm}})$ up to some terms which differ according to the algorithm. In the balanced case $m = 2n = 4r$, this gives $\mu = O(N^{1/4})$. To quantify the gap between N and μ in a compact formula, we propose a definition for an arithmetic invariant attached to the problem, defined by:

$$\Delta := \log\left(\frac{N^{r/n}}{\mu^{m/(m-n)}}\right) \tag{1.7}$$

We can derive an explicit lower bound for Δ from our heuristic analyses. The larger Δ is, the easier it is to compute a solution by our algorithms. In this respect, Δ behaves much like an inverse-density, a handy and well-studied invariant for knapsack-type problems (see e.g. [LO85, NS99]). Asymptotically, our heuristic analyses behave differently for the two algorithms, which we see when establishing a growth comparison of the error terms in our bounds.

Proposition. *(i) There exists an algorithm (Algorithm I, or Algorithm 7) which on input N and a basis of $\mathcal{M}$, computes a basis of $\overline{\mathcal{L}}$ in heuristic polynomial time under the asymptotic condition $\Delta = O(n^\ell)$ when $m = O(n^\ell)$ for every $\ell \geq 1$.*

(ii) There exists an algorithm (Algorithm II, or 8) which on input N and a basis of $\mathcal{M}$, computes a basis of $\overline{\mathcal{L}}$ in heuristic polynomial time under the asymptotic condition $\Delta = O(n^\ell)$ when $m = O(n^\ell)$ for every $\ell \geq 1$.

Along with our heuristic analyses, we also establish proven results, not relying on the Gaussian Heuristic. Note that a rigorous proof was not included in [NS99]. To do so, we rely on a discrete counting technique. For a fixed μ-small basis $\mathfrak{B}$ of $\mathcal{L}$ (sampled from some set Ω of collections of vectors) and a given integer N, we denote by $\mathcal{H}(\mathfrak{B})$ a finite sample set of hidden lattice problems constructed from $\mathfrak{B}$ and N. To an element of $\mathcal{H}(\mathfrak{B})$, one can naturally associate a hidden lattice problem with hidden lattice $\mathcal{L}$. On each of these problems, we theoretically run either Algorithm I or Algorithm II, and denote the subset of $\mathcal{H}(\mathfrak{B})$ for which Algorithm I (resp. Algorithm II) successfully computes a basis of $\overline{\mathcal{L}}$ by $\mathcal{H}_{\mathrm{I}}(\mathfrak{B})$ (resp. $\mathcal{H}_{\mathrm{II}}(\mathfrak{B})$). Since we rely on the LLL algorithm with reduction parameter $\delta \in (1/4, 1)$, our more complete notation is $\mathcal{H}_{\delta,\mathrm{I}}(\mathfrak{B})$ (resp. $\mathcal{H}_{\delta,\mathrm{II}}(\mathfrak{B})$). Our technique aims at maximizing the proportion $\#\mathcal{H}_{\delta,\mathrm{I}}(\mathfrak{B})/\#\mathcal{H}(\mathfrak{B})$ and similarly for Algorithm II. More concretely, we prove an explicit version of the following theorem, for which we refer the reader to Chapter 4, Section 4.6.

Theorem (see Theorems 4.6.1, 4.6.2)**.** *Let n, m, μ and N be as in the hidden lattice problem. Let $\delta \in (1/4, 1)$ and $\varepsilon \in (0, 1)$. There exist explicit numbers N_{I} and N_{II}, depending on $n, m, \mu, \varepsilon, \delta$, such that if $N > N_{\mathrm{I}}$ (resp. $N > N_{\mathrm{II}}$), then for every basis $\mathfrak{B}$ chosen from Ω, at least $(1-\varepsilon)\#\mathcal{H}(\mathfrak{B})$ of the hidden lattice problems constructed from $\mathcal{H}(\mathfrak{B})$ are solvable by Algorithm* I *(resp. by Algorithm* II*) by running* LLL *with reduction parameter δ.*

Variations of the HLP

Some variations of the hidden lattice problem as defined above are of interest for us. First, we study the case where given vectors lie in a small lattice modulo N only up to unknown

short "noise" vectors; we call this the *noisy hidden lattice problem* (NHLP). We notice that we can cancel the effect of the noise, by reducing the NHLP to a HLP with a "larger" (in the sense of size and dimension) hidden lattice, and apply our previous algorithms without changes. We also consider a *decisional* version (DHLP) of the hidden lattice problem, asking about the existence of a μ-small lattice $\mathcal{L}$ containing $\mathcal{M}$ modulo N. This problem, although not asking for the computation of $\overline{\mathcal{L}}$ lies at the heart of many cryptanalytic settings, and may thus be of interest to cryptanalysts. We recognize that the existence of such $\mathcal{L}$ strongly impacts the geometry of $\mathcal{M}^{\perp_N}$ (or $\mathcal{M}_N$) and, consequently, our algorithms solve the decisional version heuristically.

Applications of the HLP

Finally, we describe applications of the HLP together with some improvements implied by our Algorithm II. Our applications first show, how several computational problems all rely on the more general HLP, which can be understood as a central problem for their hardness. We describe applications to [CP19, CG20, CN19a, CNT10, BNNT11].

We first consider the approximate common divisor problem based on the Chinese Remainder Theorem (CRT-ACD). For this problem, Coron and Pereira have devised an algorithm in [CP19], which highlights a tight link to the hidden lattice problem. As second application, we consider the hidden subset sum problem from [NS99]. As already mentioned in our motivational Section 1.2.2.1, [NS99] describes a two-step algorithm similar to our Algorithm I. Recently, Coron and Gini [CG20] showed that the second step actually has exponential complexity, and propose a polynomial-time alternative. Here the first step remains unchanged and is still based on the orthogonal lattice algorithm, as in [NS99].

We also notice that the contribution described in Section 1.2.1.2 (see [CN19a]), namely our cryptanalysis of CLT13, fits in the framework of the HLP. Namely, the derived equations describe a NHLP. As pointed out in Section 1.2.1.2, due to the extra structure in the coefficients, the public vectors (i.e. the public lattice $\mathcal{M}$) can be seen to carry extra structure. In the notation of Section 1.2.1.2, one could therefore view Λ as being the *hidden* lattice, for which we computed the hidden volume. Our algorithms for NHLP thus extend the framework of the cryptanalysis of CLT13. Finally, we see that the HLP plays a central role in [CNT10, BNNT11] as well.

1.2.3 Simultaneous Diagonalization of Incomplete Matrices and Applications

1.2.3.1 A problem motivated from cryptanalysis

As already mentioned above, the cryptanalysis of CLT13 multilinear maps by Cheon et al. is an efficient attack with devastating consequences. Namely, it reveals all the secret information given access to certain public sets of encodings. Let us here in brief review the attack and raise a natural question. Let us consider for simplicity the CLT13 multilinear maps with a multilinearity degree $\kappa = 3$. Let n be the number of prime factors dividing the public modulus x_0. For the Cheon et al. attack, one considers three disjoint sets $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{C}$ of encodings, where the set $\mathcal{A} = \{\alpha_j : 1 \leq j \leq n\}$ contains n encodings of zero at level 1, the set $\mathcal{B} = \{\beta_1, \beta_2\}$ contains two encodings at level 1, and the set $\mathcal{C} = \{\gamma_k : 1 \leq k \leq n\}$ contains n encodings at level 1. At this stage, it is crucial for the elements in $\mathcal{A}$ (or in any of the sets, but for simplicity we consider them in $\mathcal{A}$) to encode zero, as then, all the product encodings of the type $\alpha\beta\gamma$

with $(\alpha, \beta, \gamma) \in \mathcal{A} \times \mathcal{B} \times \mathcal{C}$, are encodings of zero at the top-level κ, by the homomorphic multiplication property of CLT13. By public evaluations coming from zero-testing of encodings, the attacker can construct two $n \times n$ matrices W_1 and W_2 satisfying factorizations of the form $W_1 = P \cdot U_1 \cdot Q$ and $W_2 = P \cdot U_2 \cdot Q$, where P is a certain matrix of which the rows correspond to the encodings in $\mathcal{A}$, U_1 and U_2 are certain diagonal matrices of which the diagonal entries correspond to the encodings in $\mathcal{B}$, and Q is a certain matrix of which the columns correspond to the encodings in $\mathcal{C}$. Provided that at least one among W_1, W_2 is invertible over $\mathbb{Q}$ (say W_2), one can then evaluate over $\mathbb{Q}$ the matrix product:

$$W_1 \cdot W_2^{-1} = P \cdot (U_1 U_2^{-1}) \cdot P^{-1} \tag{1.8}$$

and compute the eigenvalues of $W_1 W_2^{-1}$, which reveal the diagonal entries of the product $U_1 U_2^{-1}$. From these diagonal entries, one can reveal the secret factorization of x_0 by computing greatest common divisors.

We see that the above attack requires n encodings of zero (considered to lie in $\mathcal{A}$), because the matrix P must have n rows in order to be inverted in Equation (1.8). Studying cryptanalysis of CLT13 multilinear maps along the lines of the Cheon et al. attack, when only a limited number of encodings is available to the attacker is an interesting problem. This is also true in a broader sense not specific to CLT13: algorithms with the smallest possible input length are highly desirable.

Motivated by this context, we consider the problem of decreasing the number of encodings of zero (i.e. the cardinality of $\mathcal{A}$) required for the attack. Note that this would be trivial to achieve by increasing the multilinear map degree κ. Indeed, by taking $\kappa = 4$, the set $\mathcal{A}$ could be built from products of level-1 encodings of the form $\alpha \cdot \alpha'_j$, where only a single encoding of zero α is required; then every encoding in $\mathcal{A}$ would encode zero at level 2. Therefore, we are interested in the following problem: *Is it possible to factor x_0 in polynomial time, given access to fewer encodings of zero, still with multilinearity degree $\kappa = 3$?*

Let us briefly comment on the last question. Following the above lines of attack, we can still compute the matrices $W_1 = P \cdot U_1 \cdot Q$ and $W_2 = P \cdot U_2 \cdot Q$ as above, but now the matrix P has only, say $p < n$ rows, where p is the number of encodings of zero considered in $\mathcal{A}$; therefore P is not invertible anymore and we cannot compute Equation (1.8) as above to recover the diagonal entries of $U_1 U_2^{-1}$. Note that computing linear combinations of encodings within the set $\mathcal{A}$ would not work, since the matrix P would still be of rank $p < n$. In other words, the matrix P is in some sense incomplete and some information is missing to recover the diagonal entries of $U_1 U_2^{-1}$. We can also consider the symmetric version of the problem, of having only $q < n$ encodings in the set $\mathcal{C}$, which amounts the matrix Q to have only $q < n$ columns instead of n. In both cases, the original Cheon et al. attack does not apply.

1.2.3.2 Our contribution

Motivated by the question of, for example, lowering the number of public encodings in the Cheon et al. attack, we study this problem at a more abstract level. In order to compensate the loss of rows in P, we may consider a larger number, say $t \geq 2$ of matrix factorizations. More precisely, instead of only assuming the existence of $W_1 = PU_1Q$ and $W_2 = PU_2Q$ as in the Cheon et al. attack, we start from the assumption that a sufficiently long list of matrices, say $W_a = PU_aQ$ for every $a \in \{1, \ldots, t\}$, is public, with P a matrix of size $p \times n$, $\{U_a\}_a$ diagonal $n \times n$ matrices and Q another matrix.

In more precise terms, we formulate the following computational problem from linear algebra, see Definition 5.1.1.

Definition (Problems $\mathbb{A}, \mathbb{B}, \mathbb{C}, \mathbb{D}$, see Definition 5.1.1)**.** *Let $n \geq 2, t \geq 2$ and $2 \leq p, q \leq n$ be integers. Let $\{U_a : 1 \leq a \leq t\}$ be diagonal matrices in $\mathbb{Q}^{n \times n}$. Let $\{W_a : 1 \leq a \leq t\}$ be matrices in $\mathbb{Q}^{p \times q}$ and $W_0 \in \mathbb{Q}^{p \times q}$ such that W_0 has full rank and there exist matrices $P \in \mathbb{Q}^{p \times n}$ of full rank p and $Q \in \mathbb{Q}^{n \times q}$ of full rank q, such that $W_0 = P \cdot Q$ and $W_a = P \cdot U_a \cdot Q$ for $1 \leq a \leq t$. We distinguish the following cases:*

$$(\mathbb{A}) \quad p = n \text{ and } q = n \qquad (\mathbb{B}) \quad p = n \text{ and } q < n$$
$$(\mathbb{C}) \quad p < n \text{ and } q = n \qquad (\mathbb{D}) \quad p < n \text{ and } q = p$$

In each of the four cases, the problem states as follows:

(1) Given the matrices $\{W_a : 0 \leq a \leq t\}$, compute $\{(u_{1,i}, \ldots, u_{t,i}) : 1 \leq i \leq n\}$, where for $1 \leq a \leq t$, $u_{a,1}, \ldots, u_{a,n} \in \mathbb{Q}$ are the diagonal entries of matrices $\{U_a : 1 \leq a \leq t\}$ as above.

(2) Determine whether the solution is unique.

In the above definition, we distinguish four problems that we label by $\mathbb{A}, \mathbb{B}, \mathbb{C}$ and $\mathbb{D}$, and that distinguish the sizes of the matrices P and Q with respect to n. For example, for Problem $\mathbb{A}$, the matrices P and Q have size $n \times n$, and thus Problem $\mathbb{A}$ with $t = 1$ exactly interpolates the framework of the Cheon et al. cryptanalysis against CLT13. Yet, we note a slight structural difference here. Namely, we assume that a matrix $W_0 = PQ$ is also public (we will call such a matrix a "special input"). This is not a restriction but makes calculations a little simpler. For example, Problem $\mathbb{A}$ is straightforward for any $t \geq 1$ by simultaneous diagonalization of $W_0^{-1} W_a = Q^{-1} U_a Q$ for every a. This directly reveals the tuples of diagonal entries of the matrices $\{U_a\}_a$ instead of quotients of diagonal entries, as in the Cheon et al. attack. Problems $\mathbb{B}$ and $\mathbb{C}$ applied to the CLT13 formalism then exactly correspond to the problem described above. Note that these problems are equivalent because of their symmetry in p and q: any algorithm for one solves the other upon transposing. Problem $\mathbb{D}$ considers least possible public information as both P and Q are of smaller size.

Of interest for us is to devise algorithms for Problems $\mathbb{C}$ and $\mathbb{D}$ and describe theoretical and practical values for p as a function of n. We refer to the matrices $\{W_a\}_a$ as "incomplete", as the low-rank matrices P and/or Q "steal" information. Of interest is of course the case when p is much smaller than n. For example in Problem $\mathbb{D}$, the public matrices $\{W_a\}_a$ have size $p \times p$, thus solving this problem for moderate t and p small compared to n is of interest.

Our main contribution is the description of efficient algorithms for Problems $\mathbb{C}$ and $\mathbb{D}$, and we show how to minimize the parameters p and t with respect to n. We further propose two concrete applications of our algorithms. First, we consider the multi-prime approximate common divisor problem based on Chinese Remainders, where we significantly improve on a previous algorithm by Coron and Pereira. Second, as alluded to in our motivation for these problems, we describe an improved cryptanalysis of CLT13 multilinear maps à la Cheon et al. However, we believe that our algorithms are also of independent interest and hope that more applications are to be found.

Algorithms for Problems $\mathbb{C}$ and $\mathbb{D}$

We describe practical algorithms for Problems $\mathbb{C}$ and $\mathbb{D}$, by studying these problems individually.

Our approach to Problem $\mathbb{C}$ is to use the invertibility of Q. More precisely, we can write $W_a = PU_aQ = PQQ^{-1}U_aQ = W_0Z_a$ with $Z_a = Q^{-1}U_aQ$, for every $1 \leq a \leq t$. As W_0 is not invertible, we cannot recover the matrices $\{Z_a\}_a$ directly. However, we interpret this as a system of linear equations to solve for $\{Z_a\}_a$. This system is, in general, underdetermined and does not yield the matrices $\{Z_a\}_a$ uniquely. Namely, the relation $W_a = W_0Z_a$ determines Z_a up to perturbations from elements in the kernel of W_0, which by assumption is of dimension $n - p$. That means, for every a, we can write $Z_a = Y_a + EX_a$ for a certain computable matrix Y_a, an unknown matrix X_a, and where E stands for a basis matrix for the kernel of W_0. However, noticing that the matrices $\{Z_a\}_a$ commute among each other helps us produce extra linear equations, enabling us compute such matrices $\{X_a\}_a$ efficiently. This enables to recover the matrices $\{Z_a\}_a$ uniquely, and simultaneous diagonalization eventually yields the diagonal entries of $\{U_a\}_a$. We determine exact bounds on the parameters to ensure that our system of linear equations has at least as many linear equations as variables; further, by a simple minimization problem, we justify that asymptotically this is achieved for p and t of size $O(\sqrt{n})$.

For Problem $\mathbb{D}$, we can no longer use the invertibility of Q, as it is a matrix of size $n \times p$ only. Instead, our approach is to reduce Problem $\mathbb{D}$ to Problem $\mathbb{C}$ by "augmenting" the matrix Q with well-chosen extra columns, so that it becomes invertible. Our algorithm for Problem $\mathbb{D}$ can hence be seen to function in two steps – first an augmentation technique for the matrix Q building a valid input for Problem $\mathbb{C}$, and second, running our algorithm for Problem $\mathbb{C}$ on that input. On the theoretical side, the extra columns that we augment Q with, come by considering the $n \times n$ matrix $B = QW_0^{-1}P - 1_n$, which we justify to have rank $n - p$. Any rank factorization of B, say $B = B_1B_2$ with $B_1 \in \mathbb{Q}^{n \times n-p}$ and $B_2 \in \mathbb{Q}^{n-p \times n}$, then gives rise to a suitable augmentation $Q' = [Q|B_1] \in \mathrm{GL}(n, \mathbb{Q})$ of Q. First, we prove that PB_1 is the zero matrix, which implies that $PQ' = [PQ|PB_1] = [W_0|0] =: W_0'$ is a known matrix (a special input matrix). Further, we justify that it is sufficient to compute matrices $\{V_a\}_a$ satisfying $V_a = PU_aB_1$ for every $1 \leq a \leq t$. Namely, this implies the identities $PU_aQ' = [PU_aQ|PU_aB_1] = [W_a|V_a] =: W_a'$ for every a, which are then known, and describe a valid input of Problem $\mathbb{C}$ with the augmented matrices W_0', and $\{W_a'\}_a$. The computation of such matrices $\{V_a\}_a$ is doable by standard linear algebra considerations but requires some extra information to be available beforehand and we justify that p can be set close to $(2/3)n$, which turns out to be a theoretical and practical barrier for the augmentation process. Therefore, our parameters for solving Problem $\mathbb{D}$ are a little weaker than those for Problem $\mathbb{C}$.

We provide implementations of our algorithms in SageMath and confirm our theoretical findings in practice. In summary, our first contribution is the following.

Proposition. *(i) There exists an algorithm (Algorithm 10) which on input matrices W_0 and $\{W_a\}_a$ as in Problem $\mathbb{C}$ computes the diagonal entries of the matrices $\{U_a\}_a$ under the asymptotic condition $p = O(\sqrt{n})$.*

(ii) There exists an algorithm (Algorithm 11) which on input matrices W_0 and $\{W_a\}_a$ as in Problem $\mathbb{D}$ computes the diagonal entries of the matrices $\{U_a\}_a$ under the asymptotic condition $p = O(2n/3)$.

An improved algorithm for the CRT-ACD Problem

The CRT-ACD Problem is a variant of an approximate common divisor problem based on the Chinese Remainder Theorem. We will omit a formal definition of the problem at this stage, and postpone this to Definition 2.2.7. Let $N = \prod_{i=1}^{n} p_i$ be a composite squarefree integer for certain prime numbers $\{p_i\}_i$. We should think of these primes to be rather large so that factoring N directly is intractable. Then, in simplified terms, the problem states as follows: given N and a sufficiently large set $\mathcal{S}$ of integers $\{x_s : s \in \mathcal{S}\}$ such that their reduction modulo each of the primes $\{p_i\}_i$ is "small" (which shall mean, that the unique representative in $[0, p_i) \cap \mathbb{Z}$ is somewhat small compared to p_i), factor N completely, that is, compute the prime numbers $\{p_i\}_i$. This problem is made formal by also describing the sizes of the primes $\{p_i\}_i$ and residues of the elements in $\mathcal{S}$. The choice of N and the set $\mathcal{S}$ are fixed beforehand.

In [CP19], Coron and Pereira have proposed an algorithm for the CRT-ACD Problem which works in two steps: first, a lattice-based algorithm relying on lattice reduction, and second, an algorithm which is close to the Cheon et al. attack for CLT13, as discussed above. This algorithm crucially relies on having a set $\mathcal{S}$ of size $O(n)$, where n is the number of primes dividing N. In fact, we recognize that the algorithm by Coron and Pereira relies on solving a certain instance of Problem $\mathbb{A}$, deduced from the lattice-based first step, which can be shown to succeed as soon as the size of N (or of each prime $\{p_i\}_i$) is sufficiently larger than the size of the residues of elements from $\mathcal{S}$. Thus, we contribute to an improved algorithm for the CRT-ACD Problem by modifying the input in a way that the lattice-based first step gives rise to an input of Problem $\mathbb{C}$, rather than $\mathbb{A}$. Solving this problem with our algorithm for Problem $\mathbb{C}$, allows us to squeeze the size of the public set $\mathcal{S}$ from $O(n)$ down to $O(\sqrt{n})$, which gives a quadratic improvement on the number of input samples.

Proposition. *Let $N = \prod_{i=1}^{n} p_i$ and $\mathcal{S}$ a set of CRT-ACD samples for N. There is an efficient algorithm (Algorithm 12) which factors N with $\#\mathcal{S} = O(\sqrt{n})$.*

We also illustrate our improvement by concrete parameters for instances for which N could not be factored earlier.

Improved cryptanalysis of CLT13 multilinear maps

As alluded to during the motivational part of this section, we apply our new algorithms to the Cheon et al. attack against the CLT13 scheme. We envisage to lower the number of encodings. In particular, we treat two questions: first that of lowering the number of public encodings of zero, and, that of lowering the total number of encodings. Following the description of the Cheon et al. attack above, it is nearly straightforward to see that this can be achieved by solving instances of Problem $\mathbb{C}$ and $\mathbb{D}$, instead of $\mathbb{A}$, as in the original attack. By our algorithms, we, therefore, improve the Cheon et al. attack as follows.

Proposition. *Consider the standard CLT13 notation with $x_0 = \prod_{i=1}^{n} p_i$ for prime numbers $\{p_i\}_i$ defining the encoding ring.*

(i) There exists an efficient algorithm that factors x_0 given $O(\sqrt{n})$ encodings of zero.

(ii) There exists an efficient algorithm that factors x_0 given $O(4n/3)$ encodings.

The precise algorithm description for these improvements is done in Algorithm 13. We again illustrate our algorithm by concrete parameters which could not be broken by the Cheon et al. attack.

1.2.4 Questions related to Edwards Curves

1.2.4.1 Motivation

Twisted Edwards curves over a field K are defined by the equation $ax^2 + y^2 = 1 + dx^2y^2$ for distinct elements $a \in K, d \in K \setminus \{0, 1\}$, and we denote them by $\mathcal{E}_d^a$. When $a = 1$, the curve is simply called *Edwards curve* and denoted by $\mathcal{E}_d$. Edwards curves, originally introduced by the work of Harold Edwards [Edw07], are birationally equivalent to elliptic curves. With the work of Bernstein-Lange and Bernstein et al. [BL07, BBJ+08], Edwards curves entered elliptic-curve cryptography in the early 2000's and compete still today with many other models for elliptic curves, such as Montgomery curves, Hessian curves, or Huff curves. From a purely arithmetic-geometric viewpoint, it is an interesting question to study properties of these models. This may contribute, for example, to the understanding how to construct new models for elliptic curves. In our work, we study general properties related to elliptic curves in Edwards form.

1.2.4.2 Our contribution

Abstract constructions of the Edwards and twisted Edwards models

Given an elliptic curve E defined over a field K such that $E(K)$ contains a point P of order 4, the construction by Bernstein-Lange [BL07] of an Edwards curve birationally equivalent to E is explicit. More precisely, E is assumed to be in Weierstrass form and the construction of the Edwards model, birationally equivalent to E, is explicit in the coefficients of the Weierstrass form of E and the coordinates of P. We establish this result more abstractly, by relying on the Riemann-Roch theorem. An elliptic curve over K is a pair $(E, \mathcal{O})$ where E is a non-singular curve over K of genus 1 and $\mathcal{O}$ is a K-rational point. Note that such a result would be trivial by using the construction of the Weierstrass model of $(E, \mathcal{O})$ by the Riemann-Roch theorem (see e.g. [Sil09, Chapter III, §3., Proposition 3.1]) and then invoking Bernstein-Lange's construction of the Edwards model from the Weierstrass model. We do not do this here, and our proof is independent of the Weierstrass model of $(E, \mathcal{O})$.

Theorem 1.2.1 (see Theorem 6.4.1)**.** *Let $(E, \mathcal{O})$ be an elliptic curve defined over a perfect field K of characteristic different from 2. Assume that there exists $P \in E(K)$ of order 4. Define the divisors*

$$D_1 := 2(\mathcal{O}) - 2(2P)\ ,\ D_2 := 2(\mathcal{O}) - (P) - (3P)\ \in \mathrm{Div}_K(E)$$

(i) The Riemann-Roch spaces $\mathcal{L}(D_1)$ and $\mathcal{L}(D_2)$ are one-dimensional over K.
(ii) Let $x, y \in K(E)^\times$ be generators of $\mathcal{L}(D_1)$ and $\mathcal{L}(D_2)$, respectively. The functions x^2y^2 and $x^2 + y^2 - 1$ in $K(E)^\times$ are multiples of each other, i.e. there exists $d \in K \setminus \{0, 1\}$ such that $x^2y^2 = 1 + dx^2y^2$.
(iii) The map $\phi : E(K) \to \mathbb{P}^2(K), Q \mapsto [x(Q) : y(Q) : 1]$, gives a birational equivalence from E to the Edwards curve $\mathcal{E}_d$ defined over K. In particular, x, y from (ii) are coordinate functions for the Edwards model of $(E, \mathcal{O})$.

The spaces $\mathcal{L}(D_1)$, resp. $\mathcal{L}(D_2)$ are the Riemann-Roch spaces associated with the rational divisors D_1, resp. D_2. In a sense, the divisor D_1 is associated with the 2-torsion of E, while D_2 is associated with the 4-torsion of E. By the Riemann-Roch Theorem we can show that these vector spaces have dimension 1 over K. The idea of the proof is to show that the functions

x^2y^2 and $x^2 + y^2 - 1$ lie in a 1-dimensional vector space, and are thus multiples of each other. We achieve this by explicitly computing the divisors of these functions. The map ϕ is understood as the Edwards coordinate function for E, giving a birational map between E and the Edwards model.

The twisted model of Edwards curves was introduced in [BL07]. More precisely, for squarefree $a \in K$, the curve $\mathcal{E}_d^a$ is a quadratic twist of the Edwards curve $\mathcal{E}_{d/a}$. Alternatively, twists of curves can be interpreted from a purely cohomological point of view, by using Galois cohomology. In this language, a twist of a curve $\mathcal{C}/K$ corresponds to a certain equivalence class in a cohomology group constructed from the action of the absolute Galois group on the automorphism group of $\mathcal{C}$. We recall this construction precisely in Theorem 6.4.2. When $\mathcal{C}$ is the Edwards curve $\mathcal{E}_{d/a}$, we therefore obtain a cohomological result, establishing $\mathcal{E}_d^a$ is a quadratic twist of $\mathcal{E}_{d/a}$ (we refer to Theorem 6.4.3).

Explicit construction of Edwards curves from 2-torsion points

An elliptic curve E over K with a K-rational point of order 4 is birationally equivalent to an Edwards curve $\mathcal{E}_d$ over K. The curve $\mathcal{E}_d$ and the birational map are made explicit by Bernstein-Lange [BL07, Theorem 2.1] (see also [BBJ+08, Theorem 3.3]). Not every elliptic curve has a K-rational point of order 4, thus, it is a natural question to ask how the result generalizes without this assumption. By removing this 4-torsion assumption, we are constrained to work over extensions of K. We prove the following theorem. Here, e_1, e_2, e_3 denote the x-coordinates of the $\overline{K}$-points of order 2 on E, and we use the notation $e_{ij} := e_i - e_j \in \overline{K}^\times$.

Theorem (see Theorem 6.6.5)**.** *Let E be an elliptic curve defined over a field K of characteristic not 2 and assume that $E(\overline{K})[2] \simeq \{\mathcal{O}, (e_1, 0), (e_2, 0), (e_3, 0)\}$. If $\mathcal{E}_d$ is an Edwards curve birationally equivalent to E, then*

$$d \in \left\{ \left(\frac{\sqrt{e_{12}} \pm \sqrt{e_{13}}}{\sqrt{e_{32}}} \right)^4, \left(\frac{\sqrt{e_{21}} \pm \sqrt{e_{23}}}{\sqrt{e_{31}}} \right)^4, \left(\frac{\sqrt{e_{31}} \pm \sqrt{e_{32}}}{\sqrt{e_{21}}} \right)^4 \right\} \tag{1.9}$$

Define $z_1 = e_{12}e_{13}, z_2 = e_{21}e_{23}, z_3 = e_{31}e_{32}$. The Edwards curves are defined over the extensions of K given by $K_1 := K(\sqrt{z_1})$, $K_2 := K(\sqrt{z_2})$, $K_3 := K(\sqrt{z_3})$, respectively, and birationally equivalent to E over $K(\sqrt{e_{12}}, \sqrt{e_{13}}), K(\sqrt{e_{21}}, \sqrt{e_{23}}), K(\sqrt{e_{31}}, \sqrt{e_{32}})$, respectively.

Moreover, we describe the extensions $\{K_i\}_i$ explicitly, depending on the size of $E(K)[2]$, the K-rational 2-torsion subgroup of E, (see Theorem 6.6.5). In particular, our description depends on the discriminant of E. Intuitively, we associate a pair of Edwards curves defined over the extension K_i to the 2-torsion point $(e_i, 0)$ of E, and repeat this for every $i \in \{1, 2, 3\}$. This theorem comes with several byproducts, which we now list. First, note that the formulae in Equation (1.9) do not require the knowledge of any 4-torsion point on E, which is however the case for [BL07]. This provides an easier construction of the Edwards model for E, directly from the 2-torsion.

Next, we use these formulae to refine a previous result on isomorphism classes of Edwards curves. Namely, all the Edwards curves $\mathcal{E}_d$ with d given in Equation (1.9) are isomorphic over $\overline{K}$. Isomorphisms of Edwards curves have been studied by Ahmadi and Granger

in [AG12], and we recall their main result in Proposition 6.3.5. Essentially, it says that isomorphic Edwards curves $\mathcal{E}_d$ and $\mathcal{E}_{d'}$ are related as follows:

$$d' \in \left\{ d, \frac{1}{d}, \left(\frac{1 \pm d^{1/4}}{1 \mp d^{1/4}}\right)^4, \left(\frac{1 \pm \sqrt{-1}d^{1/4}}{1 \mp \sqrt{-1}d^{1/4}}\right)^4 \right\} =: \Sigma(d) \ . \tag{1.10}$$

It is easy to deduce this result from Equation (1.9). For example, one notices that for every $i \in \{1, 2, 3\}$, two expressions within the same formula (distinguished by the sign on the numerator) are inverses of each other. Moreover, we observe that there is a natural Galois action on Edwards curves, which is compatible with that on the $\overline{K}$-points of E. For $\sigma \in G_K$, we denote by ${}^{\sigma}(\mathcal{E}_d) := \mathcal{E}_{\sigma(d)}$ the σ-conjugate Edwards curve of $\mathcal{E}_d$. By interpreting the coefficient d by Bernstein-Lange's construction functorially, i.e. as map $d : P \mapsto d_P$, we easily obtain that $d_{P^\sigma} = \sigma(d_P)$, where P^σ denotes the action of σ on the point $P \in E(\overline{K})$. We then show the following theorem, refining the result of Ahmadi and Granger (see Equation (1.10)) with the Galois-action.

Theorem (see Theorem 6.7.5)**.** *Let E be an elliptic curve over K and let $S, T \in E(\overline{K})$ be points of order 4. Let $\sigma \in \mathrm{Gal}(\overline{K}/K)$. The following hold:*

(i) if $S = \pm T^\sigma$ then $d_S = \sigma(d_T)$

(ii) if $S \neq \pm T^\sigma$ and $2S = 2T^\sigma$ then $d_S = 1/\sigma(d_T)$

(iii) if $2S \neq 2T^\sigma$ then $d_S = \left(\frac{1-\epsilon\sigma(d_T)^{1/4}}{1+\epsilon\sigma(d_T)^{1/4}}\right)^4$ for some $\epsilon \in \overline{K}$ with $\epsilon^4 = 1$.

Moreover, if $j(E) \neq 0, 1728$, then the the implications (i), (ii), (iii) are equivalences.

Finally, we describe how our formulae in Equation (1.9) relate to a modular function, considered by Edwards. Recall that in his original exposition [Edw07], Edwards introduces a modular function parametrizing Edwards curves, in a certain sense. It can be defined on the complex upper-half plane by $d(\tau) = \vartheta_2(2\tau)^4/\vartheta_3(2\tau)^4$, where $\vartheta_2(\tau) = \sum_{n\in\mathbb{Z}} q^{(n+1/2)^2}$ and $\vartheta_3(\tau) = \sum_{n\in\mathbb{Z}} q^{n^2}$ are classical theta constants derived by specializing the Jacobi theta functions at 0. This function was also considered in [KK98] and shown to be a generator for the function field of the modular curve $X_1(4)$ over $\mathbb{C}$ (see [KK98, Theorem 3]). By our formulae in Equation (1.9) we are able to obtain an expression of $d(\tau)$ in terms of the half-periods $e_1(\tau) = \wp(\omega_1/2), e_2(\tau) = \wp(\omega_2/2), e_3(\tau) = \wp((\omega_1+\omega_2)/2)$, corresponding to the Weierstrass $\wp$-function of the complex elliptic curve $\mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z})$. Such a result is for example known for the modular-λ function, which can be derived from isomorphism classes of Legendre curves $y^2 = x(x-1)(x-\lambda)$. We show for example that both functions are related via $\lambda(\tau) = d(\tau/2)$. We refer to Section 6.6.4 for the detailed results.

Rank statistics and invariants for the Edwards family

Studying the ranks of certain families of elliptic curves is of interest in many aspects. For example, Zagier and Kramarz [ZK87] study the ranks in the family $x^3 + y^3 = m$ for cubefree integers m. These curves are birationally equivalent to the elliptic curves $E_{(m)} : y^2 = x^3 - 432m^2$. Extensive computations of the rank of $E_{(m)}$ for m up to 7000 support their conjecture, claiming that curves with rank at least 2 occur with positive density. Watkins [Wat07] subsequently

extends the data from [ZK87] to $m \leq 10^7$ and shows that the density is more likely to tend to zero.

For our work, we consider the family of elliptic curves $E_{a,d}$ over the rational numbers, given by $y^2 = x^3 + 2(a+d)x^2 + (a-d)^2x$, with $a, d \in \mathbb{Z}$, $a \neq d$, and $d \neq 0, 1$. These curves are birationally equivalent to twisted Edwards curves $\mathcal{E}_d^a$, thus, we call this the *Edwards family*. Mostly, we are interested in the case $a = 1$; let $E_d := E_{1,d}$. Our first result is a complete description of the torsion subgroup of E_d, showing that most curves in the family $\{E_{1,d}\}_d$ have torsion group $\mathbb{Z}/4\mathbb{Z}$.

Theorem (see Theorem 6.9.1)**.** *Let $d \in \mathbb{Z}$ such that $d \neq 0, 1$ and let E_d be the elliptic curve defined by $y^2 = x^3 + 2(1+d)x^2 + (1-d)^2$ with $d \in \mathbb{Z}$. Then*

$$E_d(\mathbb{Q})_{\text{tors}} \simeq \begin{cases} \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z} & \text{if } d \text{ is a square} \\ \mathbb{Z}/8\mathbb{Z} & \text{if } d \text{ is not a square and } d = 1-(t^2-1)^2, \text{for some } t \in \mathbb{Z} \\ \mathbb{Z}/4\mathbb{Z} & \text{if } d \text{ is not a square and } d \neq 1-(t^2-1)^2, \text{for every } t \in \mathbb{Z} \end{cases}$$

We next derive an explicit description for invariants attached to the Edwards family, including the j-invariant of $E_{a,d}$, the discriminant and the conductor $N(E_{a,d})$ of $E_{a,d}$. When $a = 1$, we prove that $N(E_{1,d})$ can be compactly described by a relatively simple formula: up to a power of 2, $N(E_{1,d})$ is equal to the radical of $d(d-1)$. We then provide statistical data for the family of curves $\{E_{1,d}\}_d$ based on computer calculations in SageMath and Magma. While this part is in progress, we report preliminary results for 20000 curves, which we put in comparison with a bigger database for elliptic curves ([LMF20a]). Finally, we design a general simple algorithm, conditioned on the Birch and Swinnerton-Dyer Conjecture for elliptic curves, to compute the (analytic) order of the Shafarevich-Tate group for the curves in the considered family.

1.3 Roadmap of the thesis

Chapter 1 is the current section, in which we have introduced the reader to the main contributions gathered in this thesis.

In Chapter 2, we provide the reader with background material for the following chapters. We cover basics on Euclidean lattices and the arithmetic of elliptic curves, with focus on computational aspects.

In Chapter 3, we first give an overview of cryptographic multilinear maps. The second part of the chapter contains the details for our cryptanalysis against CLT13, extending a two-dimesnional lattice-attack by Gentry, Lewko and Waters.

In Chapter 4, we study the *Hidden Lattice Problem* as a generalization of a fundamental problem underlying the Nguyen-Stern algorithm for the Hidden Subset Sum Problem. We describe and compare two algorithms and discuss their relevance to more general cryptographic contexts, giving a competitive alternative to the state-of-the-art algorithms.

In Chapter 5, we study the problem of simultaneously diagonalizing certain incomplete matrices. We study this problem formally, describe algorithms for it, and discuss the theoretical and practical impact in cryptanalysis.

In Chapter 6, we study general properties of the (twisted) Edwards model of elliptic curves. We provide new explicit formulae for the construction of the Edwards model and describe properties related to the rank for related elliptic curves in a family.

CHAPTER 2

Preliminary Background

2.1 Notation

For a set $\mathcal{R}$, we denote by $\#\mathcal{R}$ its cardinality. For $r, s \in \mathbb{Z}_{\geq 1}$, we denote by $\mathcal{R}^{r\times s}$ the set of $r \times s$ matrices with entries in $\mathcal{R}$. We write 1_r for the $r \times r$ identity matrix and $0_{r\times s}$ for the zero matrix of size $r \times s$. For the latter, we drop the index $r \times s$ and simply write 0 when the size is clear from the context. For $A \in \mathcal{R}^{r\times s}$ and $B \in \mathcal{R}^{r\times s'}$, $[A|B] \in \mathcal{R}^{r\times(s+s')}$ is the augmented matrix obtained by concatenating the columns of A and B. When $\mathcal{R}$ is an associative ring with a unit 1, we denote by $\mathrm{GL}(n, \mathcal{R})$ the general linear group of $n \times n$ invertible matrices over $\mathcal{R}$ and by $\mathrm{SL}(n, \mathcal{R})$ the special linear group of $n \times n$ matrices over $\mathcal{R}$ of determinant 1.

We say that $n \in \mathbb{Z}_{\geq 1}$ has *bit size* β if $2^{\beta-1} \leq n \leq 2^\beta - 1$. For simplified asymptotic analyses we will use the notation $n \approx 2^\beta$. The same notation is used for non-rigorous approximations, such as heuristics.

We use standard Landau notation for asymptotic behaviours of real functions and complexity theroy. For two real functions f, g, we write $f(n) = O(g(n))$ if $|f(n)| \leq C|g(n)|$ for an absolute constant $C > 0$ and all sufficiently large $n \geq n_0$. We write $f(n) = \Omega(g(n))$ if $|f(n)| \geq C|g(n)|$ for an absolute constant $C > 0$ and all sufficiently large $n \geq n_0$. We write $f(n) = \Theta(g(n))$ if $C_1|g(n)| \leq f(n) \leq C_2|g(n)|$ for some absolute constants $C_1, C_2 > 0$ and all sufficiently large $n \geq n_0$. We write $f(n) = \omega(g(n))$ if for every $C > 0$, there exits n_0 such that $|f(n)| \geq C|g(n)|$ for every $n \geq n_0$.

2.2 Part I: Computational Aspects of Geometry of Numbers

This part captures background material for some techniques used in Chapters 3, 4 and 5.

Geometry of Numbers is the branch of mathematics dealing with lattices and convex bodies. We are interested in computational aspects of this theory, involving the theory of lattice reduction. For complete introductions, we refer for example to [Cas71, NV10].

2.2.1 Euclidean Lattices

For general references on Euclidean lattices, lattice reduction, and applications is cryptography, we refer the reader to the corresponding chapters in [HPS08], and [Gal12], for example.

Basic Definitions

A *Euclidean lattice*[1] is a pair $(\Lambda,\langle\cdot,\cdot\rangle)$ where Λ is a finitely generated free abelian group and $\langle\cdot,\cdot\rangle$ is a real-valued symmetric bilinear pairing $\Lambda\times\Lambda\to\mathbb{R}$. We sometimes consider *rational* lattices or *integral* lattices, in which cases the symmetric bilinear form $\langle\cdot,\cdot\rangle$ takes rational, resp. integral values. To $\langle\cdot,\cdot\rangle$ we associate the quadratic form $q:\Lambda\to\mathbb{R}$ defined by $q(x)=\langle x,x\rangle=\|x\|^2$, where $\|x\|$ is called the *norm* or *length* of x. By defining the $\mathbb{R}$-vector space $V:=\Lambda\otimes_{\mathbb{Z}}\mathbb{R}$, one views Λ as a subgroup of V, and the pairing $\langle\cdot,\cdot\rangle$ extends to a symmetric bilinear form $V\times V\to\mathbb{R}$. On V, one has all the usual notions of Euclidean geometry giving rise to a geometric interpretation to Λ.

Often and especially for computational purposes, it is more convenient to define lattices starting directly from a finite-dimensional Euclidean vector space $V\simeq\mathbb{R}^m$, equipped with the standard inner product $\langle x,y\rangle=\sum_{i=1}^m x_iy_i$ for $x=(x_i)_{1\leq i\leq m}$, $y=(y_i)_{1\leq i\leq m}\in V$. A Euclidean lattice $\Lambda\subseteq V$ is then an additive discrete subgroup of $(V,+,0)$. *Discreteness* of Λ means that for every $x\in\Lambda$, there is an open neighbourhood U such that $\Lambda\cap U=\{x\}$ (one often takes U an open disc, i.e. for every $x\in\Lambda$, there exists $\varepsilon>0$ such that the intersection of Λ and the disc $U:=B(x,\varepsilon)$ centered at x and of radius ε only contains x). Euclidean lattices are also *co-compact*, meaning that the quotient $(\Lambda\otimes\mathbb{R})/\Lambda$ is compact. This quotient is identified with the torus $(\mathbb{R}/\mathbb{Z})^n$ where $n\geq 0$ is the dimension of the $\mathbb{R}$-vector space $\Lambda\otimes\mathbb{R}$. Thereby, it follows that $\Lambda\simeq\mathbb{Z}^n$. Lattices are specified by their bases: one shows that $\Lambda\subseteq V$ is a lattice if and only if there exists a collection $\mathfrak{B}=\{b_1,\ldots,b_n\}\subseteq V$ of $\mathbb{R}$-linearly independent vectors such that

$$\Lambda=\sum_{i=1}^n\mathbb{Z}b_i=\left\{\sum_{i=1}^n a_ib_i : a_i\in\mathbb{Z}\ ,\ 1\leq i\leq n\right\}\ .$$

There is an isomorphism of abstract groups $\Lambda\simeq\mathbb{Z}^n$, and we call n the *rank* of Λ, and m the *dimension* of Λ, and one has $0\leq n\leq m$. We say that Λ has *full rank* if $n=m$. The set $\mathfrak{B}$ is a $\mathbb{Z}$*-basis* (or simply a basis) of Λ. We often use *basis matrices*, obtained by writing the basis vectors either in the rows or columns of a matrix. When $n\geq 2$, there are infinitely many bases for Λ. Two matrices $S,T\in\mathbb{R}^{n\times m}$ (where the vectors are written in rows) generate the same lattice Λ if and only if there is exists a unimodular transformation $U\in\mathrm{GL}(n,\mathbb{Z})$ such that $S=UT$. We define

$$\mathrm{Vol}(\Lambda)=\sqrt{\det(BB^T)}\ ,$$

where B is any basis matrix of Λ. For a full-rank lattice, this simplifies to $\mathrm{Vol}(\Lambda)=|\det(B)|$. This definition is independent of the choice of the basis, and we call $\mathrm{Vol}(\Lambda)$ the *volume of* Λ (or *determinant of* Λ). Geometrically, the volume of a lattice is that of one of its fundamental parallelepipeds, defined for a basis $\mathfrak{B}=\{b_1,\ldots b_n\}$ of Λ as the set

$$\mathcal{F}(\mathfrak{B})=\left\{\sum_{i=1}^n a_ib_i : a_i\in[0,1)\ ,\ 1\leq i\leq n\right\}\ .$$

The volume of Λ is the volume of $\mathcal{F}(\mathfrak{B})$. In the language above, one defines the volume of Λ more directly as follows: the quotient group $(\Lambda\otimes\mathbb{R})/\Lambda\simeq\mathbb{R}^n/\Lambda$ is a Lie group and with a canonical Λ-invariant measure; then $\mathrm{Vol}(\Lambda)$ is the total measure of $\mathbb{R}^n/\Lambda$. *Hadamard's Inequality* gives an upper bound on $\mathrm{Vol}(\Lambda)$.

[1]Throughout this thesis, we simply write *lattice* for a Euclidean lattice. We will not consider any other class of more general lattices, such as structured lattices in rings of integers of number fields.

Proposition 2.2.1 (Hadamard's Inequality)**.** *Let Λ be a lattice. For any basis $\mathfrak{B}$ of Λ, one has:*

$$\mathrm{Vol}(\Lambda) \leq \prod_{b \in \mathfrak{B}} \|b\| \ .$$

Equality holds if and only if the basis vectors are pairwise orthogonal.

We say that Λ' is a *sublattice* of Λ if $\Lambda' \subseteq \Lambda$ and Λ' is a lattice. Then Λ' has rank $n' \leq n$. If $n' = n$, we say that Λ' is a *full rank sublattice*. In this case, the index $(\Lambda : \Lambda')$ as subgroups is finite and equals $\mathrm{Vol}(\Lambda')/\mathrm{Vol}(\Lambda)$. In particular, $\mathrm{Vol}(\Lambda') \geq \mathrm{Vol}(\Lambda)$.

For $1 \leq i \leq n$, let $\lambda_i(\Lambda)$ denote the infimum of the real numbers R such that the closed ball of radius R around $0 \in \Lambda$ contains at least i linearly independent vectors of Λ. The numbers $\lambda_i(\Lambda)$, for $1 \leq i \leq n$, are called the *successive minima* of Λ. They are achieved in Λ, that is, there exist linearly independent lattice points $x_1, ...x_n \in \Lambda$ such that $\|x_i\| = \lambda_i(\Lambda)$ for all $1 \leq i \leq n$. It is clear that $\lambda_1(\Lambda) \leq \lambda_2(\Lambda) \leq \ldots \leq \lambda_n(\Lambda)$. These are important invariants describing the geometry of Λ, and understanding how large these numbers are is often very useful. Minkowski's theorems give upper bounds in terms of Hermite's constant. For an integer $n \geq 1$, *Hermite's constant* γ_n (in dimension n) is defined by $\sup \lambda_1(\Lambda)^2/\mathrm{Vol}(\Lambda)^{2/n}$, where the supremum is taken over all rank-n lattices Λ. Computing the exact value of γ_n is not easy and is only known for $1 \leq n \leq 8$ and $n = 24$. For example, $\gamma_2 = \sqrt{4/3}$. Therefore finding good approximations for γ_n is a natural question. For example, one has: $\gamma_n \leq (4/3)^{(n-1)/2} = \gamma_2^{n-1}$ (Hermite's Inequality) and $\gamma_n \leq 1 + n/4$, for every $n \geq 1$. The following upper bound is linear in n:

$$\gamma_n \leq \frac{2}{3}n \quad , \quad n \geq 2 \ . \tag{2.1}$$

Asymptotically, one has $n/(2\pi e) \leq \gamma_n \leq n/(\pi e + o(1))$ as n tends to infinity.

Minkowski's First Theorem and *Minkowski's Second Theorem* give the following upper bounds on the first minimum of Λ.

Theorem 2.2.2 (Minkowski)**.** *Let $\Lambda \subseteq \mathbb{R}^m$ be a lattice of rank n.*

(i) (Minkowski's First Theorem)

$$\lambda_1(\Lambda) \leq \sqrt{\gamma_n} \cdot \mathrm{Vol}(\Lambda)^{1/n}$$

(ii) (Minkowski's Second Theorem, [NV10, Chapter 2, Theorem 5])

$$\left(\prod_{i=1}^{r} \lambda_i(\Lambda) \right)^{1/r} \leq \sqrt{\gamma_n} \cdot \mathrm{Vol}(\Lambda)^{1/n} \ , \ 1 \leq r \leq n \ .$$

Minkowski's First Theorem is a consequence of the classical Minkowski Convex Body Theorem and a result by Blichfeldt (see e.g. Theorem 4 and Lemma 8 in [NV10, Chapter 2]). Since the successive minima can be unbalanced, one cannot have general upper bounds for the other minima separately. However, Minkowski's Second Theorem gives an upper bound for the geometric mean of the first consecutive successive minima of Λ. Note also that for $r = 1$ it gives Minkowski's First Theorem.

Dual lattices

For a lattice $\Lambda \subseteq \mathbb{Q}^m$, its *dual lattice*[2] $\Lambda^\vee$ is defined as the group $\mathrm{Hom}_{\mathbb{Z}}(\Lambda, \mathbb{Z})$ of $\mathbb{Z}$-linear maps $\Lambda \to \mathbb{Z}$. Equivalently, $\Lambda^\vee$ is identified with the set of vectors $x \in \mathbb{Q}^m$, for which the inner product with elements of Λ is an integer, that is, there is an isomorphism:

$$\{v \in \mathbb{Q}^m : \langle x, v\rangle \in \mathbb{Z},\ \forall x \in \Lambda\} \xrightarrow{\sim} \Lambda^\vee = \mathrm{Hom}_{\mathbb{Z}}(\Lambda, \mathbb{Z})\ .$$

A proof of this fact can be read off from [Con, Theorem 2.7] by taking $m = 1$, $M = \Lambda$, $R = \mathbb{Z}$ and $K = \mathbb{Q}$ therein. The argument generalizes for $m > 1$: an isomorphism is given by $v \mapsto \phi_v \in \Lambda^\vee$, where ϕ_v maps x to $\langle x, v\rangle$. In the literature, the dual lattice is often directly defined by $\{v \in \mathbb{Q}^m : \langle x, v\rangle \in \mathbb{Z},\ \forall x \in \Lambda\}$, which from a practical point of view is better suited, since elements are indeed vectors. Under this identification, there is a natural pairing $\Lambda \times \Lambda^\vee \to \mathbb{Z}$, $(v, w) \mapsto \langle v, w\rangle$.

Intuitively, dual lattices have "inverse" behaviours. For example, if $B \in \mathbb{Q}^{n\times m}$ is a basis matrix for Λ (again, with basis vectors as rows), then $B^\vee := (B^T(BB^T)^{-1})^T$ is a basis matrix for $\Lambda^\vee$, called the *dual basis* of B. For lattices of full rank, this gives $B^\vee = (B^{-1})^T$. In particular, dilating Λ by a non-zero constant k, shrinks the dual lattice by k, i.e. $(k\Lambda)^\vee = k^{-1}\Lambda^\vee$. This also implies $\mathrm{Vol}(\Lambda^\vee) = 1/\mathrm{Vol}(\Lambda)$ and $(\Lambda^\vee)^\vee = \Lambda$. The successive minima $\{\lambda_i(\Lambda^\vee)\}_i$ of $\Lambda^\vee$ are related to those of Λ by the following transference theorem due to Banaszczyk.

Theorem 2.2.3 ([Ban93], Theorem 2.1)**.** *For every lattice $\Lambda \subseteq \mathbb{R}^m$ of rank m, one has for all $1 \leq j \leq m$, the inequality*

$$1 \leq \lambda_j(\Lambda)\lambda_{m-j+1}(\Lambda^\vee) \leq m\ .$$

In cryptographic contexts one often considers, for a matrix $A \in \mathbb{Z}^{n\times m}$ and an integer q (typically a prime number), the lattices $\Lambda^{\perp_q}(A) := \{x \in \mathbb{Z}^m : Ax \equiv 0 \pmod q\}$ and $\Lambda_q(A) := \{x \in \mathbb{Z}^m : x \equiv A^T v \pmod q,\ v \in \mathbb{Z}^n\}$. They both contain $q\mathbb{Z}^m$ and are for this reason referred to as *q-ary lattices*. It follows that they have full rank m. These lattices are dual up to the scaling factor q, that is, $\Lambda_q(A) = q(\Lambda^{\perp_q}(A))^\vee$ and $\Lambda^{\perp_q}(A) = q\Lambda_q(A)^\vee$ (see e.g. [Mic11, Exercise 10]). This can be established directly, by writing down basis matrices for $\Lambda_q(A)$ and $\Lambda^{\perp_q}(A)$. Note the analogy with Definition 4.3.1 in Chapter 4.

2.2.2 Lattice Reduction

Some lattice bases are much more interesting to deal with than other bases. These bases consist of almost orthogonal and short vectors. The theory of *lattice reduction* aims at computing such bases, called *reduced bases*. There are various algorithms for lattice reduction, which differ mainly in their output quality and efficiency. To give some insight, consider the *Shortest Vector Problem* (or shortly, SVP). It asks to compute a vector of norm $\lambda_1(\Lambda)$ in a lattice Λ. This is a hard problem, and lattice reduction algorithms typically solve this problem only approximately (i.e. they solve the milder problem Approx-SVP, versus Exact-SVP). In particular, they compute a vector of norm $c\lambda_1(\Lambda)$ for $c > 1$. The closer one wishes c to be to 1, the more costly the algorithm will be. The two algorithms we will discuss below achieve c exponential in the rank of Λ in polynomial time (LLL-based algorithms), or else, smaller c in exponential time (BKZ-based algorithms).

[2] sometimes also called *polar* or *reciprocal* lattice

Lagrange-Gauss and LLL reduction

For a lattice of rank 2, one uses *Lagrange-Gauss* reduction, a 2-dimensional generalization of Euclid's algorithm for computing greatest common divisors. On input a basis $\{b_1, b_2\}$ of Λ, the algorithm returns a reduced basis $\{b_1', b_2'\}$ of Λ with $\|b_i'\| = \lambda_i(\Lambda)$ for $i = 1, 2$. In particular, it solves SVP exactly. For higher ranks, a generalization is the polynomial-time *LLL reduction* algorithm (LLL), due to Lenstra, Lenstra and Lovász in 1982, [LLL82]. We refer to [NV10] for a complete survey. We recall the definition of *LLL-reduced bases* here (see e.g. [LLL82]). For a basis $\mathfrak{B} = \{b_i : 1 \leq i \leq n\}$ of Λ, denote by $\{b_i^* : 1 \leq i \leq n\}$ the associated Gram-Schmidt orthogonalization, that is:

$$b_i^* = b_i - \sum_{1 \leq j < i} \mu_{i,j} b_j^* \quad , \quad 1 \leq i \leq n \, ,$$

where $\mu_{i,j} = \langle b_i, b_j^* \rangle / \langle b_j^*, b_j^* \rangle$ with $\langle \cdot, \cdot \rangle$ denoting the standard Euclidean inner product.

Definition 2.2.4. *Let $\mathfrak{B} = \{b_i : 1 \leq i \leq n\}$ be a basis for Λ and $\{b_i^* : 1 \leq i \leq n\}$ the associated Gram-Schmidt orthogonal basis. Let $\delta \in (1/4, 1]$. Then $\mathfrak{B}$ is said* δ-LLL reduced *if it satisfies the conditions:*

(i) (Size Condition) $|\mu_{i,j}| \leq 1/2$, for all $1 \leq j < i \leq n$.

(ii) (Lovász Condition) $\|b_i^*\|^2 \geq (\delta - \mu_{i,i-1}^2)\|b_{i-1}^*\|^2$, for all $2 \leq i \leq n$.

Note that there are equivalent ways to state the Lovász Condition. The two conditions in the definition make the resulting basis vectors reasonably short and close to being orthogonal. The reduction parameter $\delta \in (1/4, 1]$ is often set to be $3/4$, while 0.99 in practice. We refer to the algorithm producing δ-LLL reduced bases as δ-LLL. In the case $\delta = 1$, the algorithm is not guaranteed to run in polynomial time. The LLL-algorithm has many good properties and we will rely on the following theorem.

Theorem 2.2.5 (LLL)**.** *Let Λ be a lattice in $\mathbb{R}^m$ of positive rank n. Let $\{b_i : 1 \leq i \leq n\}$ be a basis of Λ. Let $\delta \in (1/4, 1)$. The* δ-LLL *algorithm computes a* δ-LLL *reduced basis $\{b_i' : 1 \leq i \leq n\}$ of Λ. In particular, the vectors $\{b_i' : 1 \leq i \leq n\}$ satisfy*

$$\|b_j'\| \leq c^{(n-1)/2} \lambda_i(\Lambda)$$

for all $1 \leq j \leq i \leq n$, and where $c = 1/(\delta - 1/4)$.

Note that when $i = j = 1$ in Theorem 2.2.5, one obtains that

$$\|b_1'\| \leq c^{(n-1)/2} \|x\| \quad , \quad \forall \, x \in \Lambda \setminus \{0\} \, . \tag{2.2}$$

We will always use the LLL algorithm on integer lattices, i.e. contained in $\mathbb{Z}^m$. In this case, the following formulae are used to bound the complexity of LLL. Let $\{b_i\}_i$ be a basis of $\Lambda \subseteq \mathbb{Z}^m$ with vectors of norm at most $X \in \mathbb{Z}_{\geq 2}$. By [Gal12, Corollary 17.5.4], LLL computes, on input $\{b_i\}_i$, a reduced basis of Λ, in $O(n^5 m \log(X)^3)$ bit operations. In [NS09], using the L^2-variant of the LLL algorithm, the complexity was improved to

$$O(n^4 m (n + \log(X)) \log(X)) \, , \tag{2.3}$$

which is only quadratic in $\log(X)$ (hence the name L^2) and relies on naive integer multiplication. We also refer to [NSV11] for a variant of LLL with complexity quasi-linear in $\log(X)$.

BKZ Reduction

Another frequently used lattice reduction algorithm is the *BKZ algorithm*, a block-variant of LLL, designed upon the notion of *block Korkin-Zolotarev reduced bases*, [Sch87, SE94]. See also [GHGKN06b] and [GN08a]. This type of reduction is generally much stronger than LLL-reduction. Besides a reduction parameter, the BKZ-algorithm also uses a block-size $2 \leq \beta \leq n$ (for a lattice of rank n). After LLL-reducing the input basis, it relies on an enumeration subroutine which iteratively finds a shortest vector in certain local projected lattices (thus solving an Exact-SVP Problem). Therefore the running time is exponential in n. In general, a block-size $\beta \leq 25$ is very efficient, but beyond that barrier, the subroutine becomes less practical and the running time increases in large dimensions. In [CN11], the authors describe BKZ 2.0 based on several optimizations, allowing to run higher block-sizes. In general, the running time remains exponential in n. For a fixed block-size and when restricted to polynomially many iterations, the running time of BKZ is polynomial time, and the output quality of the reduced basis can be guaranteed as shown in [LN20]. We will for this thesis rely on a heuristic complexity, detailed in the subsequent paragraph. However, it is known that lattice reduction behaves much better in practice than what theory predicts, [GN08b].

Heuristics

We will sometimes rely on heuristic arguments, such as heuristic analyses for our algorithms for lattices. Whenever we make a heuristic analysis later, we assume that a lattice reduction algorithm outputs a basis $\{b'_i\}_i$ of Λ with

$$\|b'_i\| \leq \iota^n \lambda_i(\Lambda) \quad , \quad 1 \leq i \leq n \, , \tag{2.4}$$

where $\iota > 1$ is the *root Hermite factor* depending on the reduction algorithm[3]. By Theorem 2.2.5, we have $\iota^n = c^{(n-1)/2}$ for the LLL algorithm, and $\iota^n = 1/2(\gamma_\beta)^{\frac{n-1}{\beta-1}}(i+3)^{1/2}$ for the BKZ algorithm with block-size β (see [Sch87]).

Heuristically, we can bound the complexity of BKZ from below, by means of an upper bound on ι. Namely, a root Hermite factor ι is (heuristically) achieved within time at least $2^{\Theta(1/\log(\iota))}$ by using BKZ with block-size $\beta = \Theta(1/\log(\iota))$, see [HPS11a].

For a full-rank lattice $\Lambda \subseteq \mathbb{R}^n$ and a measurable body $\mathcal{K} \subseteq \mathbb{R}^n$, the *Gaussian Heuristic* (see e.g. [HPS08, Section 6.5.3]) predicts that $\#(\Lambda \cap \mathcal{K})$ is approximately $\mathrm{Vol}(\mathcal{K})/\mathrm{Vol}(\Lambda)$. When applied to Euclidean balls, one therefore heuristically approximates $\lambda_1(\Lambda)$ by

$$\sqrt{\frac{n}{2\pi e}} \cdot \mathrm{Vol}(\Lambda)^{1/n} \, .$$

For a "random" lattice Λ, this is proven with overwhelming probability, [Ajt06]. Although there exists a precise mathematical formulation of "random" lattices using the Haar measure, we will not enter in this discussion. For a precise setting, we refer to [Ajt06]. As noticed in [MO90], the number of integral points in high-dimensional spheres strongly depends on the location of its center, therefore the heuristic can be rejected in some cases. When considering

[3]Sometimes we also write the condition as $\|b'_i\| \leq 2^{\iota n} \lambda_i(\Lambda)$, in which case the root Hermite factor is 2^ι

"random" lattices, we will heuristically assume all the minima to be approximately equal, that is:

$$\lambda_k(\Lambda) \approx \sqrt{\frac{n}{2\pi e}} \cdot \mathrm{Vol}(\Lambda)^{1/n} \ , \ 1 \leq k \leq n \ . \tag{2.5}$$

We will sometimes also rely on Equation (4.32) when Λ is not of full rank. Moreover, for an even simpler analysis, we sometimes drop the factor in front of $\mathrm{Vol}(\Lambda)^{1/n}$.

2.2.3 Computational Problems of cryptographic interest

2.2.3.1 The Hidden Subset Sum Problem

For an integer $n \geq 1$, the *classical subset sum problem* [LO85] asks, given n integers $\alpha_1, \ldots, \alpha_n$ and a target sum s, to compute weights $x_1, \ldots, x_n \in \{0,1\}$ such that $\sum_{i=1}^{n} \alpha_i x_i = s$. It is customary to study this problem with a modulus N, which is public. To attack the problem, it is common to turn it into a (shortest vector) lattice problem, by constructing a certain lattice depending on $\{\alpha_i\}_i$ and s (and N if available). Computing $\{x_i\}_i$ then corresponds to the computation of a somewhat short vector in that lattice (see e.g. [LO85]). The *hidden subset sum problem* (HSSP) is a vector variant of the classical subset sum problem and states as follows (see [NS99, CG20]):

Definition 2.2.6 (Hidden subset sum problem, HSSP)**.** *Let $n, m \in \mathbb{Z}_{\geq 1}$ with $n \leq m$, and $N \in \mathbb{Z}$. Let $v \in \mathbb{Z}^m$ be such that $v \equiv \sum_{i=1}^{n} \alpha_i x_i \pmod N$ with $\alpha_i \in \mathbb{Z}$ and $x_i \in \{0,1\}^m$, for all $1 \leq i \leq n$.*

The HSSP *states as follows: given v and N, compute vectors $\{x_i\}_i \in \{0,1\}^m$ and integers $\{\alpha_i\}_i$ such that $v \equiv \sum_{i=1}^{n} \alpha_i x_i \pmod N$.*

In contrast to the classical subset sum problem, the weights $\{\alpha_i\}_i$ are hidden. In small dimensions, the problem is not uniquely solvable, but when n, m are large, the solution is likely unique. The origins of this problem lie in protocols based on the discrete logarithm problem, such as a fast random generation method by Boyko, Peinado and Venkatesan [BPV98]. The security of this generator depends on the HSSP, which was formally studied by Nguyen and Stern in 1999, [NS99]. They provide a lattice-based algorithm for the hidden subset sum problem, relying on the *orthogonal lattice*, introduced as a strong tool for cryptanalysis in [NS97]. In [CG20], Coron and Gini notice that the algorithm in [NS99] has exponential time complexity in large dimensions, and provide a polynomial-time variant.

We consider the HSSP mostly in Chapter 4, in the context of the Hidden Lattice Problem, but already mention it in Chapter 3 as analogy to our problem of study.

2.2.3.2 Approximate Common Divisor Problems

The goal of *approximate common divisor problems* (ACD) is to reveal a secret (prime) integer p, not from *exact multiples of p* (in which case the problem would be easy), but "approximate" multiples of p, that is, integers of the form $q_i p + r_i$ (for $i \in \mathcal{S}$, some finite set of indices), where $\{r_i\}_i$ are somewhat "small" integers, referred to as "noise". Formally, the problem is be stated including parameters describing the sizes of the involved integers.

Approximate common divisor problems have gained a lot of interest in cryptography and were first used to build a fully homomorphic encryption scheme [VDGHV10]. In [CH13], Cohn and Heninger study generalizations of the approximate common divisor problem via

lattices. The paper [GGM16] surveys and compares algorithms for the ACD Problem based on lattices. Namely, many lattice attacks against the ACD Problem have been introduced [CN12, CNT12, DT14] after that a first lattice attack was studied in [VDGHV10].

The CRT-ACD Problem

In this work, we consider the multi-prime version (CRT-ACD Problem) from [CP19], interpreted as a factorization problem with constraints based on the Chinese Remainder Theorem (CRT). We formally state the CRT-ACD Problem as follows (see e.g. [CP19, Definition 3]):

Definition 2.2.7 (CRT-ACD Problem)**.** *Let $n, \eta, \rho \in \mathbb{Z}_{\geq 1}$. Let $p_1, \ldots, p_n$ be distinct η-bit prime numbers and $N = \prod_{i=1}^{n} p_i$. Consider a non-empty finite set $\mathcal{S}$ of integers in $\mathbb{Z} \cap [0, N)$ such that for every $x \in \mathcal{S}$:*

$$x \equiv x_i \pmod{p_i}, \quad 1 \leq i \leq n$$

integers $x_i \in \mathbb{Z}$ satisfying $|x_i| \leq 2^{\rho}$.

The CRT-ACD problem *states as follows: given the set $\mathcal{S}$, the integers η, ρ and N, factor N completely, i.e. reveal the prime numbers $p_1, \ldots, p_n$.*

To keep track of the parameters, one also refers to the problem as (γ, η, ρ)-CRT-ACD. In the literature, many definitions also mention probability distributions according to which the integers are chosen. For ACD Problems, this is typically the uniform distribution. An instance of the (γ, η, ρ)-CRT-ACD Problem is thus created as follows: one chooses $\{x_1, \ldots, x_n\}$ uniformly distributed in $[-2^{\rho}, 2^{\rho}] \cap \mathbb{Z}$, and for every vector $(x_1, \ldots, x_n)$ one computes, by the Chinese Remainder Theorem, $x \in \mathbb{Z} \cap [0, N)$ such that $x \equiv x_i \pmod{p_i}$ for every $1 \leq i \leq n$. Repeating this for many vectors $(x_1, \ldots, x_n)$ gives rise to the set $\mathcal{S}$ in Definition 2.2.7, consisting of the CRT-representations x associated to $(x_1, \ldots, x_n)$.

The size η of the primes $\{p_i\}_i$ is large, so that direct factorization of N is (classically) intractable. The size of ρ is typically much smaller. Also, the larger the set $\mathcal{S}$ is, the more information is available and the easier the problem is.

The algorithm of Coron and Pereira [CP19]

Coron and Pereira propose an algorithm for the CRT-ACD Problem for the case $\#\mathcal{S} = n + 1$. Their algorithm proceeds in two steps, which are referred to as, the "orthogonal lattice attack" following [NS99] and the "algebraic attack" following [CHL$^+$15]. We briefly review their algorithm and refer to [CP19, Section 4.3] for the original description.

Let $\mathcal{S} = \{x_1, \ldots, x_n, y\}$ and $x = (x_1, \ldots, x_n) \in \mathcal{S}^n$. Then, the vector $b = (x, y \cdot x) \in \mathbb{Z}^{2n}$ is public, and by the Chinese Remainder Theorem, letting $x \equiv x^{(i)} \pmod{p_i}$ and $y \equiv y^{(i)} \pmod{p_i}$ for all $1 \leq i \leq n$, one has

$$b \equiv \sum_{i=1}^{n} c_i (x^{(i)}, y^{(i)} x^{(i)}) =: \sum_{i=1}^{n} c_i b^{(i)} \pmod{N}$$

for some integers $c_1, \ldots, c_n$. If the vectors $\{x^{(i)}\}_i$ are $\mathbb{R}$-linearly independent, then so are $\{b^{(i)}\}_i$ and generate a $2n$-dimensional lattice $\mathcal{L}$ of rank n. Importantly, by Definition 2.2.7, the vectors $\{b^{(i)}\}_i$ are reasonably short vectors, of ℓ_2-norm approximately $2^{2\rho}$, where ρ is considered much smaller than η.

Step 1: Orthogonal lattice attack. The first step of the algorithm is, in rough terms, to reveal (a basis of) the lattice $\mathcal{L}$ from the knowledge of b and N. This is precisely what the "orthogonal lattice attack" from [NS99] aims at: on input b and N, it computes a basis of the completion $\overline{\mathcal{L}} = \mathcal{L}_{\mathbb{Q}} \cap \mathbb{Z}^{2n}$ of $\mathcal{L}$, where $\mathcal{L}_{\mathbb{Q}}$ denotes the $\mathbb{Q}$-span of $\mathcal{L}$, that is $\sum_{i=1}^{n} \mathbb{Q}b^{(i)}$. To achieve this, one performs lattice reduction on the lattice $(\mathbb{Z}b)^{\perp_N}$ of vectors $v \in \mathbb{Z}^{2n}$ such that $\langle v, b\rangle \equiv 0 \pmod N$. The lattice $\mathcal{L}^{\perp}$ is a sublattice thereof. Lattice reduction on $(\mathbb{Z}b)^{\perp_N}$ reveals a sublattice of $\mathcal{L}^{\perp} \subseteq (\mathbb{Z}b)^{\perp_N}$, from which the completion of $\mathcal{L}$ is revealed by computing the orthogonal complement. The parameters are chosen accordingly, and one essentially requires the condition $2\rho < \eta$.

Step 2: Algebraic attack. Upon finding a basis $\{b'^{(i)}\}_i$ of $\overline{\mathcal{L}}$, the authors proceed with techniques from linear algebra. The idea is to imitate the attack of Cheon et al. against the CLT13 multilinear map scheme (see [CHL+15] and Chapter 3), by computing the eigenvalues of a well-chosen matrix. More precisely, let us denote by $\{b'^{(i)}\}_i$ the basis of $\overline{\mathcal{L}}$ computed in Step 1. This basis is related to the basis $\{b^{(i)}\}_i$ of $\mathcal{L}$ via an (unknown) invertible base change matrix $Q \in \mathbb{Q}^{n\times n}$. We denote the basis matrix of $\overline{\mathcal{L}}$ computed by the orthogonal lattice attack as $[W_0|W_1] \in \mathbb{Z}^{n\times 2n}$, with blocks $W_0 \in \mathbb{Z}^{n\times n}$ and $W_1 \in \mathbb{Z}^{n\times n}$. By the design of the vectors $\{b^{(i)}\}_i$, defined by $b^{(i)} = (x^{(i)}, y^{(i)}x^{(i)})$, one has by letting $P = [x^{(1)}|\dots|x^{(n)}] \in \mathbb{Z}^{n\times n}$ with columns $\{x^{(i)}\}_i$, the following matrix relations

$$W_0 = P \cdot Q \ , \ W_1 = P \cdot U_1 \cdot Q \tag{2.6}$$

where U_1 is $n \times n$ diagonal with entries $\{y^{(i)}\}_i$. The matrix W_0 is invertible (over $\mathbb{Q}$) and one computes the eigenvalues $\{y^{(i)}\}_i$ of

$$W_1 W_0^{-1} = P U_1 P^{-1} \, .$$

Using $y \equiv y^{(i)} \pmod{p_i}$ for every $1 \le i \le n$, one factors N by the gcd-computations $\gcd(y - y^{(i)}, N)$ for all i.

Algorithm. In summary, the algorithm is as follows.

Algorithm 1 Algorithm for the CRT-ACD Problem [CP19] with $\#\mathcal{S} = n + 1$

Parameters: The CRT-ACD parameters (Definition 2.2.7)
Input: An integer $N = \prod_{i=1}^{n} p_i$ and a set $\mathcal{S}$ as in Definition 2.2.7 with $\#\mathcal{S} = n + 1$
Output: The prime factors $\{p_i : 1 \le i \le n\}$ of N

1: Write the elements in $\mathcal{S}$ as $x_1, \dots, x_n, y$ and let $x = (x_1, \dots, x_n)$. Construct the vector $b = (x, y \cdot x) \in \mathbb{Z}^{2n}$ from $\mathcal{S}$
2: Run the "orthogonal lattice attack" on b and N and denote by $[W_0|W_1] \in \mathbb{Z}^{n\times 2n}$ the computed basis of the lattice $\overline{\mathcal{L}}$, where $\mathcal{L}$ is the lattice generated by $\{b^{(i)} : 1 \le i \le n\}$
3: Compute the eigenvalues $\{y^{(i)}\}_i$ of $W_1 W_0^{-1}$
4: **for** $1 \le i \le n$ **do**
5: Compute and return $\gcd(y - y^{(i)}, N)$
6: **end for**

In Chapter 4, we consider the CRT-ACD Problem as an application of the Hidden Lattice Problem and will improve the first step of the algorithm of [CP19]. In particular, we design a

new algorithm to reveal the lattice $\overline{\mathcal{L}}$ from b and N. In Chapter 5, we will again consider the CRT-ACD Problem, and mainly improve the second step of the algorithm of [CP19]. Together with a meaningful rewriting of the problem statement, we propose an efficient algorithm working with $\#\mathcal{S} = O(\sqrt{n})$, instead of $\#\mathcal{S} = O(n)$.

2.3 Part II: Computational Aspects of Elliptic Curves

In this part, we gather some background on elliptic curves, which we consider in Chapter 6. The theory of elliptic curves occupies a fundamental place in number theory and cryptography. We refer to [Sil09, Sil94, Hus04, Was03] for complete introductions.

2.3.1 Basic Definitions

An *elliptic curve* E defined over a field K is a pair $(E, \mathcal{O})$ where E is a non-singular projective curve over K of genus 1 and $\mathcal{O}$ a K-rational point on E. The point $\mathcal{O}$ is often omitted and we write E/K to say that E is defined over K. A rational function on E is a map $\phi : E(K) \to \mathbb{P}^1(K)$ and the set (which is a field) of rational functions is denoted by $K(E)$, called the *function field of* E. One says that $P \in E(K)$ is a pole of ϕ if $\phi(P) = \infty$. Otherwise, P is a regular point and its image under ϕ is $[\phi(P) : 1]$. Using the Riemann-Roch theorem for curves (see e.g. [Sil09, Chapter II, §5, Theorem 5.4]), one shows that there exist rational functions x, y on E inducing the (affine) long *Weierstrass equation* for E:

$$y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6 \ ,$$

with $a_1, a_2, a_3, a_4, a_6 \in K$. Since elliptic curves are projective, we will always implicitly mean the projective closure of the above equation, given by the equation $Y^2Z + a_1XYZ + a_3YZ^2 = X^3 + a_2X^2Z + a_4XZ^2 + a_6Z^3$, with $[X : Y : Z] \in \mathbb{P}^2(K)$. Then $\mathcal{O} = [0 : 1 : 0]$ is the unique point with $Z = 0$, the *point at infinity*. To pass to affine coordinates one sets $(x, y) = (X/Z, Y/Z)$ with $Z \neq 0$. For simplicity, we mostly work with affine models.

The set of K-points $(x, y) \in K^2$ satisfying the Weierstrass equation of E, together with $\mathcal{O}$, is denoted by $E(K)$. It carries the structure of an abelian group with neutral element $\mathcal{O}$. The group law is written additively and we denote it by $+$. There are explicit formulas for the addition law in terms of the Weierstass equation. We refer to the Weierstrass equation as a *model* for E; there are other models. In many cryptographic applications, one uses the *Montgomery* or *Edwards* model, [Mon87, Edw07]. In Chapter 6 we study properties of the Edwards model and more background will be provided therein. When $\mathrm{char}(K) \neq 2, 3$, E can be given by a *short Weierstrass equation* $y^2 = x^3 + a_4x + a_6$.

We often work with $K = \mathbb{Q}$. By Mordell's Theorem, $E(\mathbb{Q})$ is finitely generated, therefore, there exists a non-negative integer $r = r(E)$ and a finite group $T = T(E)$ such that $E(\mathbb{Q}) \simeq \mathbb{Z}^r \oplus T$. We call r the *algebraic rank of* E or the *Mordell-Weil rank*; it is the $\mathbb{Z}$-rank of $E(\mathbb{Q})$, also denoted by $\mathrm{rk}_{\mathbb{Z}}(E(\mathbb{Q}))$. The group $T =: E(\mathbb{Q})_{\mathrm{tors}}$ is the *torsion subgroup of* E, consisting of points of finite order, that is, $\cup_{m \geq 1} E[m]$, where $E[m]$ denotes the m-torsion subgroup of $E(\mathbb{Q})$, containing the points P such that $mP = \mathcal{O}$, the points of *order* m. Computing the rank of E is generally a more involved problem and is related to the conjecture of Birch and Swinnerton-Dyer, which we discuss below. Mazur's Theorem classifies possible torsion subgroups for E over $\mathbb{Q}$: $E(\mathbb{Q})_{\mathrm{tors}} \simeq \mathbb{Z}/n\mathbb{Z}$ for $n = 1, \ldots, 10$ or $n = 12$, or $E(\mathbb{Q})_{\mathrm{tors}} \simeq \mathbb{Z}/2n\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ for $n = 1, \ldots, 4$.

Invariants of elliptic curves. The *discriminant* $\Delta(E) \in K^\times$ of E is essentially the discriminant of the cubic polynomial in x and the property of E being non-singular is equivalent to $\Delta(E)$ being non-zero. The *j-invariant* $j(E)$ of E is defined by $c_4^3/\Delta(E)$, where c_4 is a classical constant, explicit in the coefficients $\{a_i\}_i$ of the Weierstrass equation. The j-invariant classifies isomorphism classes of elliptic curves over $\overline{K}$. Elliptic curves over K can be isomorphic over $\overline{K}$ without being isomorphic over K. In this case, they are *twists* of each other. In particular, we say that E/K is a *quadratic twist* of E'/K if there exists an isomorphism $E \to E'$ defined over a quadratic extension of K and E and E' are non-isomorphic over K.

Let $K = \mathbb{Q}$. If the coefficients defining the Weierstrass form of E are integral, one can reduce them modulo a prime p. The resulting curve is an elliptic curve over $\mathbb{F}_p$ if and only if $p \nmid \Delta(E)$, in which case p is called a prime of *good reduction*. Otherwise, the reduced curve is singular and p is called a prime of *bad reduction*. The primes of bad reduction are therefore finitely many. The behaviour of the reduction of E at bad primes is measured by the *conductor of* E, denoted by $N(E)$. It can be computed by $N(E) = \prod_{p|\Delta'(E)} p^{f_p(E)}$ where $\{f_p(E)\}_p$ are certain integer exponents, described explicitly in terms of *Kodaira symbols*, classifying the possible reduction types of E over local fields. Here, $\Delta'(E)$ denotes the *minimal discriminant of* E, which can be computed efficiently using an algorithm of Tate (see [Sil94, Chapter IV, §9]). More details are found in [Sil94, Chapter IV]. The conductor appears in many formulas, as for example, in the context of modular forms and L-functions for elliptic curves.

2.3.2 L-functions of elliptic curves

L-functions of elliptic curves are generating functions recording information about the reduction of the elliptic curve modulo every prime number, and allow to produce global information about $E(\mathbb{Q})$.

Let $E/\mathbb{Q}$ be an elliptic curve given by an integral Weierstrass equation. For a prime number p of good reduction, let[4] $a_p := a_p(E) = p + 1 - \#E(\mathbb{F}_p)$, where $\#E(\mathbb{F}_p)$ is the number of points on the modulo-p-reduction of E. By Hasse's Theorem, $|a_p| \leq 2\sqrt{p}$. Schoof's point counting algorithm [Sch85] efficiently computes $\#E(\mathbb{F}_p)$, thus a_p is efficiently computable. For primes p of bad reduction, we distinguish the cases of *additive, split multiplicative* and *non-split multiplicative* reduction, and define:

$$a_p := a_p(E) = \begin{cases} 0 & E \text{ has additive reduction at } p \\ 1 & E \text{ has split multiplicative reduction at } p \\ -1 & E \text{ has non split multiplicative reduction at } p \end{cases}.$$

We define the *Hasse-Weil L-series* (or L-function) of E by the following *Euler product*, defined for $s \in \mathbb{C}$:

$$L(E,s) = \prod_{p|\Delta(E)} (1 - a_p p^{-s})^{-1} \cdot \prod_{p\nmid\Delta(E)} (1 - a_p p^{-s} + p^{1-2s})^{-1} \tag{2.7}$$

Using Hasse's Theorem, one shows that $L(E,s)$ converges for $\Re(s) \geq 3/2$. The modularity theorem [TW95, Wil95] shows that it has meromorphic continuation to all of $\mathbb{C}$. In particular, the function is defined at $s = 1$. This is the important case in the Birch and Swinnerton-Dyer

[4] This notation is not to confuse with the labeling of the coefficients of the Weierstrass equation.

Conjecture. One can show that $L(E,s)$ can be written as

$$L(E,s) = \sum_{n\geq 1} \frac{a_n}{n^s} \ ,$$

where $a_n := a_n(E)$ is defined by the following arithmetic function:

$$\begin{cases} a_1 = 1 \\ a_{p^k} = a_p a_{p^{k-1}} - \chi(p) p a_{p^{k-2}} \quad \text{if } p \text{ is prime and } \ k \geq 2 \end{cases} , \tag{2.8}$$

where $\chi(p) = 1$ if $p \nmid N(E)$ and $\chi(p) = 0$ if $p \mid N(E)$. For composite $n = \prod_{i=1}^{s} p_i^{k_i}$, we extend a_n multiplycatively and set $a_n = \prod_{i=1}^{s} a_{p^{k_i}}$.

2.3.3 The Birch and Swinnerton-Dyer Conjecture

Birch and Swinnerton-Dyer [BSD65] formulated a conjecture about the algebraic rank of E, by extensive computer calculations (which is quite impressive given the time it dates back). Today it is one of the most important open problems in number theory and is part of the Millenium Problems [Wil06]. In modern language, it can be stated in two parts. The first part gives a conjectural formula for the algebraic rank of E in terms of the L-series of E. The second part conjectures an explicit formula depending on other invariants of E, that we introduce briefly below:

- One denotes by Ω_E the *real period of* E, defined by $\int_{\mathbb{R}} |\omega(E)|$, where $\omega(E) = dx/(2y + a_1 x + a_3)$ is the invariant differential on E (given by the usual Weierstrass equation).

- One defines the *regulator* R_E *of* E as follows. By the Mordell-Weil Theorem, the group $E(\mathbb{Q})/E(\mathbb{Q})_{\text{tors}}$ is isomorphic to $\mathbb{Z}^r$, where r is the rank of E. It is called the *Mordell-Weil lattice*. It is equipped with the Néron-Tate height pairing $\langle \cdot, \cdot \rangle$ on $E(\mathbb{Q})$ defined by $\langle P, Q \rangle = \frac{1}{2}(\hat{h}(P+Q) - \hat{h}(P) - \hat{h}(Q))$, for points P and Q, and where $\hat{h}$ is the canonical height function on E. The regulator of E is the volume[5] of the Mordell-Weil lattice $E(\mathbb{Q})/E(\mathbb{Q})_{\text{tors}}$. Denoting by $P_1, \ldots, P_r$ a $\mathbb{Z}$-basis of $E(\mathbb{Q})/E(\mathbb{Q})_{\text{tors}}$, it is equal to $R_E = |\det(A)|$ where $A = (\langle P_i, P_j \rangle)_{i,j=1,\ldots,r}$. Note that when $r = 0$ then $R_E = 1$, and when $r = 1$ then $R_E = \hat{h}(P)$ where P is a generator for $E(\mathbb{Q})/E(\mathbb{Q})_{\text{tors}}$.

- One defines the *global Tamagawa number* c_E *of* E as follows. For a prime p, let $E_0(\mathbb{Q}_p)$ denote the subgroup of $E(\mathbb{Q}_p)$ of points such that their reduction in $E(\mathbb{F}_p)$ is non-singular. The *local Tamagawa number at* p is defined by the integer $c_{E,p} = \#(E(\mathbb{Q}_p)/E_0(\mathbb{Q}_p))$. In particular, if E has good reduction at p, then $c_{E,p} = 1$. The global Tamagawa number is defined by $c_E = \prod_{p|\Delta(E)} c_{E,p}$.

Conjecture 2.3.1 (Birch and Swinnerton-Dyer, BSD Conjecture)**.** *Let E be an elliptic curve defined over $\mathbb{Q}$.*

(i) The order of vanishing $\text{ord}_{s=1} L(E,s)$ *of* $L(E,s)$ *at* $s = 1$ *is equal to* $\text{rk}_{\mathbb{Z}}(E(\mathbb{Q}))$*, the* algebraic *rank of* E*. The quantity* $\text{ord}_{s=1} L(E,s)$ *is called the* analytic rank *of* E*.*

[5] This is defined in exactly the same way as in Section 2.2.1.

(ii) The Shafarevich-Tate group $\text{Ш}(E)$ *is finite and*

$$\frac{L^{(r)}(E,1)}{r!} = \frac{\#\text{Ш}(E)\cdot\Omega_E\cdot R_E\cdot c_E}{(\#E(\mathbb{Q})_{\text{tors}})^2}\,, \tag{2.9}$$

where $r := \text{ord}_{s=1}L(E,s)$ *and* $L^{(r)}(E,1)$ *denotes the* r*th derivative of* $L(E,s)$ *evaluated at* $s=1$*.*

We will introduce the Shafarevich-Tate group of E in the following section. part *(i)* of the conjecture is often referred to as the *weak* BSD Conjecture. It says that the leading term in the Taylor expansion around $s=1$ of $L(E,s)$ is of the form $c(s-1)^r$, where $r=\text{rk}_{\mathbb{Z}}(E(\mathbb{Q}))$ and c is a non-zero constant. The constant c is described in part *(ii)*. In Section 6.9 of Chapter 6 we will rely on Conjecture 2.3.1 to approximate the order of the Shafarevich-Tate group of E, based on a truncation method for the L-series of E.

Part of the BSD Conjecture is known in special cases. Kolyvagin [Kol88] showed that if $\text{ord}_{s=1}L(E,s)\leq 1$ then point (i) of BSD Conjecture holds and $\text{Ш}(E)$ is finite. Bhargava, Skinner and Zhang [BSZ14] proved that point (i) of the BSD Conjecture holds for at least 66% of elliptic curves over $\mathbb{Q}$, ordered by height.

2.3.4 Shafarevich-Tate group, Selmer group and isogeny-descent

We now define the Shafarevich-Tate group and Selmer group of E, following [Sil09, Chapter X]. These groups are often encountered when computing rational points on E, and appear in the proof of the weak Mordell-Weil theorem, saying that $E(\mathbb{Q})/mE(\mathbb{Q})$ is finite for every integer m.

Shafarevich-Tate group. Let E be an elliptic curve over $\mathbb{Q}$. A *principal homogeneous space* for E is a pair (C,μ) where $C/\mathbb{Q}$ is a smooth curve and $\mu : C(\mathbb{Q})\times E(\mathbb{Q})\to C(\mathbb{Q})$ is a free and transitive group action of $E(\mathbb{Q})$ on $C(\mathbb{Q})$. Two principal homogenous spaces (C,μ) and (C',μ') are equivalent if there is a $\mathbb{Q}$-isomorphism $C\to C'$ which is compatible with the group action of $E(\mathbb{Q})$ on $C(\mathbb{Q})$ and $C'(\mathbb{Q})$. For the sequel, we just write C instead of (C,μ). The set of principal homogeneous spaces for $E/\mathbb{Q}$ up to equivalence is called the *Weil-Châtelet group* of $E/\mathbb{Q}$, denoted by $\text{WC}(E/\mathbb{Q})$. The group $E(\mathbb{Q})$ acts on itself, giving up to equivalence, the trivial class $[E]$ inside $\text{WC}(E/\mathbb{Q})$. A criterion to know which homogeneous spaces are trivial is given in [Sil09, Chapter X, Proposition 3.3]: $[C]$ is trivial in $\text{WC}(E/\mathbb{Q})$ (i.e. $[C]=[E]$) if and only if $C(\mathbb{Q})\neq\emptyset$. Let $G_{\mathbb{Q}}=\text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ be the absolute Galois group of $\mathbb{Q}$. By [Sil09, Chapter X, §3, Theorem 3.6], we identify $\text{WC}(E/\mathbb{Q})$ with a certain cohomology group:

$$\text{WC}(E/\mathbb{Q})\xrightarrow{\sim} H^1(G_{\mathbb{Q}},E(\mathbb{Q}))\,. \tag{2.10}$$

We will omit the construction of this bijection; it says that two equivalent homogenous spaces on the left correspond to the same class of a certain 1-cocycle on the right. More background on Galois cohomology is found in [Sil09, Appendix B]. Importantly, this bijection also gives a group structure to $\text{WC}(E/\mathbb{Q})$.

The Shafarevich-Tate group and Selmer group of E are described in terms of local considerations. Let p be a finite or infinite place of $\mathbb{Q}$. Let $G_{\mathbb{Q}_p}=\text{Gal}(\overline{\mathbb{Q}}_p/\mathbb{Q}_p)$ be the absolute

Galois group of $\mathbb{Q}_p$, acting on $E(\overline{\mathbb{Q}_p})$. For every p, we choose an embedding $\overline{\mathbb{Q}} \hookrightarrow \overline{\mathbb{Q}_p}$, inducing an embedding $\overline{\mathbb{Q}} \hookrightarrow \overline{\mathbb{Q}_p}$. Hence we have an injection $G_{\mathbb{Q}_p} \hookrightarrow G_{\mathbb{Q}}, \sigma \mapsto \sigma|_{\overline{\mathbb{Q}}}$. This induces a restriction map in cohomology $\mathrm{res}_p : H^1(G_{\mathbb{Q}}, E(\overline{\mathbb{Q}})) \to H^1(G_{\mathbb{Q}_p}, E(\overline{\mathbb{Q}_p}))$.

Definition 2.3.2. *The* Shafarevich-Tate group of E *is the subgroup of* $\mathrm{WC}(E/\mathbb{Q})$ *given by*

$$Ш(E/\mathbb{Q}) := \ker\left\{ \mathrm{WC}(E/\mathbb{Q}) \xrightarrow{\prod_p \mathrm{res}_p} \prod_p \mathrm{WC}(E/\mathbb{Q}_p) \right\} .$$

This is the group that we denoted by $Ш(E)$ in Conjecture 2.3.1. The elements of $Ш(E/\mathbb{Q})$ are those equivalence classes of principal homogeneous spaces $C/\mathbb{Q}$ for E such that the equivalence class of $C/\mathbb{Q}_p$ is trivial for every p, i.e. $C/\mathbb{Q}$ is everywhere *locally trivial*. By the above fact, this is equivalent to $C(\mathbb{Q}_p) \neq \emptyset$ for every p. In view of the bijection in Equation (2.10), it is direct to define $Ш(E/\mathbb{Q})$ from a cohomological viewpoint.

It is not known whether $Ш(E/\mathbb{Q})$ is finite, except in special cases. This is a conjecture [Sil09, Chapter X, §5, Conjecture 4.13]. If $Ш(E/\mathbb{Q})$ is finite then its order is a square. This follows because of the existence of an alternating bilinear pairing on $Ш(E/\mathbb{Q})$, due to Cassels and Tate (see [Sil09, Chapter X, §5, Theorem 4.14]). Remark that this is not true in general for abelian varieties.

ϕ-Selmer group. We now define the ϕ-Selmer group of E, for an isogeny of elliptic curves ϕ. Let $\phi : E \to E'$ be an isogeny of elliptic curves, a morphism such that $\phi(\mathcal{O}) = \mathcal{O}$. For example, one can let $\phi = [m]$, scalar multiplication by m, for every $m \in \mathbb{Z}$. We let $E(\mathbb{Q})[\phi] := \ker(\phi)$. There is a short exact sequence of $G_{\mathbb{Q}}$-modules $0 \to E(\mathbb{Q})[\phi] \to E(\mathbb{Q}) \xrightarrow{\phi} E'(\mathbb{Q}) \to 0$. The associated long exact sequence in cohomology is

$$0 \to E(\mathbb{Q})[\phi] \to E(\mathbb{Q}) \xrightarrow{\phi} E'(\mathbb{Q}) \xrightarrow{\delta} H^1(G_{\mathbb{Q}}, E(\mathbb{Q})[\phi]) \to H^1(G_{\mathbb{Q}}, E(\mathbb{Q})) \to H^1(G_{\mathbb{Q}}, E'(\mathbb{Q})) ,$$

where δ is the connecting homomorphism. We extract the fundamental short exact sequence

$$0 \to E'(\mathbb{Q})/\phi(E(\mathbb{Q})) \xrightarrow{\delta} H^1(G_{\mathbb{Q}}, E(\mathbb{Q})[\phi]) \to H^1(G_{\mathbb{Q}}, E(\mathbb{Q}))[\phi] \to 0 . \tag{2.11}$$

Repeating the same for the $G_{\mathbb{Q}_p}$-module $E(\mathbb{Q}_p)$ and $E'(\mathbb{Q}_p)$, gives the exact sequences, for every p:

$$0 \to E'(\mathbb{Q}_p)/\phi(E(\mathbb{Q}_p)) \xrightarrow{\delta} H^1(G_{\mathbb{Q}_p}, E(\mathbb{Q}_p)[\phi]) \to H^1(G_{\mathbb{Q}_p}, E(\mathbb{Q}_p))[\phi] \to 0 . \tag{2.12}$$

Note that in (2.11) and (2.12) the last terms are identified with $\mathrm{WC}(E/\mathbb{Q})[\phi]$ and $\mathrm{WC}(E/\mathbb{Q}_p)[\phi]$, respectively, via Equation (2.10). Both sequences induce the commutative diagram:

$$\begin{array}{ccccccccc}
0 & \longrightarrow & \dfrac{E'(\mathbb{Q})}{\phi(E(\mathbb{Q}))} & \longrightarrow & H^1(G_{\mathbb{Q}}, E[\phi]) & \longrightarrow & \mathrm{WC}(E/\mathbb{Q})[\phi] & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow{\scriptstyle \prod_p \mathrm{res}_p} & \overset{f}{\searrow} & \downarrow{\scriptstyle \prod_p \mathrm{res}_p} & & \\
0 & \longrightarrow & \prod_p \dfrac{E'(\mathbb{Q}_p)}{\phi(E(\mathbb{Q}_p))} & \longrightarrow & \prod_p H^1(G_{\mathbb{Q}_p}, E[\phi]) & \longrightarrow & \prod_p \mathrm{WC}(E/\mathbb{Q}_p)[\phi] & \longrightarrow & 0
\end{array}$$

where we have abbreviated $E[\phi]$ for both $E(\mathbb{Q})[\phi]$ and $E(\mathbb{Q}_p)[\phi]$. The left-most vertical map is induced by the inclusion $E'(\mathbb{Q}) \hookrightarrow E'(\mathbb{Q}_p)$.

Definition 2.3.3. *The ϕ-*Selmer group of E *is defined as the subgroup of $H^1(G_{\mathbb{Q}}, E[\phi])$ given by the kernel of f (the diagonal arrow):*

$$\mathrm{Sel}^{(\phi)}(E/\mathbb{Q}) := \ker\left\{ H^1(G_{\mathbb{Q}}, E(\mathbb{Q})[\phi]) \to \prod_p \mathrm{WC}(E/\mathbb{Q}_p) \right\} .$$

The Snake Lemma applied to the diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \dfrac{E'(\mathbb{Q})}{\phi(E(\mathbb{Q}))} & \longrightarrow & H^1(G_{\mathbb{Q}}, E[\phi]) & \longrightarrow & \mathrm{WC}(E/\mathbb{Q})[\phi] & \longrightarrow & 0 \\ & & \downarrow & & \vdots\, f & & \downarrow \prod_p \mathrm{res}_p & & \\ 0 & \longrightarrow & 0 & \longrightarrow & \prod_p \mathrm{WC}(E/\mathbb{Q}_p)[\phi] & \xrightarrow{\mathrm{id}} & \prod_p \mathrm{WC}(E/\mathbb{Q}_p)[\phi] & \longrightarrow & 0 \end{array}$$

yields the following fundamental exact sequence, relating the ϕ-Selmer group and the ϕ-torsion subgroup of the Shafarevich-Tate group of E:

$$0 \to \frac{E'(\mathbb{Q})}{\phi(E(\mathbb{Q}))} \to \mathrm{Sel}^{(\phi)}(E/\mathbb{Q}) \to Ш(E/\mathbb{Q})[\phi] \to 0 .$$

The ϕ-Selmer group is finite (see [Sil09, Chapter X, §4, Theorem 4.2]) and is effectively computable in practice. This is the main ingredient in the proof of the weak Mordell-Weil Theorem: applied to $\phi = [m] : P \mapsto mP$ it yields the finiteness of $E(\mathbb{Q})/mE(\mathbb{Q})$. Remark that the same procedure can be repeated with the dual isogeny $\phi^\vee : E' \to E$, and one similarly defines $\mathrm{Sel}^{(\phi^\vee)}(E')$.

Isogeny descent. The computation of the ϕ-Selmer group of E becomes very explicit when ϕ is a 2-isogeny $\phi : E \to E'$. This is called a *descent* via 2-isogenies. Namely, in this case the principal homogeneous spaces can be written down explicitly. We state it in the following theorem.

Theorem 2.3.4 (Theorem 4.9, Chapter X.4 in [Sil09])**.** *Let $E : y^2 = x^3 + ax^2 + bx$ over $\mathbb{Q}$ and $E' : y^2 = x^3 - 2ax^2 + (a^2 - 4b)x$ over $\mathbb{Q}$ which is 2-isogenous to E via $\phi : E \to E'$ with $\ker(\phi) = \{\mathcal{O}, (0,0)\}$. Consider the set $\mathcal{S} = \{\infty\} \cup \{p \in \mathbb{N} \text{ prime} : p \mid 2b(a^2 - 4b)\}$ and*

$$\mathbb{Q}(\mathcal{S}) := \{\lambda \in \mathbb{Q}^\times/(\mathbb{Q}^\times)^2 : v_p(\lambda) \equiv 0 \bmod 2 \ \forall p \notin \mathcal{S}\} ,$$

a subgroup of $\mathbb{Q}^\times/(\mathbb{Q}^\times)^2$. For every $\lambda \in \mathbb{Q}(\mathcal{S})$, define the homogenous space for E:

$$C_\lambda : \lambda W^2 = \lambda^2 - 2a\lambda Z^2 + (a^2 - 4b)Z^4 \tag{2.13}$$

(i) There is an exact sequence $0 \to E'(\mathbb{Q})/\phi(E(\mathbb{Q})) \xrightarrow{\delta} \mathbb{Q}(\mathcal{S}) \to \mathrm{WC}(E/\mathbb{Q})[\phi]$, where δ is the connecting homomorphism sending $(X, Y) \mapsto X \bmod (\mathbb{Q}^\times)^2$, $\mathcal{O} \mapsto 1 \bmod (\mathbb{Q}^\times)^2$, $(0,0) \mapsto a^2 - 4b \bmod (\mathbb{Q}^\times)^2$, and $\mathbb{Q}(\mathcal{S}) \to \mathrm{WC}(E/\mathbb{Q})[\phi]$ sends λ to the equivalence class of C_λ.

(ii) The ϕ-Selmer group of E is $\mathrm{Sel}^{(\phi)}(E/\mathbb{Q}) = \{\lambda \in \mathbb{Q}(\mathcal{S}) : C_\lambda(\mathbb{Q}_p) \neq \emptyset \ \forall p \in \mathcal{S}\}$.

(iii) The map $\psi_\lambda : C_\lambda(\mathbb{Q}) \to E'(\mathbb{Q})$, $(Z, W) \mapsto (\lambda/Z^2, -\lambda W/Z^3)$ satisfies $\delta(\psi_\lambda(P)) \equiv \lambda \pmod{(\mathbb{Q}^\times)^2}$ for $P \in C_\lambda(\mathbb{Q})$.

The quartic polynomial defining C_λ describes a hyperelliptic curve with affine coordinates (Z, W). In Chapter 6, we rely on this proposition in the context of elliptic curves birationally equivalent to Edwards curves.

CHAPTER 3

Cryptanalysis of a Multilinear Map Scheme

Cryptographic multilinear maps were introduced relatively recently and find many interesting applications. We first give a high-level introduction to cryptographic multilinear maps. We next give a brief overview of the CLT13 Scheme with its main functionalities and discuss known attacks against the scheme, such as the powerful Cheon et al. attack from Eurocrypt 2015. Our contributions in this chapter are based on an attack described by Gentry, Lewko and Waters [GLW14] against CLT13. This is a simple lattice attack in dimension 2, and a countermeasure was described. We propose a new attack based on higher dimension lattice reduction that breaks the countermeasure from [GLW14] for a wide range of parameters and extends the 2-dimensional attack. Combined with the Cheon et al. attack, our new attack leads to the recovery of all the secret parameters of CLT13, assuming that low-level encodings of almost zero plaintexts are available to the attacker. We conclude this chapter by showing how to apply our attack against concrete constructions based on composite-order CLT13.

Section 3.3 is based on joint work [CN19a] with Jean-Sébastien Coron, which has been published in the proceedings of Asiacrypt 2019. We closely follow the exposition of [CN19a].

3.1 Introduction

3.1.1 Multilinear Maps in Cryptography

The revolutionary work "New Directions in Cryptography" by Diffie and Hellman [DH82] in 1976 is the foundation of public-key cryptography: the Diffie-Hellman key exchange protocol, enabling Alice and Bob to communicate securely over an unsecure channel. About thirty years later, Joux [Jou04] made use of bilinear pairings (on elliptic curves) in order to generalize the Diffie-Hellman key exchange protocol to three users, say Alice, Bob and Charlie (by following conventional names). This construction works as follows. Let $G = \langle g \rangle$ and G_T be cyclic groups (G_T is said a target group) and $e : G \times G \to G_T$ a symmetric non-degenerate bilinear pairing. Upon agreement of Alice, Bob and Charlie on g, they compute $A = g^a$, $B = g^b$ and $C = g^c$ for secretely chosen integers a, b, c, respectively. Once those values are publicly transmitted, Alice, Bob and Charlie can compute the common shared key $g_{abc} := e(g, g)^{abc} \in G_T$. Namely, upon receiving B and C, Alice can compute g_{abc} as $e(B, C)^a = e(g^b, g^c)^a$, and very similarly can do Bob and Charlie. An eavesdropper would need to recover a, given g and g_{abc}, known as the bilinear computational Diffie-Hellman problem, as a direct generalization

of the computational Diffie-Hellman problem.

Multilinear maps attempt to generalize bilinear pairings to a higher-degree multilinearity. In [BS03], Boneh and Silverberg survey about cryptographic multilinear maps and their applications in cryptography. More precisely, following [BS03, Definition 2.1], a symmetric multilinear map of multilinearity-degree $\kappa \in \mathbb{Z}_{\geq 2}$ is a map $e : G^\kappa \to G_T$ for cyclic groups G and G_T, such that

(1) G and G_T have the same prime order

(2) for every $a_1, \ldots, a_\kappa \in \mathbb{Z}$ and $g_1, \ldots, g_\kappa \in G$: $e(g_1^{a_1}, \ldots, g_\kappa^{a_\kappa}) = e(g_1, \ldots, g_\kappa)^{\prod_{1 \leq i \leq \kappa} a_i}$

(3) e is non-degenerate, in the sense that, if g is a generator of G, then $e(g, \ldots, g)$ is a generator of G_T

In order this object to be of cryptographic interest, one requires further that

(4) the groups G and G_T admit efficiently computable group operations

(5) e is efficiently computable

(6) the discrete logarithm problem in G is hard

Under these conditions, e is called a *cryptographic κ-linear map*. For instance, these assumptions are all valid for bilinear pairings, realized as the Weil or Tate pairing on elliptic curves over sufficiently large finite fields. By "efficiently computable" we mean that the corresponding operations can be carried out by algorithms whose running time is polynomial in the input length. Solving the discrete logarithm problem in a finite cyclic group G of order q means, given $g \in G$ and $g^a \in G$ for some $a \in \{0, \ldots, q-1\}$, to compute a. It is enough to ensure the hardness of this problem in G only, as it implies the hardness in G_T. Namely, writing $h = g^a \in G$ and letting $h^\star = e(h, g \ldots, g)$ and $g^\star = e(g, \ldots, g)$, one has that $h^\star = (g^\star)^a \in G_T$, i.e. a is also a discrete logarithm in G_T. Note that the definition above restricts to *prime order* groups G, G_T because it is known that the discrete logarithm problem in a composite order group is reduced to the discrete logarithm problem in its prime order subgroups, by an algorithm of Pohlig-Hellman. For simplicity, we only treat the *symmetric* case; the *asymmetric* case consists in having different source groups $G_1, \ldots, G_\kappa$ of the same prime order.

The paper [BS03] however points to a rather pessimistic conclusion on the existence of higher-degree multilinear maps from algebraic geometry, as for example, generalizations of elliptic curve bilinear pairings.

3.1.2 Graded Encoding Schemes

A breakthrough was obtained in 2009 by Gentry's fully homomorphic encryption scheme, [Gen09]. Fully homomorphic encryption (FHE) allows to add and multiply ciphertexts, with the underlying plaintexts being added and multiplied, without the need of being decrypted. In 2013, this led Garg, Gentry and Halevi to the observation that FHE ciphertexts behave like group exponents for multilinear maps, and formally introduced the notion of a *graded encoding scheme*, [GGH13a]. Informally, one here considers a family of cyclic groups $G_0, \ldots, G_\kappa$ of the same prime order q, carrying efficiently computable group operations, together with bilinear pairings $e_{i,j} : G_i \times G_j \to G_{i+j}$ for all i, j such that $i + j \leq \kappa$. If $g_{i+j} = e(g_i, g_j)$,

then, addition of exponents within G_i corresponds to $g_i^a g_i^b = g_i^{a+b}$, whereas multiplication (whenever $i + j \leq \kappa$) corresponds to $e(g_i^a, g_j^b) = g_{i+j}^{ab}$.

We omit the precise definition of a κ-graded encoding system and refer to [GGH13a, Definition 2]. In light of their conceptual differences with "ideal" multilinear maps as proposed in [BS03], the construction of graded encoding schemes is viewed as an *approximation* of cryptographic multilinear maps. Roughly speaking, the authors consider rings instead of groups in order to be able to both add and multiply. An encoding of a ring element will be of the form g^a. From this definition, the authors in [GGH13a] propose the first plausible construction of a cryptographic multilinear map, based on *ideal lattices*. Although graded encoding schemes present some conceptual differences, they achieve similar goals than cryptographic multilinear maps, raising imminent interest in the cryptographic community.

Shortly after, two more constructions of graded encoding schemes appeared in the literature. In 2013, Coron, Lepoint and Tibouchi [CLT13] present an analogous construction (called CLT13) but over the integers, based on the fully homomorphic encryption scheme from [VDGHV10]. Two years later, a third family was published by Gentry, Gorbunov and Halevi [GGH15] (referred to as GGH15), based on the LWE Problem with matrix encodings.

In this chapter, we mainly work with the CLT13 Scheme. We postpone a more detailed overview of this construction to the next section.

3.1.3 Some applications of Multilinear Maps

The vast interest for building efficient multilinear maps is certainly due to their large number of applications. Here we mention some without however entering in their details.

In light of the above discussion on Joux's 3-party Diffie-Hellman key exchange [Jou04], the most straightforward application of multilinear maps probably is a *non-interactive one-round multiparty Diffie-Hellman key exchange*. Following the construction of Joux with a cryptographic n-linear map, one directly constructs a one-round key exchange protocol for $n + 1$ users. In [GGH13a], this construction is also adapted to work for graded encoding schemes.

Other applications include different types of encryption, such as *attribute-based encryption*, *witness encryption*, and *functional encryption*. See for example, [GGH$^+$13d, GGSW13, GGH$^+$13c]. Another strong application of multilinear maps is *indistinguishability obfuscation* (iO), which followed almost immediately after the GGH13 scheme was introduced (see [GGH$^+$13c]). At rough level, iO aims at publishing programs whose functionality depends on certain secrets, without the source code of the programs revealing those secrets. More formally, an *obfuscator* for a certain program $\mathcal{C}$ can be thought of as a function $\mathcal{O}$ which outputs a modified program $\mathcal{O}(\mathcal{C})$ while not changing the functionality of $\mathcal{C}$ (i.e. $\mathcal{O}(\mathcal{C})$ and $\mathcal{C}$ compute the same output for all inputs), and such that for any two programs $\mathcal{C}$ and $\mathcal{C}'$ with the same functionality, one cannot distinguish between their obfuscations $\mathcal{O}(\mathcal{C})$ and $\mathcal{O}(\mathcal{C}')$ (i.e. $\mathcal{O}(\mathcal{C})$ and $\mathcal{O}(\mathcal{C}')$ are computationally indistinguishable). Indistinguishability obfuscation is commonly formalized in the language of *matrix branching programs* following [GGH13a]. We will recall this notion in more detail at the end of this chapter (see Section 3.3.10).

3.2 The CLT13 multilinear map

We review some details of the CLT13 multilinear map scheme over the integers, by Coron, Lepoint and Tibouchi (see [CLT13]).

3.2.1 The CLT13 multilinear Map

Let $n \in \mathbb{Z}_{\geq 1}$, which we regard as a dimension for the CLT13 scheme. This integer should be set large enough to ensure correctness and security. The instance generation of CLT13 generates n distinct secret "large" prime numbers $p_1, \ldots, p_n$ of bit size η, and publishes the modulus $x_0 = \prod_{i=1}^{n} p_i$. We let γ denote the bit size of x_0; therefore $\gamma \approx n \cdot \eta$. One also generates n distinct secret "small" prime numbers $g_1, \ldots, g_n$ of bit size α. We do not make the words "large" and "small" precise here, but the reader should interpret a significant size difference between α and η. The plaintext ring is composite, i.e. a plaintext is an element $m = (m_1, \ldots, m_n)$ of the ring $\mathbb{Z}/G\mathbb{Z} \simeq \bigoplus_{i=1}^{n} \mathbb{Z}/g_i\mathbb{Z}$ where $G = \prod_{i=1}^{n} g_i$.

Let $\kappa \in \mathbb{Z}_{\geq 1}$ be the multilinearity degree, i.e. we construct a κ-linear map. For $k \in \{1, \ldots, \kappa\}$, an encoding at level k of the plaintext m is an integer $c \in \mathbb{Z}$ such that

$$c \equiv \frac{r_i g_i + m_i}{z^k} \pmod{p_i}, \text{ for all } 1 \leq i \leq n \tag{3.1}$$

for "small" random integers $\{r_i\}_i$ of bit size ρ; we refer to $\{r_i\}_i$ as the *noise* in the encodings. Here z is a random secret integer which is invertible modulo x_0 and which is the same for all encodings. We refer to encodings of level $k = \kappa$ as *top-level encodings*, or encodings at the *last level*. Remark that since the primes $\{p_i\}_i$ are secret, the user cannot directly encode messages using Equation (3.1). For this reason, the public parameters also include a certain set of level-zero encodings of random messages, and the user generates a level-zero encoding by computing a random subset-sum of those public level-zero encodings.

The CLT13 multilinear map supports homomorphic properties. From Equation (3.1), it is clear that two encodings at the same level can be added, and the underlying plaintexts get added in $\mathbb{Z}/G\mathbb{Z}$. For example, let c_1 be an encoding of $m^{(1)} = (m_i^{(1)})_i$ and c_2 an encoding of $m^{(2)} = (m_i^{(2)})_i$ at the same level $k \in \{1, \ldots, \kappa\}$, then $c_1 + c_2 \equiv (r_i g_i + m_i^{(1)} + m_i^{(2)})/z^k \pmod{p_i}$ for all $1 \leq i \leq n$, where $r_i := r_i^{(1)} + r_i^{(2)}$ is the noise corresponding to the addition of the encodings c_1 and c_2, with self-explaining notation; thus, $c_1 + c_2$ encodes $m^{(1)} + m^{(2)}$ at level k. In particular, the sum of κ level-1 encodings is an encoding at level κ. Similarly, the product of two encodings at level i and j gives an encoding of the product plaintexts at level $i + j$, as long as the numerators in Equation (3.1) do not grow too large, i.e. they must remain smaller than each p_i.

Let us now describe how zero-testing for a top-level encoding c works. The instance generation publishes the zero-testing parameter p_{zt}, defined as

$$p_{zt} := \sum_{i=1}^{n} h_i z^{\kappa} (g_i^{-1} \bmod p_i) \frac{x_0}{p_i} \bmod x_0 , \tag{3.2}$$

where $\{h_i\}_i \subseteq \mathbb{Z}$ are "small" random integers of bit size n_h (denoted by β in [CLT13]), and $g_i^{-1} \bmod p_i$ is the multiplicative inverse of g_i in $(\mathbb{Z}/p_i\mathbb{Z})^{\times}$ for $1 \leq i \leq n$. Zero-testing of a top-level encoding c then consists in multiplying c by p_{zt} and checking whether the result is significantly smaller than x_0. More precisely, we publicly compute:

$$\omega := p_{zt} \cdot c \bmod x_0 \equiv \sum_{i=1}^{n} h_i (r_i + m_i (g_i^{-1} \bmod p_i)) \frac{x_0}{p_i} \pmod{x_0} \tag{3.3}$$

and we consider that c encodes the zero message if ω (i.e. the representative in $(-x_0/2, x_0/2] \cap \mathbb{Z}$) is "small" compared to x_0. Namely, if $m_i = 0$ for all $1 \leq i \leq n$, that is, m is the zero-message, then we obtain from Equation (3.3):

$$\omega \equiv \sum_{i=1}^{n} h_i r_i \frac{x_0}{p_i} \pmod{x_0} ,$$

and since $\{h_i\}_i$ and $\{r_i\}_i$ are "small", the resulting ω will be "small" compared to x_0. More precisely, let ρ_∞ be the maximum bit size of the noise integers $\{r_i\}_i$ in the encodings. Then $\{h_i r_i x_0/p_i\}_i$ have bit size roughly $\gamma - \eta + n_h + \rho_\infty$, and letting

$$\nu := \eta - n_h - \rho_\infty , \tag{3.4}$$

they have bit size roughly $\gamma - \nu$ bits. Consequently, when $m_i = 0$ for all i, ω has bit size roughly $\gamma - \nu$ bits; whereas when $m_i \neq 0$ for some i, we expect that ω is of full size modulo x_0, that is, γ bits. The parameter ν in Equation (3.4) corresponds to the number of bits that can be extracted from zero-testing; namely from Equation (3.3), the ν most significant bits of ω only depend on the plaintext messages m_i, and not on the noise r_i. We refer to [CLT13, Lemma 3] for precise bounds needed for correct zero-testing. On input κ, the instance generation of CLT13 outputs the parameters $(n, \eta, \alpha, \rho, n_h, \nu, x_0, p_{zt})$ together with a strong randomness extractor.

Note that to obtain a proper zero-testing procedure, one needs to use a *vector* of n elements p_{zt}; namely using a single p_{zt} there exist encodings c with $m_i \neq 0$ while $p_{zt} \cdot c$ is "small" modulo x_0. In this case, an integer matrix $H = (h_{ij})_{ij} \in \mathbb{Z}^{n \times n}$ with small infinity norm is sampled instead of the vector $(h_i)_i \in \mathbb{Z}^n$. For simplicity, in this chapter, we mainly consider a single p_{zt}, as it is usually the case in constructions over CLT13 multilinear maps. We refer to [CLT13, Section 3.1] for the precise setting of the parameters.

3.2.2 Cryptanalysis of CLT13 multilinear maps

3.2.2.1 Cryptanalysis of Multilinear Maps

We now review some important cryptanalysis of the CLT13 scheme, and provide details for the powerful attack proposed by Cheon et al. [CHL+15]. First, let us briefly recall some of the history for the cryptanalysis of current multilinear map constructions.

Over the last years, numerous powerful attacks have appeared against all three proposals of multilinear maps [GGH13a, CLT13, GGH15]. An important family of attacks against multilinear maps are so-called "zeroizing attacks", which recover in polynomial time the secret parameters from encodings of zero, using linear algebra. For the non-interactive multipartite Diffie-Hellman key exchange, the zeroizing attack proposed by Cheon et al. [CHL+15] at Eurocrypt 2015 recovers all secret parameters of CLT13; the attack can be extended to encoding variants where encodings of zero are not directly available [CGH+15]. Shortly after, a new proposal [CLT15] replacing CLT13 and claiming tentative fixes against the Cheon et al. attack, appeared, but was shown insecure, in [CFL+16]. The zeroizing attack from [HJ16] also breaks the Diffie-Hellman key exchange over GGH13. Finally, the key exchange over GGH15 was also broken in [CLLT16], using an extension of the Cheon et al. zeroizing attack.

Even though direct multipartite key exchange protocols are broken for the three known families of multilinear maps, more complex constructions are not necessarily broken. For example, for indistinguishability obfuscation (iO), low-level encodings of zero are generally not available, hence preventing zeroizing attacks. However the Cheon et al. attack against CLT13 was extended in [CGH+15] to matrix branching programs where the input can be partitioned into three independent sets. The attack was extended in [CLLT17] to branching programs without a simple input partition structure, using a so-called tensoring technique. For obfuscation based on GGH13, Miles, Sahai and Zhandry [MSZ16] introduced "annihilation attacks" that break a certain class of matrix branching programs; the attack was later extended in [CGH17] to break the [GGH+13b] obfuscation under GGH13, using a variant of the input partitioning attack. Finally, Chen, Vaikuntanathan and Wee described in [CVW18] an attack against iO over GGH15. In general, the above attacks only apply against branching programs with a simple structure, and breaking more complex constructions (such as dual-input branching programs) is currently infeasible.

3.2.2.2 The Cheon et al. attack against CLT13 with encodings of zero

Clearly, one way to break CLT13 is by factoring x_0, which is doable in classic sub-exponential time via ECM [Len87], or in quantum polynomial time using Shor's algorithm [Sho97]. The Cheon et al. attack factors x_0 in polynomial time on a classical computer, when sufficiently many encodings of zero are available[1] to the attacker. From the factorization of x_0, all the other secret parameters can be computed. This breaks the CLT13-multiparty Diffie-Hellman key exchange. The attack has been extended to matrices of encodings in [CGH+15], which we will not recall here.

We now recall how the attack works; we refer to [CHL+15] for full details. For simplicity of exposition, we consider trilinear maps, i.e. the case $\kappa = 3$; the attack is easily extended to $\kappa > 3$. Consider disjoint sets $\mathcal{A} = \{\alpha_j : 1 \leq j \leq n\}$, $\mathcal{B} = \{\beta_1, \beta_2\}$ and $\mathcal{C} = \{\gamma_k : 1 \leq k \leq n\}$ of encodings at level 1 and where all encodings in $\mathcal{A}$ encode zero. Therefore, there are $\#\mathcal{A} = n$ public encodings of zero and $\#(\mathcal{A} \cup \mathcal{B} \cup \mathcal{C}) = 2n + 2$ encodings in total. Following the CLT13-notation introduced above, we write

$$\alpha_j \equiv \alpha_{ji}/z \pmod{p_i}\ ,\ \beta_a \equiv \beta_{ai}/z \pmod{p_i}\ ,\ \gamma_k \equiv \gamma_{ki}/z \pmod{p_i}$$

for all $i, j, k \in \{1, \ldots, n\}$ and $a \in \{1, 2\}$. Because $\alpha_j \beta_a \gamma_k$ encodes zero at level 3 (since the encodings in $\mathcal{A}$ encode zero), correct zero-testing ensures that the zero-test equations $\omega_{jk}^{(a)} = p_{zt}(\alpha_j \beta_a \gamma_k)$, given by the matrix notations

$$\omega_{jk}^{(a)} = \sum_{i=1}^{n} p_{zt,i} \alpha_{ji} \beta_{ai} \gamma_{ki} = \begin{bmatrix} \alpha_{j1} \cdots \alpha_{jn} \end{bmatrix} \begin{bmatrix} \beta_{a1} p_{zt,1} & & \\ & \ddots & \\ & & \beta_{an} p_{zt,n} \end{bmatrix} \begin{bmatrix} \gamma_{k1} \\ \vdots \\ \gamma_{kn} \end{bmatrix}$$

with $p_{zt,i} = (h_i (g_i^{-1} \bmod p_i) \cdot x_0 / p_i$ for $1 \leq i \leq n$ defining the zero-test parameter, hold over $\mathbb{Z}$ instead of $\mathbb{Z}/x_0\mathbb{Z}$. Writing these relations out for all indices $(j, k) \in \{1, \ldots, n\}^2$, the $n \times n$ matrices

$$W_a := (\omega_{jk}^{(a)})_{1 \leq j,k \leq n}\ , \quad a \in \{1, 2\} \tag{3.5}$$

[1] Such encodings are for instance public in the rerandomization procedure of CLT13.

satisfy the matrix equalities

$$W_a = P \cdot U_a \cdot Q \ , \quad a \in \{1, 2\} \tag{3.6}$$

for secret matrices P, Q of full rank n (corresponding to encodings of $\mathcal{A}$ and $\mathcal{C}$, respectively) and diagonal matrices U_1, U_2 containing the elements $\{\beta_{ai} p_{zt,i} : 1 \leq i \leq n\}$. If at least one of W_1, W_2 is invertible over $\mathbb{Q}$ (say W_2), the attacker computes the eigenvalues of the product

$$W_1 \cdot W_2^{-1} = P(U_1 U_2^{-1})P^{-1} \ ,$$

by factoring the characteristic polynomial (over $\mathbb{Q}$). By similarity, these eigenvalues coincide with those of $U_1 U_2^{-1}$ which are $\{\beta_{1i}/\beta_{2i} : 1 \leq i \leq n\}$. These ratios are now enough to factor x_0 completely: namely, writing $\beta_{1i}/\beta_{2i} = x_i/y_i$ for coprime integers x_i, y_i and using that $\beta_a \equiv \beta_{ai}/z \pmod{p_i}$, one obtains $x_i\beta_2 - y_i\beta_1 \equiv (x_i\beta_{2i} - y_i\beta_{1i})/z \equiv 0 \pmod{p_i}$ for $1 \leq i \leq n$, and therefore, one reveals the primes $\{p_i\}_i$ by computing $\gcd(x_i\beta_2 - y_i\beta_1, x_0) = p_i$ (this holds with high probability, as it is likely that, the only $1 \leq j \leq n$ such that this gcd is p_j is $j = i$). In summary, this attack recovers all secret prime factors $\{p_i\}_i$ of x_0 in polynomial time, given as input the set $\mathcal{A}$ of level-one encodings of zero and the sets $\mathcal{B}$ and $\mathcal{C}$. The running time of this algorithm is polynomial time and relies on basic algorithms from linear algebra. We give a summary in Algorithm 2.

Algorithm 2 Cheon et al. attack against CLT13 with encodings of zero

Parameters: The CLT13 parameters with $\kappa = 3$
Input: Disjoint sets $\mathcal{A}, \mathcal{B}, \mathcal{C}$ of encodings at level one, where the encodings in $\mathcal{A}$ encode zero, and $\#\mathcal{A} = \#\mathcal{C} = n$ and $\#\mathcal{B} = 2$; and $x_0 = \prod_{i=1}^{n} p_i$
Output: The prime factors $\{p_i : 1 \leq i \leq n\}$ of x_0

1: **for** $a \in \{1, 2\}$ **do**
2: Construct the matrix $W_a \in \mathbb{Z}^{n \times n}$ as in Equation (3.5)
3: **end for**
4: Compute the eigenvalues $\{\beta_{1i}/\beta_{2i}\}_i$ of $W_1 W_2^{-1}$ (assuming invertibility of W_2)
5: **for** $1 \leq i \leq n$ **do**
6: Compute coprime integers x_i, y_i such that $\beta_{1i}/\beta_{2i} = x_i/y_i$
7: Compute and return $\gcd(x_i\beta_2 - y_i\beta_1, x_0)$, where $\mathcal{B} = \{\beta_1, \beta_2\}$
8: **end for**

3.3 Cryptanalysis of CLT13 with Independent Slots

3.3.1 Introduction

Many constructions based on multilinear maps require independent slots in the plaintext so that multiple computations can be performed in parallel over the slots. Such constructions are usually based on CLT13 multilinear maps, since CLT13 inherently provides a composite encoding space $\bigoplus_{i=1}^{n} \mathbb{Z}/g_i\mathbb{Z}$ for small prime numbers $\{g_i\}_i$. However, a vulnerability was identified at Crypto 2014 by Gentry, Lewko and Waters, with a 2-dimensional lattice-based attack, and the authors have suggested a simple countermeasure. In this section, we identify

an attack based on higher dimension lattice reduction that breaks the author's countermeasure for a wide range of parameters. Combined with the Cheon et al. attack recalled in Section 3.2.2.2, this leads to the recovery of all the secret parameters of CLT13, assuming that low-level encodings of almost zero plaintexts are available. We show how to apply our attack against various constructions over composite-order CLT13. For the [FRS17] construction, our attack enables to recover the secret CLT13 plaintext ring for a certain range of parameters; however, breaking the indistinguishability of the branching program remains an open problem.

Multilinear maps with independent slots. Many constructions based on multilinear maps require independent slots in the plaintext so that multiple computations can be performed in parallel over the slots when evaluating the multilinear map. For example, [GLW14] and [GLSW15] use independent slots to obtain improved security reductions for witness encryption and obfuscation. Multilinear maps with independent slots were also used in the circuit based constructions of [AB15, Zim15]. The construction from [FRS17], which gives a powerful technique for preventing zeroizing attacks against iO, is also based on multilinear maps with independent slots.

The CLT13 multilinear map scheme inherently supports a composite integer encoding space, with a plaintext ring $\mathbb{Z}/G\mathbb{Z} \simeq \bigoplus_{i=1}^{n} \mathbb{Z}/g_i\mathbb{Z}$ for small secret primes g_i's and $G = g_1 \cdots g_n$. For example, in the construction from [FRS17], every branching program works independently modulo each g_i. In that case, the main difference with the original CLT13 is that the attacker can obtain encodings of subring elements which are zero modulo all g_i's except one; for example, in [FRS17] this would be done by carefully choosing the input so that all branching programs would evaluate to zero except one. Whereas in the original CLT13 construction, one never provides encodings of subring elements; instead one uses an "all-or-nothing" approach: either the plaintext element is zero modulo all $\{g_i\}_i$, or it is non-zero modulo all $\{g_i\}_i$ (with high probability).

The attack and countermeasure from [GLW14]. At Crypto 2014, Gentry, Lewko and Waters observed that using CLT13 with independent slots leads to a simple lattice attack in dimension 2, which efficiently recovers the (secret) plaintext ring $\bigoplus_{i=1}^{n} \mathbb{Z}/g_i\mathbb{Z}$ (see [GLW14, Appendix B]). Namely, when using CLT13 with independent slots, the attacker can obtain encodings where all slots are zero modulo g_i except one. For example, for a matrix branching program evaluation as in [FRS17], the result of the program evaluation could have the form:

$$A(x) \equiv \sum_{i=1}^{n} h_i \cdot (r_i + m_i \cdot (g_i^{-1} \bmod p_i)) \cdot \frac{x_0}{p_i} \pmod{x_0}$$

where $m_i = 0$ for all i except $m_{i_0} \neq 0$ for some $1 \leq i_0 \leq n$. This implies:

$$g_{i_0} A(x) \equiv h_{i_0}(r_{i_0} g_{i_0} + m_{i_0}) \frac{x_0}{p_{i_0}} + \sum_{i \neq i_0} g_{i_0} h_i r_i \frac{x_0}{p_i} \pmod{x_0}$$

and therefore $g_{i_0} A(x) \bmod x_0$ is "small" (significantly smaller than x_0). Since g_{i_0} is very small, we can then recover g_{i_0} using lattice reduction in dimension 2, while normally the $\{g_i\}_i$ are secret in CLT13. Moreover, once we know g_{i_0}, we can simply multiply the evaluation by g_{i_0} to obtain a "small" result, even if the evaluation of the branching program is non-zero

modulo g_{i_0}; in particular, this cancels the effect of the protection against input partitioning from [FRS17].

The countermeasure considered in [GLW14, Appendix B] is to give many "buddies" to each g_i, so that we do not have a plaintext element which is non-zero modulo a single isolated g_i. Then, either an encoding is 0 modulo g_i and all its prime buddies g_{i_0}, or it is (with high probability) non-zero modulo all of them. In other words, instead of using individual $\{g_i\}_i$ to define the plaintext slots, every slot is defined modulo a product of θ primes for some $1 \leq \theta < n$. Therefore, we obtain a total of $\lfloor n/\theta \rfloor$ plaintext slots (instead of n). While the above attack can be extended by multiplying $A(x)$ by the θ corresponding primes, for large enough θ the right-hand side of the equation is not "small" anymore and the attack is thwarted.

3.3.2 Our contributions

We identify an attack based on higher dimension lattice reduction that breaks the countermeasure from [GLW14, Appendix B] for a wide range of parameters, with significant impact on the security of CLT13 multilinear maps with independent slots.

Analysis of the attack from [GLW14]. We first provide a theoretical study of the above attack, in order to derive a precise bound on θ as a function of the CLT13 parameters, where θ is the number of primes $\{g_i\}_i$ for each plaintext slot. Note that such an explicit bound was not given in [GLW14]. We argue that, when ν denotes the number of bits that can be extracted from zero-testing in CLT13, the 2-dimensional lattice attack requires the following simplified condition on the parameters:

$$\alpha\theta < \frac{\nu}{2}\,, \tag{3.7}$$

where α is the bit size of the primes $\{g_i\}_i$.

Breaking the countermeasure from [GLW14]. Our main contribution in this chapter is to extend the 2-dimensional attack mentioned above and to break the countermeasure for larger values of θ. Our attack is based on lattice reduction in higher dimensions, by using a variant of the orthogonal lattice attack exploited in [NS99] for solving the hidden subset sum problem. In this extension, we use ℓ encodings $\{c_j : 1 \leq j \leq \ell\}$ where the corresponding plaintexts have only θ non-zero components modulo the primes $\{g_i\}_i$ (instead of $\ell = 1$ in the attack from [GLW14]). Constructing and reducing a certain public lattice in dimension $\ell + 1$, we show that our attack requires the simplified condition $\left(1 + \frac{1}{\ell}\right)\alpha\theta < \nu$ for the parameters. Therefore, for moderately large values of ℓ, this entails the simpler condition $\alpha\theta < \nu$, which gives an improvement of (3.7) by a factor 2.

In the same vein, we show how to further improve this condition by considering products of encodings of the form $c_j \cdot d_k$ for $1 \leq j \leq \ell$ and $1 \leq k \leq d$, where, as previously, the plaintexts of the encodings $\{c_j\}_j$ have only θ non-zero components modulo the primes $\{g_i\}_i$. In that case, using a variant of the previous lattice attack, this time in dimension $\ell + d$, the bound improves to:

$$\alpha\theta = O(\nu^2)$$

The above bound also applies when a vector of zero-testing elements is available, instead of a single zero-test parameter p_{zt}. While the original attack from [GLW14] recovers the secret plaintext ring of CLT13, we additionally recover the plaintext messages $\{m_j : 1 \leq j \leq \ell\}$ for

the encodings $\{c_j : 1 \leq j \leq \ell\}$, up to a scaling factor.

We provide in Section 3.3.8 the result of practical experiments for our attack. For the original parameters of [CLT13], our attack takes a few seconds for $\theta = 40$, and a few hours for θ as large as 160, while the original attack from [GLW14] only works for $\theta = 1$. In summary, our attack is more powerful than the attack in [GLW14], as it allows much larger values for θ, and additionally recovers secret information about the plaintext messages. Finally, we propose a set of secure parameters for CLT13 multilinear maps that prevents our extended attack. For example, for $\lambda = 80$ bits of security, we recommend to take $\theta \geq 1789$.

Recovering all the secret parameters of CLT13. For the range of parameters derived in our extended high-dimensional attack, we show how to combine our algorithm with the Cheon et al. attack from [CHL$^+$15], in order to reveal all secret parameters of CLT13. More precisely, when intermediate-level encodings of partially zero messages are available, our approach consists in applying our lattice attack to generate intermediate-level encodings of zero; then the Cheon et al. attack is applied on these newly-created encodings of zero, to recover all secret parameters.

Application to CLT13-based constructions. Finally, we show how our attack affects the parameter selection of several schemes based on CLT13 multilinear maps with independent slots, namely the constructions from [GLW14, GLSW15, Zim15] and [FRS17]. For the [FRS17] construction, our lattice attack enables to recover the secret CLT13 plaintext ring for a certain range of parameters; however, breaking the indistinguishability of the branching program remains an open problem.

3.3.3 Basic Attack against CLT13 with Independent Slots

In this section, we review the basic attack identified in [GLW14] based on 2-dimensional lattice reduction, as well as the proposed countermeasure. We then provide a detailed analysis of the parameters for this attack.

3.3.3.1 The basic attack from [GLW14]

When using CLT13 with independent slots, the attacker can obtain encodings of plaintext elements where all slots are zero modulo g_i except one. For example, in the [FRS17] construction where each branching program works modulo g_i, the attacker can choose the input so that the resulting evaluation is 0 modulo all primes $\{g_i\}_i$ except one, say g_1, without loss of generality. Let c be a level-κ encoding of a plaintext $m = (m_1, \ldots, m_n)$ where $m_i = 0$ for all $2 \leq i \leq n$. From Equation (3.3), we then obtain the following zero-testing evaluation:

$$\omega \equiv h_1 \cdot m_1 \cdot (g_1^{-1} \bmod p_1) \cdot \frac{x_0}{p_1} + \sum_{i=1}^{n} h_i \cdot r_i \cdot \frac{x_0}{p_i} \pmod{x_0}$$

which implies, by multiplying by g_1:

$$g_1 \cdot \omega \equiv h_1 \cdot m_1 \cdot \frac{x_0}{p_1} + \sum_{i=1}^{n} g_1 \cdot h_i \cdot r_i \cdot \frac{x_0}{p_i} \pmod{x_0}\,.$$

As the integers h_i and r_i are "small", one sees that $g_1 \cdot \omega \bmod x_0$ is significantly smaller than x_0. Using a lattice reduction in dimension 2 for a well-chosen lattice, this implies that we can recover g_1 and similarly the other $\{g_i\}_i$, while normally the $\{g_i\}_i$ are secret in CLT13. This eventually recovers the plaintext ring for CLT13. We analyze the attack below.

The proposed countermeasure. To prevent the attack, the following countermeasure was suggested by the authors: instead of using individual primes $\{g_i\}_i$ to define the plaintext slots, every slot is defined modulo a product of θ primes $\{g_i\}_i$, where $2 \leq \theta < n$. Therefore, a plaintext element cannot be non-zero modulo a single prime g_i; it has to be non-zero modulo at least θ primes $\{g_i\}_i$. Assuming for simplicity that θ divides n, this gives a total of n/θ plaintext slots instead of n.

Therefore, the original CLT13-plaintext ring $R = \mathbb{Z}/g_1\mathbb{Z} \times \cdots \times \mathbb{Z}/g_n\mathbb{Z}$ can be rewritten as $R = \bigoplus_{j=1}^{n/\theta} R_j$, where for all $1 \leq j \leq n/\theta$, the subrings R_j are such that $R_j \simeq \bigoplus_{i=1}^{\theta} \mathbb{Z}/g_{(j-1)\theta+i}\mathbb{Z}$. We can assume that the attacker can obtain encodings of random subring plaintexts in R_j for any $1 \leq j \leq n/\theta$. In that case, the attacker obtains an encoding c of $m = (m_1, \ldots, m_n) \in R$ where $m_i \equiv 0 \pmod{g_i}$ for all $i \in \{1, \ldots, n\} \setminus \{(j-1)\theta + 1, \ldots, j\theta\}$. In that case we will say that m has non-zero support of length θ. Following the lines of attack and assuming that the zero components are in the first slot, one would now multiply by the product $g_1 \ldots g_\theta$ instead of only g_1. Consequently the product $g_1 \cdots g_\theta \cdot \omega$ grows and $g_1 \ldots g_\theta$ can eventually not be recovered by lattice reduction anymore.

3.3.3.2 Analysis of the basic attack

We now analyze the attack from [GLW14] in more details, and derive an explicit bound on the parameter θ, as a function of the other CLT13 parameters.

Given an integer $1 \leq \theta < n$ (the attack mentioned in Section 3.3.3.1 is obtained for $\theta = 1$), we consider a message having non-zero support of length θ; without loss of generality, we can assume that it is of the form $m = (m_1, \ldots, m_n) \in \mathbb{Z}^n$ with $0 \leq m_i < g_i$ such that $m_i = 0$ for $\theta + 1 \leq i \leq n$, i.e. we assume that the non-zero support of m is located in the first slot. We consider a level-κ encoding c of m, satisfying

$$c \equiv \frac{r_i g_i + m_i}{z^\kappa} \pmod{p_i}, \quad 1 \leq i \leq n$$

with integers $\{r_i\}_i$ of bit size ρ_∞. From zero-testing, we then obtain from Equation (3.3):

$$\omega \equiv p_{zt} \cdot c \equiv \sum_{i=1}^{\theta} h_i (g_i^{-1} \bmod p_i) m_i \frac{x_0}{p_i} + \sum_{i=1}^{n} h_i r_i \frac{x_0}{p_i} \pmod{x_0}$$

By multiplying by $g := \prod_{i=1}^{\theta} g_i$ this leads to:

$$g\omega \equiv \sum_{i=1}^{\theta} h_i m_i \frac{g}{g_i} \frac{x_0}{p_i} + \sum_{i=1}^{n} g h_i r_i \frac{x_0}{p_i} \pmod{x_0} \tag{3.8}$$

Letting $u = \sum_{i=1}^{\theta} h_i m_i (g/g_i)(x_0/p_i) + \sum_{i=1}^{n} g h_i r_i (x_0/p_i)$, this simplifies to

$$g\omega \equiv u \pmod{x_0} \tag{3.9}$$

Since the integers h_i and r_i are "small" in order to ensure correct zero-testing, the integer u is "small" in comparison to x_0. More precisely, the proposition below shows that if $g \cdot u$ is a bit smaller than x_0, then we can recover g and u by lattice reduction in dimension 2.

Proposition 3.3.1. *Let $g, \omega, u \in \mathbb{Z}_{\geq 1}$ and $x_0 \in \mathbb{Z}_{\geq 1}$ be such that $g\omega \equiv u \pmod{x_0}$ and $\gcd(\omega, x_0) = \gcd(u, g) = 1$. Assume that $g \cdot u < x_0/10$. Given ω and x_0 as input, one can recover g and u in polynomial time.*

Proof. Without loss of generality we can assume $g \leq u$, as otherwise we can apply the algorithm with $u\omega^{-1} \equiv g \pmod{x_0}$. Let $B \in \mathbb{Z}_{\geq 1}$ such that $u \leq Bg \leq 2u$. When the bit size of g and u is unknown, such a B can be found by exhaustive search in polynomial time. We consider the lattice $\mathcal{L} \subseteq \mathbb{Z}^2$ of vectors (Bx, y) such that $x\omega \equiv y \pmod{x_0}$. From $g\omega \equiv u \pmod{x_0}$ it follows that $\mathcal{L}$ contains the vector $v = (Bg, u)$. We show that v is a shortest non-zero vector in $\mathcal{L}$. By Minkowski's Theorem (Theorem 2.2.2), we have $\lambda_1(L) \leq \sqrt{2\mathrm{Vol}(\mathcal{L})}$. From Hadamard's Inequality (Proposition 2.2.1) together with $\mathrm{Vol}(\mathcal{L}) = Bx_0$, we obtain:

$$\lambda_2(\mathcal{L}) \geq \frac{\mathrm{Vol}(\mathcal{L})}{\lambda_1(\mathcal{L})} \geq \frac{\sqrt{\mathrm{Vol}(\mathcal{L})}}{\sqrt{2}} = \frac{\sqrt{Bx_0}}{\sqrt{2}} > \sqrt{5Bgu} \geq \sqrt{5}u.$$

Moreover, we have that $\|v\| = ((Bg)^2 + u^2)^{1/2} \leq \sqrt{5}u$. This implies that $\|v\| < \lambda_2(\mathcal{L})$ and consequently v is a multiple of a shortest non-zero vector in $\mathcal{L}$: we write $v = kv'$ with $\|v'\| = \lambda_1(\mathcal{L})$, and $k \in \mathbb{Z}\backslash\{0\}$. Letting $v' = (Bv_1, v_2)$, we have $g = kv_1$ and $u = kv_2$. Hence k divides both g and u. Since $\gcd(g, u) = 1$ one has $k = \pm 1$. This shows that v is a shortest non-zero vector of $\mathcal{L}$. By running Lagrange-Gauss reduction on the matrix of row vectors:

$$\begin{bmatrix} B & \omega \\ 0 & x_0 \end{bmatrix}$$

one obtains in polynomial time a length-ordered basis $\{b_1, b_2\}$ of $\mathcal{L}$ satisfying $\|b_1\| = \lambda_1(\mathcal{L})$ and $\|b_2\| = \lambda_2(\mathcal{L})$, which enables to recover g and u. □

Using the same notations as in Section 3.2.1, $g = \prod_{i=1}^{\theta} g_i$ has approximate bit size $\theta \cdot \alpha$, while u has an approximate bit size $\gamma - \eta + n_h + \rho_\infty + \theta\alpha$. From the condition $g \cdot u < x_0/10$ of Proposition 3.3.1, we obtain, by dropping the term $\log_2(10)$, the simplified condition $\gamma - \eta + n_h + \rho_\infty + \theta\alpha + \theta\alpha < \gamma$. Writing as $\nu = \eta - n_h - \rho_\infty$ as in Section 3.2.1 for the number of bits that can be extracted during zero testing, the attack works under the condition

$$2\alpha\theta < \nu\,. \tag{3.10}$$

3.3.4 An extended attack against CLT13 with Independent Slots

We now describe a high-dimensional lattice reduction attack revealing the secret plaintext ring of CLT13 as for [GLW14], but improving the bound on θ derived in Equation (3.10).

Outline of our new attack. The main difference between our new attack and Section 3.3.3.1 is that this time we work with several messages instead of a single one, thus relying on high-dimensional lattice reduction instead of dimension 2. As a result, our new attack not only improves the bound (3.10), but also recovers multiples of the underlying plaintext messages,

which was not the case for [GLW14].

Let $\ell \geq 1$ be an integer. Assume that we have ℓ level-κ encodings $\{c_j\}_j$ of plaintext elements $m_j = (m_{j1}, \ldots, m_{jn})$ for $1 \leq j \leq \ell$, where each message has non-zero support of length θ. Without loss of generality, we can assume that the first θ components of each m_j are non-zero, that is, $m_{ji} = 0$ for all $\theta + 1 \leq i \leq n$ and all $1 \leq j \leq \ell$. We consider the zero-testing evaluations $\omega_j = p_{zt} \cdot c_j \bmod x_0$ of these encodings, which gives as previously:

$$\omega_j \equiv \sum_{i=1}^{\theta} h_i (r_{ji} + m_{ji}(g_i^{-1} \bmod p_i)) \frac{x_0}{p_i} + \sum_{i=\theta+1}^{n} h_i r_{ji} \frac{x_0}{p_i} \pmod{x_0}, \quad 1 \leq j \leq \ell$$

for integers $\{r_{ji}\}_{i,j}$. We can now rewrite the above equation as:

$$\omega_j \equiv \sum_{i=1}^{\theta} \alpha_i \cdot m_{ji} + R_j \pmod{x_0}, \quad 1 \leq j \leq \ell \tag{3.11}$$

for some integers $\{\alpha_i\}_i$, where for each ω_j, the integer R_j is significantly smaller than x_0.

We can view Equation (3.11) as an instance of a "noisy" hidden subset sum problem. In [NS99], the authors consider the hidden subset sum problem, as introduced in Definition 2.2.6. Given a positive integer N, and a vector $v = (v_1, \ldots, v_\ell) \in \mathbb{Z}^\ell$ with entries in $\{0, \ldots, N-1\}$, find integers $\alpha_1, \ldots, \alpha_n \in \{0, \ldots, N-1\}$ such that there exist vectors $x_1, \ldots, x_n \in \mathbb{Z}^\ell$ with entries in $\{0, 1\}$ satisfying:

$$v \equiv \alpha_1 x_1 + \alpha_2 x_2 + \cdots + \alpha_n x_n \pmod{N}$$

In our case, the weights $\alpha_1, \ldots, \alpha_n$ are hidden as in [NS99], but for each equation we have an additional hidden noisy term R_j. Moreover, the weights $\alpha_i = h_i(g_i^{-1} \bmod p_i)x_0/p_i$ have a special structure, instead of being random in [NS99]. Thanks to this special structure, using a variant of the orthogonal lattice approach from [NS99], we can recover the secret product $g = g_1 \cdots g_\theta$ and the plaintext elements $\{m_{ji}\}_i$ up to a scaling factor.

A preliminary result on lattices. Based on Hadamard's Inequality (Proposition 2.2.1), we prove the following simple lemma.

Lemma 3.3.2. *Let $1 \leq n \leq d$ be integers and let $L \subseteq \mathbb{Z}^d$ be a lattice of rank $n \geq 2$. Let $x_1, \ldots, x_{n-1} \in L$ be linearly independent. Then for every vector $y \in L$ not in the linear span of $x_1, \ldots, x_{n-1}$, one has $\|y\| \geq \mathrm{Vol}(L)/\prod_{i=1}^{n-1} \|x_i\|$.*

Proof. Since $x_1, \ldots, x_{n-1}, y \in L$ are linearly independent, they generate a rank-n sublattice L' of L and hence $\mathrm{Vol}(L) \leq \mathrm{Vol}(L')$ as $\mathrm{Vol}(L)$ divides $\mathrm{Vol}(L')$. By Hadamard's Inequality, $\mathrm{Vol}(L) \leq \mathrm{Vol}(L') \leq \|y\| \cdot \prod_{i=1}^{n-1} \|x_i\|$. The bound follows. □

3.3.5 Our first lattice-based attack

Setting. In this section, we describe our first attack based on a variant of the hidden subset-sum problem. We consider plaintext elements $m_1, \ldots, m_\ell \in \mathbb{Z}^n$ and write m_{ji} for the i-th entry of the j-th message, where $0 \leq m_{ji} < g_i$ for all $1 \leq i \leq n$ and $1 \leq j \leq \ell$. As previously,

we assume that $m_{ji} = 0$ for all $\theta+1 \leq i \leq n$. We write M for the matrix of row vectors m_j for $1 \leq j \leq \ell$; and we will denote its columns by $\hat{m}_i$ for $1 \leq i \leq n$, that is, $M = \left[\hat{m}_1 \,\middle|\, \cdots \,\middle|\, \hat{m}_n\right] \in \mathbb{Z}^{\ell \times n}$. By construction, the vectors $\{\hat{m}_i : \theta+1 \leq i \leq n\}$ are all zero. We also assume that for all $1 \leq i \leq \theta$, $\hat{m}_i \not\equiv 0 \pmod{g_i}$. For $1 \leq j \leq \ell$, we let c_j denote an encoding of m_j at the last level κ, satisfying $c_j \equiv z^{-\kappa}(r_{ji}g_i + m_{ji}) \pmod{p_i}$ for all $1 \leq i \leq n$, and where where $r_{ji} \in \mathbb{Z}$ are ρ_∞-bit integers. Letting $c = (c_j)_{1 \leq j \leq \ell}$, this gives a vector equation over $\mathbb{Z}^\ell$:

$$c \equiv z^{-\kappa}\left(g_i r_i + \hat{m}_i\right) \pmod{p_i}, \quad 1 \leq i \leq n \tag{3.12}$$

for $r_i = (r_{ji})_{1 \leq j \leq \ell}$. Let p_{zt} be the zero-testing parameter, as defined in Equation (3.2). From zero-testing we obtain the following equations:

$$\omega_j \equiv c_j \cdot p_{zt} \equiv \sum_{i=1}^{\theta} h_i m_{ji} (g_i^{-1} \bmod p_i)\frac{x_0}{p_i} + \sum_{i=1}^{n} h_i r_{ji}\frac{x_0}{p_i} \pmod{x_0},\ 1 \leq j \leq \ell$$

which can be rewritten as $\omega_j \equiv \sum_{i=1}^{\theta} \alpha_i m_{ji} + R_j \pmod{x_0}$, where we use the shorthand notations:

$$\alpha_i := h_i (g_i^{-1} \bmod p_i)\frac{x_0}{p_i} \quad , \quad 1 \leq i \leq \theta \tag{3.13}$$

$$R_j := \sum_{i=1}^{n} h_i r_{ji}\frac{x_0}{p_i} \quad , \quad 1 \leq j \leq \ell$$

As a vector equation, this reads:

$$\omega \equiv p_{zt} \cdot c \equiv \sum_{i=1}^{\theta} \alpha_i \hat{m}_i + R \pmod{x_0} \tag{3.14}$$

with $\omega = (\omega_j)_{1 \leq j \leq \ell}$; the vectors $\hat{m}_i$, for $1 \leq i \leq \theta$, are as above, and $R = (R_j)_{1 \leq j \leq \ell}$. The components of R have approximate bit size $\rho_R = \gamma - \eta + n_h + \rho_\infty$. Using, as in Section 3.2.1, $\nu := \eta - n_h - \rho_\infty$ as the number of bits that can be extracted, we have therefore $\rho_R = \gamma - \nu$. As explained above, Equation (3.14) is similar to an instance of the hidden subset sum problem, so we describe a variant of the orthogonal lattice attack considered in [NS99], which recovers the secret CLT13 plaintext ring and the hidden plaintexts $\{\hat{m}_i : 1 \leq i \leq \theta\}$, up to a scaling factor. For the sequel, we assume that $g_1, \ldots, g_\theta$ are distinct, and that $\gcd(g_i, h_i x_0/p_i) = 1$, for every $1 \leq i \leq \theta$.

The orthogonal lattice $\mathcal{L}$. To ω and x_0, we associate the lattice $\mathcal{L} := \mathcal{L}(\omega, x_0)$ of vectors $(Bu, v) \in \mathbb{Z}^{\ell+1}$, with $u \in \mathbb{Z}^\ell$ and $v \in \mathbb{Z}$, such that (u, v) is orthogonal to $(\omega, 1)$ modulo x_0, where $B \in \mathbb{Z}_{\geq 1}$ is a constant, which serves as scaling factor and that will be determined later in our analysis. Since $\mathcal{L}$ contains the sublattice $x_0\mathbb{Z}^{\ell+1}$, it has full-rank $\ell+1$. Note that this lattice is known (i.e. we can easily construct a basis for it) since ω and x_0 are given. Our attack is based on the fact that $\mathcal{L}$ contains a rank-ℓ sublattice $\mathcal{L}'$, generated by reasonably short vectors $\{(Bu_i, v_i) : 1 \leq i \leq \ell\}$ of $\mathcal{L}$, which can be used to reveal the secret product

$$g := \prod_{i=1}^{\theta} g_i \,.$$

More precisely, for every vector $(Bu, v) \in \mathcal{L}$, we obtain from Equation (3.14) and the linearity of the inner product:

$$\langle u, \omega \rangle + v \equiv \sum_{i=1}^{\theta} \alpha_i \langle u, \hat{m}_i \rangle + \langle u, R \rangle + v \equiv 0 \pmod{x_0}$$

and therefore, the vector $(\langle u, \hat{m_1} \rangle, \ldots, \langle u, \hat{m}_\theta \rangle, \langle u, R \rangle + v)$ is orthogonal modulo x_0 to the vector $a = (\alpha_1, \ldots, \alpha_\theta, 1)$. To obtain balanced components, we use another scaling constant $C \in \mathbb{Z}_{\geq 1}$ and consider the vector:

$$p_{u,v} := (C\langle u, \hat{m_1} \rangle, \ldots, C\langle u, \hat{m}_\theta \rangle, \langle u, R \rangle + v)$$

Following the original orthogonal lattice attack from [NS99], if a vector (Bu, v) from $\mathcal{L}$ is short enough, then the associated vector $p_{u,v} = (Cx, y)$ will also be short, and if (x, y) becomes shorter than a shortest non-zero vector orthogonal to a modulo x_0, we must have $p_{u,v} = 0$, which implies $\langle u, \hat{m}_i \rangle = 0$ for all $1 \leq i \leq \theta$. We will see that in our setting, because of the specific structure of the coefficients $\{\alpha_i\}_i$ from Equation (3.13), we only get $\langle u, \hat{m}_i \rangle \equiv 0 \pmod{g_i}$ for all $1 \leq i \leq \theta$. Therefore, by applying lattice reduction to $\mathcal{L}$, we expect to recover the lattice Λ of vectors u which are orthogonal to all $\hat{m}_i$ modulo g_i; since by assumption $\hat{m}_i \not\equiv 0 \pmod{g_i}$ for all $1 \leq i \leq \theta$, the lattice $\Lambda_i := \{u \in \mathbb{Z}^\ell : \langle u, \hat{m}_i \rangle \equiv 0 \pmod{g_i}\}$ has volume g_i, and since $g_1, \ldots, g_\theta$ are distinct primes, the lattice $\Lambda = \bigcap_{i=1}^{\theta} \Lambda_i$ has volume equal to $\prod_{i=1}^{\theta} g_i = g$. In particular, any basis for this lattice reveals g by computing its determinant.

The lattice $\mathcal{A}^\perp$. Henceforth, we must study the short vectors in the lattice of vectors orthogonal to a modulo x_0. More precisely, we consider the lattice $\mathcal{A}^\perp$ of vectors (Cx, y) in $\mathbb{Z}^{\theta+1}$, such that (x, y) is orthogonal to $a = (\alpha_1, \ldots, \alpha_\theta, 1)$ modulo x_0; therefore $p_{u,v} \in \mathcal{A}^\perp$. The lattice $\mathcal{A}^\perp$ has full-rank $\theta + 1$ and satisfies $\mathrm{Vol}(\mathcal{A}^\perp) = C^\theta x_0$. Namely, we have an abstract group isomorphism $\mathcal{A}^\perp \simeq (C\mathbb{Z})^\theta \oplus x_0\mathbb{Z}$, sending (Cx, y) to $(Cx, \langle x, a \rangle + y)$. As mentioned previously, the components $\{\alpha_i\}_i$ of a have a particular structure. Namely, by Equation (3.13), we have:

$$g_i \cdot \alpha_i \equiv h_i \cdot \frac{x_0}{p_i} \pmod{x_0}$$

for all $1 \leq i \leq \theta$. As a consequence, $\mathcal{A}^\perp$ contains the θ linearly independent short vectors $q_i = (0, \ldots, 0, Cg_i, 0, \ldots, 0, -s_i)$, where $s_i := h_i \cdot x_0/p_i$. Using $C := 2^{\rho_R - \alpha}$, we obtain the approximate bound $\|q_i\| \simeq C \cdot 2^\alpha$.

Let us now derive a condition on the Euclidean norm of $p_{u,v}$ so that it belongs to the sublattice of $\mathcal{A}^\perp$ generated by the short vectors $\{q_i : 1 \leq i \leq \theta\}$. From Lemma 3.3.2, if $\|p_{u,v}\| < \mathrm{Vol}(\mathcal{A}^\perp)/\prod_{i=1}^{\theta} \|q_i\|$, then $p_{u,v}$ must belong to the linear span generated by the $\{q_i\}_i$; further since by assumption, the g_i's are distinct primes and $\gcd(s_i, g_i) = 1$ for all $1 \leq i \leq \theta$, this implies that it must belong to the sublattice generated by the $\{q_i\}_i$. In that case, we have:

$$\langle u, \hat{m}_i \rangle \equiv 0 \pmod{g_i}, \quad 1 \leq i \leq \theta \tag{3.15}$$

From $\mathrm{Vol}(\mathcal{A}^\perp) = C^\theta x_0$ and $\|q_i\| \simeq C2^\alpha$, the condition $\|p_{u,v}\| < \mathrm{Vol}(\mathcal{A}^\perp)/\prod_{i=1}^{\theta} \|q_i\|$ gives the simplified condition:

$$\|p_{u,v}\| < 2^{\gamma - \alpha\theta} \tag{3.16}$$

Short vectors in $\mathcal{L}$. We now study the short vectors of $\mathcal{L}$; more precisely, we explain that $\mathcal{L}$ contains ℓ linearly independent short vectors of norm approximately $2^{\rho_R+\alpha\theta/\ell}$. We show that these vectors can be derived from the lattice Λ of vectors $u \in \mathbb{Z}^\ell$ satisfying Equation (3.15). This is a full-rank lattice of dimension ℓ and volume $g = \prod_{i=1}^{\theta} g_i$, with $g \simeq 2^{\alpha\theta}$. Therefore, we heuristically expect Λ to contains ℓ linearly independent vectors of norm approximately $(\mathrm{Vol}\Lambda)^{1/\ell} \simeq 2^{\alpha\theta/\ell}$. We show that from any short $u \in \Lambda$, we can generate a vector (u, v) with small v, and orthogonal to $(\omega, 1)$ modulo x_0, and consequently a short vector $(Bu, v) \in \mathcal{L}$. For this, we write $\langle u, \hat{m_i}\rangle = k_i g_i$ with $k_i \in \mathbb{Z}$, and we have:

$$\begin{aligned}\langle u, \omega\rangle + v &\equiv \sum_{i=1}^{\theta} \alpha_i \langle u, \hat{m_i}\rangle + \langle u, R\rangle + v \equiv \sum_{i=1}^{\theta} k_i \cdot g_i \cdot \alpha_i + \langle u, R\rangle + v \pmod{x_0} \\ &\equiv \sum_{i=1}^{\theta} k_i \cdot s_i + \langle u, R\rangle + v \pmod{x_0}\end{aligned}$$

It suffices to let $v := -\langle u, R\rangle - \sum_{i=1}^{\theta} k_i s_i$ to obtain $\langle u, \omega\rangle + v \equiv 0 \pmod{x_0}$; the vector (u, v) is then orthogonal to $(\omega, 1)$ modulo x_0, and thus $(Bu, v) \in \mathcal{L}$. We obtain $|v| \simeq \|u\| \cdot 2^{\rho_R}$; therefore letting $B := 2^{\rho_R}$, we get $\|(Bu, v)\| \simeq 2^{\rho_R}\|u\|$. In summary, the lattice $\mathcal{L}$ contains a sublattice $\mathcal{L}'$ of rank ℓ, generated by ℓ vectors of norm roughly $2^{\rho_R+\alpha\theta/\ell}$. That the recovered vectors are indeed linearly independent is the content of the following lemma.

Lemma 3.3.3. *Let $\{(Bu_j, v_j) : 1 \le j \le \ell+1\}$ be a basis of the lattice $\mathcal{L}$ and assume that the vectors $\{p_{u_j,v_j} : 1 \le j \le \ell\}$ belong to the sublattice of $\mathcal{A}^\perp$ generated by $\{q_i : 1 \le i \le \theta\}$. Then the vectors $\{u_j : 1 \le j \le \ell\}$ are $\mathbb{R}$-linearly independent.*

Proof. Let $\mathcal{B} = \{(Bu_j, v_j) : 1 \le j \le \ell+1\}$ be a basis of $\mathcal{L}$. By contradiction, assume that the vectors $\{u_j : 1 \le j \le \ell\}$ are linearly dependent. For every $1 \le j \le \ell$, we consider the vector p_{u_j,v_j} associated with (Bu_j, v_j). Since $\{p_{u_j,v_j} : 1 \le j \le \ell\}$ belong to the lattice generated by $\{q_i : 1 \le i \le \theta\}$, there exist integers $\{\beta_{ij}\}_{i,j}$ such that $p_{u_j,v_j} = \sum_{i=1}^{\theta} \beta_{ij} q_i$ for every $1 \le j \le \ell$. The definition of the vectors $\{q_i : 1 \le i \le \theta\}$ gives $p_{u_j,v_j} = (C\beta_{1j}g_1, \ldots, C\beta_{\theta j}g_\theta, -\sum_{i=1}^{\theta} \beta_{ij}s_i)$ for every $1 \le j \le \ell$, and from the definition of p_{u_j,v_j}, we conclude, by equalizing the components, the relations $\beta_{ij}g_i = \langle u_j, \hat{m_i}\rangle$ and $-\sum_{i=1}^{\theta} \beta_{ij}s_i = \langle u_j, R\rangle + v_j$ for every $1 \le j \le \ell, 1 \le i \le \theta$. Combining both relations leads to $v_j = -\sum_{i=1}^{\theta} \frac{s_i}{g_i}\langle u_j, \hat{m_i}\rangle - \langle u_j, R\rangle$, for $1 \le j \le \ell$. Thus, we see that if $\{u_j : 1 \le j \le \ell\}$ are linearly dependent over $\mathbb{R}$, then so are $\{(Bu_j, v_j) : 1 \le j \le \ell\}$, which contradicts the fact that $\mathcal{B}$ is a basis of $\mathcal{L}$. □

Recovering $g = \prod_{i=1}^{\theta} g_i$. In this part, we establish heuristic conditions on the parameters, under which lattice reduction is expected to reveal $g = \prod_{i=1}^{\theta} g_i$. We rely on the heuristic bounds from Section 2.2.2, and for an even simpler analysis, we omit the factor $\sqrt{(\ell+1)/(2\pi e)}$, which comes from the Gaussian Heuristic in Equation (4.5). We denote the *root Hermite factor* by $2^{\iota(\ell+1)}$, for some positive constant ι depending on the lattice reduction algorithm.

Therefore, by applying lattice reduction to $\mathcal{L}$, we expect that the first ℓ vectors $\{(Bu_j, v_j) : 1 \le j \le \ell\}$ of a reduced basis belong to $\mathcal{L}'$ and have norm approximately:

$$\|(Bu_j, v_j)\| \simeq 2^{\rho_R+\alpha\theta/\ell} \cdot 2^{\iota(\ell+1)}\ ,\ 1 \le j \le \ell\ . \tag{3.17}$$

With $C = 2^{\rho_R - \alpha}$ as suggested before, we have approximately $\|p_{u_i,v_i}\| \simeq \|(Bu_i, v_i)\|$ for all $1 \leq i \leq \ell$. From the condition given by Equation (3.16), we obtain that $u_i \in \Lambda$ if $\|p_{u_i,v_i}\| < 2^{\gamma - \alpha\cdot\theta}$; therefore combining with Equation (3.17) we get the approximate condition:

$$\rho_R + \frac{\alpha\theta}{\ell} + \iota(\ell+1) < \gamma - \alpha\theta \; .$$

Using $\rho_R = \gamma - \nu$ where ν is the number of bits that can be extracted from zero-testing, this condition becomes

$$\alpha\theta\left(1 + \frac{1}{\ell}\right) + \iota(\ell+1) < \nu \; . \tag{3.18}$$

In summary, when (3.18) is satisfied, we expect to recover a basis $\{u_i : 1 \leq i \leq \ell\}$ of Λ; then since $\mathrm{Vol}(\Lambda) = g = \prod_{i=1}^{\theta} g_i$, the absolute value of the determinant of the basis matrix reveals g. Let us add a remark concerning this line of attack. In fact, our algorithm only reveals a full-rank sublattice of Λ, instead of Λ. Thus we would only recover a non-trivial multiple of g, instead of g. However, in the results of our implementation, the revealed sublattice is always equal to Λ, and thus we recover g precisely.

From Equation (3.18), we observe that ℓ can be kept relatively small (say $\ell \simeq 10$), as larger values of ℓ would not significantly improve the bound; this implies that the lattice dimension $\ell + 1$ on which LLL is applied can be kept relatively small. Moreover for LLL, experiments show that $2^{\iota} \simeq 1.021$ so that ι is approximately 0.03, and therefore for such small values of ℓ, the term $\iota(\ell+1)$ is negligible. Thus we can use the simpler approximate bound for our attack:

$$\alpha\theta < \nu \tag{3.19}$$

Note that this gives a factor 2 improvement compared to the previous bound given by Equation (3.10), following the attack of [GLW14]. In the next subsection we will see how to get a much more significant improvement, with an asymptotic bound $\alpha\theta = O(\nu^2)$.

A proven variant. The above analysis is *heuristic* only, in that, it is subjected to heuristic bound for lattice reduction. Below we describe a proven variant that can recover a vector u such that $\langle u, \hat{m}_i \rangle \equiv 0 \pmod{g_i}$ for all $1 \leq i \leq \theta$, using LLL. Although we only recover a single vector u instead of a lattice basis, this will be enough when combined with the Cheon et al. attack to recover all secret parameters of CLT13 (see Section 3.2.2.2). As the proof of the proposition is lengthy, we postpone it to Section 3.3.11, as appendix.

Proposition 3.3.4. *Let $\ell, \theta \in \mathbb{Z}_{\geq 1}$, $x_0 \in \mathbb{Z}_{\geq 1}$ and let $g_i \in \mathbb{Z}_{\geq 2}$ be distinct α-bit prime numbers for $1 \leq i \leq \theta$ and some $\alpha \in \mathbb{Z}_{\geq 1}$. For $1 \leq i \leq \theta$, let $\alpha_i \in \mathbb{Z}$ such that $g_i \cdot \alpha_i \equiv s_i \pmod{x_0}$, for $s_i \in \mathbb{Z}$ satisfying $|s_i| \leq 2^{\rho_R}$, for some $\rho_R \in \mathbb{Z}_{\geq 1}$ and assume that $\gcd(g_i, s_i) = 1$. For $1 \leq i \leq \theta$, let $\hat{m}_i \in \mathbb{Z}^{\ell}$ be vectors with entries in $[0, g_i) \cap \mathbb{Z}$ such that $\hat{m}_i \not\equiv 0 \pmod{g_i}$, and let $R \in \mathbb{Z}^{\ell}$ such that $\|R\|_{\infty} \leq 2^{\rho_R}$. Let $\omega \in \mathbb{Z}^{\ell}$ such that $\omega \equiv \sum_{i=1}^{\theta} \alpha_i \hat{m}_i + R \pmod{x_0}$. Assume that*

$$\alpha\theta\left(1 + \frac{1}{\ell}\right) + \frac{\ell + \theta}{2} + \log_2(\ell\sqrt{\ell+1}\cdot\theta) + 4 < \log_2(x_0) - \rho_R \; . \tag{3.20}$$

Given the integers $\ell, \theta, \rho_R, x_0$ and the vector ω, one can recover in polynomial time a vector $u \in \mathbb{Z}^{\ell}$ such that $\langle u, \hat{m}_i \rangle \equiv 0 \pmod{g_i}$ for all $1 \leq i \leq \theta$, satisfying

$$\|u\| \leq 2^{\ell/2}\sqrt{\ell(\ell+1)}\left(\prod_{i=1}^{\theta} g_i\right)^{1/\ell} \; .$$

We remark that by replacing $\log_2(x_0) - \rho_R$ by $\gamma - \rho_R = \nu$, we recover, up to additional logarithmic terms, the approximate bound established in Equation (3.18).

3.3.6 Extended Orthogonal Lattice Attack

We shall now describe an extended attack that significantly improves the bound on θ established in Equation (3.19). Let $\ell, d \geq 1$ be integers. As previously, we will assume that an attacker has encodings $\{c_j\}_j$ of plaintext elements $m_j = (m_{j1}, \ldots, m_{jn})$ for $1 \leq j \leq \ell$, where only the first θ components of each m_j are non-zero, that is, $m_{ji} = 0$ for $\theta + 1 \leq i \leq n$. However, we assume that these encodings are at level $\kappa - 1$, and that we also have an additional set of d level-1 encodings $\{c'_k : 1 \leq k \leq d\}$ of plaintext elements $x_k = (x_{k1}, \ldots, x_{kn})$ for $1 \leq k \leq d$. By computing the top-level κ product encodings, we can therefore obtain the following zero-testing evaluations:

$$\omega_{jk} \equiv (c_j \cdot c'_k) \cdot p_{zt} \equiv \sum_{i=1}^{\theta} h_i m_{ji} x_{ki} (g_i^{-1} \bmod p_i) \frac{x_0}{p_i} + \sum_{i=1}^{n} h_i r_{jki} \frac{x_0}{p_i} \pmod{x_0} \tag{3.21}$$

for some integers r_{jki}. Since every encoding $\{c_j\}_j$ encodes a message with non-zero support of length θ, the product encodings $\{c_j c'_k\}_{j,k}$ maintain their zero slots. Note that the same remains valid if the encodings $\{c_j\}_j$ are at even lower levels, because they can be raised to level $\kappa - 1$ without removing their zero slots. As previously, we rewrite Equation (3.21) as:

$$\omega_{jk} \equiv \sum_{i=1}^{\theta} \alpha_{ik} m_{ji} + R_{jk} \pmod{x_0}$$

where we have written $\alpha_{ik} = h_i x_{ki} (g_i^{-1} \bmod p_i) \frac{x_0}{p_i}$ for $1 \leq i \leq \theta$ and $1 \leq k \leq d$, and $R_{jk} = \sum_{i=1}^{n} h_i r_{jki} x_0 / p_i$ for all $1 \leq j \leq \ell$ and $1 \leq k \leq d$. As before, for $1 \leq i \leq \theta$, we denote by $\hat{m}_i \in \mathbb{Z}^\ell$ the vector with components m_{ji} for $1 \leq j \leq \ell$, and similarly ω_k and R_k the corresponding vectors in $\mathbb{Z}^\ell$. We assume that $\hat{m}_i \not\equiv 0 \pmod{g_i}$ for all i. The previous equation can then be rewritten as:

$$\omega_k \equiv \sum_{i=1}^{\theta} \alpha_{ik} \hat{m}_i + R_k \pmod{x_0} \tag{3.22}$$

The difference with Equation (3.14) from our first lattice attack is that the vectors $\{\hat{m}_i\}_i$ now satisfy d equations for $1 \leq k \leq d$, instead of a single equation, as was the case in Section 3.3.5. With more constraints on the vectors $\{\hat{m}_i\}_i$, we can therefore break the countermeasure from [GLW14] for much higher values of θ. In order to derive a condition on the parameters, we proceed as previously. Namely, the lattices considered in Section 3.3.5 now admit natural higher-dimensional analogues.

The orthogonal lattice $\mathcal{L}$. As previously, for a scaling constant $B \in \mathbb{Z}_{\geq 1}$, we consider the lattice $\mathcal{L}$ of vectors $(Bu, v) \in \mathbb{Z}^{\ell+d}$, with $u \in \mathbb{Z}^\ell$ and $v \in \mathbb{Z}^d$, such that (u, v) is orthogonal to the d vectors $\{(\omega_k, e_k) : 1 \leq k \leq d\}$ modulo x_0, where $e_k \in \mathbb{Z}^d$ is the kth unit vector for $1 \leq k \leq d$. This gives for all $1 \leq k \leq d$ and all $(Bu, v) \in \mathcal{L}$, writing $v = (v_1, \ldots, v_d)$:

$$\langle u, \omega_k \rangle + v_k \equiv \sum_{i=1}^{\theta} \alpha_{ik} \langle u, \hat{m}_i \rangle + \langle u, R_k \rangle + v_k \equiv 0 \pmod{x_0}$$

and therefore $(\langle u, \hat{m_1}\rangle, \ldots, \langle u, \hat{m}_\theta\rangle,\ \langle u, R_1\rangle + v_1, \ldots, \langle u, R_d\rangle + v_d)$ is orthogonal over $\mathbb{Z}/x_0\mathbb{Z}$ to the coefficient-vectors $a_k = (\alpha_{1k}, \ldots, \alpha_{\theta k}, e_k)$, for $1 \leq k \leq d$. Again, using a scaling constant $C \in \mathbb{Z}_{\geq 1}$, we let $p_{u,v} = (C\langle u, \hat{m_1}\rangle, \ldots, C\langle u, \hat{m}_\theta\rangle,\ \langle u, R_1\rangle + v_1, \ldots, \langle u, R_d\rangle + v_d)$.

The lattice $\mathcal{A}^\perp$. To bound the norm of the vector $p_{u,v}$, we must study the short vectors in the lattice of vectors orthogonal to the vectors a_k modulo x_0 (instead of single vector a). As previously, we consider the lattice $\mathcal{A}^\perp$ of vectors $(Cx, y) \in \mathbb{Z}^{\theta+d}$ such that (x, y) is orthogonal to the d vectors $\{a_k : 1 \leq k \leq d\}$ modulo x_0; therefore $p_{u,v} \in \mathcal{A}^\perp$. The lattice $\mathcal{A}^\perp$ has full-rank $\theta + d$ and volume $C^\theta x_0^d$. As previously, the components $\{\alpha_{ik}\}_{ik}$ in the vectors $\{a_k\}_k$ have a special structure, since they satisfy the congruence relations

$$g_i \cdot \alpha_{ik} \equiv h_i \cdot x_{ik} \cdot \frac{x_0}{p_i} \pmod{x_0}$$

for all $1 \leq i \leq \theta$ and $1 \leq k \leq d$. Therefore, letting $s_{ik} = h_i \cdot x_{ik} \cdot x_0/p_i$, the lattice $\mathcal{A}^\perp$ contains the θ short vectors $q_i = (0, \ldots, 0, Cg_i, 0, \ldots, 0, -s_{i1}, \ldots, -s_{id})$ for $1 \leq i \leq \theta$. Using $C = 2^{\rho_R - \alpha}$, we get as previously $\|q_i\| \simeq C \cdot 2^\alpha$.

We now derive a bound on $\|p_{u,v}\|$ so that $p_{u,v}$ belongs to the sublattice generated by the θ vectors $\{q_i : 1 \leq i \leq \theta\}$. We expect a reduced basis of $\mathcal{A}^\perp$ to have the first θ vectors with approximately the same norm as the vectors $\{q_i : 1 \leq i \leq \theta\}$, and to have the last d vectors with norm U satisfying $(C \cdot 2^\alpha)^\theta \cdot U^d \simeq \mathrm{Vol}(\mathcal{A}^\perp)$. Using $\mathrm{Vol}(\mathcal{A}^\perp) = C^\theta x_0^d$, this gives $U \simeq x_0/2^{\alpha\theta/d}$. This implies that, heuristically, if $\|p_{u,v}\| < U$, then $p_{u,v}$ must belong to the sublattice generated by the θ vectors $\{q_i : 1 \leq i \leq \theta\}$. As previously, in that case we have that for all $1 \leq i \leq \theta$:

$$\langle u, \hat{m}_i\rangle \equiv 0 \pmod{g_i}\,. \tag{3.23}$$

Short vectors in $\mathcal{L}$. Let us now study the short vectors of $\mathcal{L}$; as previously, we can argue that $\mathcal{L}$ contains ℓ linearly independent short vectors of Euclidean norm approximately $2^{\rho_R + \alpha\theta/\ell}$, which can be derived from the the lattice Λ of vectors $u \in \mathbb{Z}^\ell$ satisfying Equation (3.23). Similarly to the previous case, because the lattice Λ heuristically contains ℓ linearly independent vectors of Euclidean norm approximately $(\det \Lambda)^{1/\ell} \simeq 2^{\alpha\theta/\ell}$, the lattice $\mathcal{L}$ contains ℓ linearly independent vectors of norm roughly $2^{\rho_R + \alpha\theta/\ell}$. Therefore, by applying lattice reduction to the lattice $\mathcal{L}$, we expect that the first ℓ vectors $\{(Bu_i, v_i) : 1 \leq i \leq \ell\}$ of the basis have norm roughly:

$$\|(Bu_i, v_i)\| \simeq B \cdot 2^{\alpha\theta/\ell} \cdot 2^{\iota(\ell+d)}$$

where $2^{\iota(\ell+d)}$ is the Hermite factor. Putting $B = 2^{\rho_R}$ and $C = 2^{\rho_R - \alpha}$, this now asymptotically amounts to $\|p_{u_i,v_i}\| \simeq \|(Bu_i, v_i)\|$. From the condition $\|p_{u_i,v_i}\| < U$, we get the condition $\rho_R + \frac{\alpha\theta}{\ell} + \iota(\ell + d) < \gamma - \frac{\alpha\theta}{d}$, which further gives, using $\rho_R = \gamma - \nu$:

$$\alpha\theta \cdot \left(\frac{1}{\ell} + \frac{1}{d}\right) + \iota(\ell + d) < \nu \tag{3.24}$$

Remark that with $d = 1$, Equation (3.24) gives Equation (3.18). Since, as bivariate function of ℓ, d, the left-hand side of (3.24) is concave and symmetric in both variables ℓ and d, the optimum is to take $\ell = d$. With this choice, we obtain the bound:

$$\frac{2\alpha\theta}{\ell} + 2\iota\ell < \nu \tag{3.25}$$

Recovering $g = \prod_{i=1}^{\theta} g_i$. When the above condition on the parameters is satisfied, then, as previously, we heuristically expect to recover a basis $\{u_i : 1 \leq i \leq \ell\}$ of the lattice Λ. As $\mathrm{Vol}(\Lambda) = g = \prod_{i=1}^{\theta} g_i$, the absolute value of the determinant of the basis matrix reveals then reveals the secret product g. In particular, it follows that the attack requires $\ell > 2\alpha\theta/\nu$, and we must therefore have the following bound on ι:

$$\iota < \frac{\nu^2}{4\alpha\theta}$$

Heuristically, achieving a Hermite factor of $2^{\iota 2\ell}$ requires complexity $2^{\Theta(1/\iota)}$ using BKZ reduction with block-size $\beta = \Theta(1/\iota)$, [HPS11b]. The attack has therefore complexity $2^{\Theta(\alpha\theta/\nu^2)}$; the attack has therefore (heuristic) polynomial-time complexity under the condition:

$$\alpha\theta = O(\nu^2)$$

This gives a significant improvement of our previous bound given by Equation (3.19). Conversely, one expects that the attack is prevented under the condition:

$$\theta = \omega\left(\frac{\nu^2}{\alpha}\log\lambda\right) \tag{3.26}$$

In Section 3.3.8 we provide a concrete set of parameters for CLT13 multilinear maps with independent slots. We will see that Condition (3.26) requires a much higher value for θ than the condition $2\theta\alpha \geq \nu$ for preventing the attack from [GLW14], described in Section 3.3.3. For instance, for $\lambda = 80$ bits of security, the bound $2\theta\alpha \geq \nu$ already holds for $\theta = 2$, while a concrete application of Condition (3.26) requires $\theta \geq 1789$.

Analogy of the attacks. We remark that our extended attacks share similarities with the 2-dimensional attack from Section 3.3.3. For $\ell, d \in \mathbb{Z}_{\geq 1}$, our extended lattice attack works by reducing the $(\ell + d)$-dimensional lattice

$$\mathcal{L}_{(\ell,d)} = \{(Bu, v) \in \mathbb{Z}^{\ell} \times \mathbb{Z}^{d} : \langle (u, v), (\omega_k, e_k)\rangle \equiv 0 \pmod{x_0}, 1 \leq k \leq d\} ,$$

where $B \in \mathbb{Z}_{\geq 1}$ is fixed. With this notation, the three attacks work by reducing the lattices $\mathcal{L}_{(1,1)}$, $\mathcal{L}_{(\ell,1)}$ and $\mathcal{L}_{(\ell,d)}$, respectively. Note that $\mathcal{L}_{(1,1)}$ is the lattice $\{(Bu, v) \in \mathbb{Z}^2 : u\omega + v \equiv 0 \pmod{x_0}\}$. For the extended attacks, the $\ell \times \ell$ top-left submatrix of a reduced basis of $\mathcal{L}_{(\ell,d)}$ (divided by B) has determinant $\pm g$. Note that this coincides with the 2-dimensional case $\ell = d = 1$: the first entry (divided by B) of the first vector in a reduced basis equals $\pm g$ (i.e. a "1×1 submatrix" of determinant $\pm g$). As such, our higher-dimensional attacks are consistent generalizations of the 2-dimensional attack.

Summary. We have described a lattice-based attack, which under the condition $\alpha\theta = O(\nu^2)$, and given as input a collection of encodings (or products of encodings) of messages with non-zero support of length θ, outputs the secret plaintext ring of CLT13. More precisely, our extended lattice attack with the improved bound $\alpha\theta = O(\nu^2)$ can be described as follows. We provide in [CN19b] the source code in Sage [S+17].

Algorithm 3 Lattice Attack against CLT13 with Independent Slots

Parameters: Integers $\ell, d \geq 1$, $0 < \theta \leq n$, and the CLT13 parameters

Input: Sets of level-κ encodings $\{c_j \cdot c'_k \bmod x_0 : 1 \leq j \leq \ell, 1 \leq k \leq d\}$ where c_j encodes a message of non-zero support of length θ

Output: A basis of the lattice of vectors orthogonal to $\{\hat{m}_i\}_i$ modulo $\{g_i\}_i$, and $g = \prod_{i=1}^{\theta} g_i$

1: **for** $1 \leq k \leq d$ **do**
2: Compute the vector $\omega_k \in \mathbb{Z}^\ell$ with $(\omega_k)_j = c_j \cdot c'_k \cdot p_{zt} \bmod x_0$
3: **end for**
4: Let $B = 2^{\rho_R}$ and compute a LLL-reduced basis of the lattice $\mathcal{L}_{(\ell,d)} \subseteq \mathbb{Z}^{\ell+d}$ of vectors $\{(Bu, v) \in \mathbb{Z}^\ell \times \mathbb{Z}^d : \langle (u, v), (\omega_k, e_k) \rangle \equiv 0 \pmod{x_0}, 1 \leq k \leq d\}$, where $e_k \in \mathbb{Z}^d$ is the kth unit vector for $1 \leq k \leq d$. Denote by $\{(Bu_j, v_j) : 1 \leq j \leq \ell + d\}$ the LLL-reduced basis
5: Form the matrix $U \in \mathbb{Z}^{\ell \times \ell}$ of vectors $\{u_j : 1 \leq j \leq \ell\}$
6: Return U and $|\det(U)|$

Variant with multiple p_{zt}. In many concrete constructions based on composite order multilinear maps, intermediate-level encodings of almost zero plaintexts are not necessarily available. We refer to Section 3.3.10 for the application of our attacks to concrete constructions. In order to get around this assumption, we consider a variant of the above attack, where we have multiple zero-testing elements p_{zt} instead of a single one. Namely, as described in [CLT13], in order to get a proper zero-testing procedure, one needs to use a vector of n elements p_{zt}. We denote by $p_{zt,k}$ for $1 \leq k \leq n$ those zero-testing elements:

$$p_{zt,k} = \sum_{i=1}^{n} h_{ik} z^\kappa (g_i^{-1} \bmod p_i) \frac{x_0}{p_i} \bmod x_0$$

for corresponding integers h_{ik}. As previously, we assume that we have encodings $\{c_j\}_j$ of plaintext elements $m_j = (m_{j1}, \ldots, m_{jn})$ for $1 \leq j \leq \ell$, where only the first θ components of each m_j are non-zero, that is, $m_{ji} = 0$ for $\theta + 1 \leq i \leq n$. We can now assume that these encodings are at the last level κ. Thanks to the multiple zero-testing elements, we can therefore obtain the following zero-testing evaluations:

$$\omega_{jk} \equiv c_j \cdot p_{zt,k} \equiv \sum_{i=1}^{\theta} h_{ik} m_{ji} (g_i^{-1} \bmod p_i) \frac{x_0}{p_i} + \sum_{i=1}^{n} h_{ik} r_{jki} \frac{x_0}{p_i} \pmod{x_0}$$

for some integers $\{r_{jki}\}$, which is similar to Equation (3.21) with $h_{ik} = h_i \cdot x_{ki}$. Therefore, the same attack applies and the secret $g = \prod_{i=1}^{\theta} g_i$ can be recovered in heuristic polynomial-time under the condition $\alpha\theta = O(\nu^2)$.

3.3.7 Revealing information about the plaintext elements

We show that our attack not only reveals the secret CLT13 plaintext ring, but also information about the secret plaintext elements $\{\hat{m}_i : 1 \leq i \leq \theta\}$. Namely, the orthogonal lattice attack also constructs a matrix U of rows $\{u_j : 1 \leq j \leq \ell\}$ orthogonal to the vectors $\{\hat{m}_i : 1 \leq i \leq \theta\}$ modulo the primes $\{g_i\}_i$ (i.e. a basis of the lattice Λ, following the notation of the previous section). We now explain how one can use the matrix U in order to recover scalar multiples

of the plaintext vectors $\{\hat{m}_i : 1 \leq i \leq \theta\}$. We describe two algorithms: one by computing the factorization of g and hence with sub-exponential running time, and one without computing the factorization of g and polynomial running time.

A sub-exponential-time algorithm. Our first algorithm relies on factoring $g = \prod_{i=1}^{\theta} g_i$ to reveal the primes $\{g_i\}_i$ in first place; this is feasible if these primes are small enough. For example, for the concrete parameters proposed in [CLT13], the $\{g_i\}_i$ are 80-bit primes; which makes the factorization of g straightforward. From the knowledge of a basis matrix U of the lattice of vectors u with $\langle u, \hat{m}_i \rangle \equiv 0 \pmod{g_i}$ for all $1 \leq i \leq \theta$, it suffices to compute the $\mathbb{Z}/g_i\mathbb{Z}$-kernel of the $\ell \times \ell$ matrix U_{g_i}, by which we denote the reduction of $U \bmod g_i$; assuming that $\hat{m}_i \not\equiv 0 \pmod{g_i}$, we have that $\ker(U_{g_i})$ has dimension 1 over over $\mathbb{F}_{g_i}$ and therefore, we recover a non-trivial multiple $\lambda_i \hat{m}_i$ of the original messages $\hat{m}_i$ modulo g_i, for $1 \leq i \leq \theta$. Using the ECM factorization algorithm [Len87], the factorization of $g = \prod_{i=1}^{\theta} g_i$ can be computed in time $\exp(c\sqrt{\alpha \ln \alpha})$ for some positive constant c and where α is the bit size of each prime g_i. The cost of this attack is clearly dominated by the factorization of g, which gives a sub-exponential time algorithm.

Algorithm 4 Compute (with factoring) multiples of hidden CLT13 plaintexts

Parameters: Integer $\ell \geq 1$, $0 < \theta \leq n$ and the CLT13 parameters

Input: $g = \prod_{i=1}^{\theta} g_i$ and a basis matrix $U \in \mathbb{Z}^{\ell \times \ell}$ of the lattice of vectors u with $\langle u, \hat{m}_i \rangle \equiv 0 \pmod{g_i}$ with $\hat{m}_i \not\equiv 0 \pmod{g_i}$, for all $1 \leq i \leq \theta$

Output: vectors $\lambda_i \hat{m}_i \in \mathbb{Z}^{\ell}$ with $\lambda_i \in \mathbb{Z}$, $\lambda_i \not\equiv 0 \pmod{g_i}$ for $1 \leq i \leq \theta$

1: Factor g using the ECM algorithm and return $\{g_i : 1 \leq i \leq \theta\}$
2: **for** $1 \leq i \leq \theta$ **do**
3: Form the matrix $U_{g_i} \in (\mathbb{Z}/g_i\mathbb{Z})^{\ell \times \ell}$, a representative of U modulo g_i
4: Compute and return a basis vector of the $\mathbb{Z}/g_i\mathbb{Z}$-kernel of U_{g_i}
5: **end for**

A polynomial-time algorithm. Alternatively, we may avoid the factorization of g and obtain a polynomial time algorithm which reveals a non-trivial multiple of a vector $\hat{m}$, such that $\hat{m} \equiv \hat{m}_i \pmod{g_i}$ for all $1 \leq i \leq \theta$. In other terms, $\hat{m}$ is the image of the vectors $\{\hat{m}_i\}$ under the componentwise Chinese Remainder isomorphism $(\prod_{i=1}^{\theta} \mathbb{Z}/g_i\mathbb{Z})^{\ell} \to (\mathbb{Z}/g\mathbb{Z})^{\ell}$. To do so, we can compute the integer right kernel of the matrix $U(g) := [U \mid g \cdot 1_{\ell}]$, where 1_{ℓ} denotes the identity matrix in dimension ℓ. We obtain the following corresponding proposition.

Proposition 3.3.5. *Let $\ell, \theta \in \mathbb{Z}_{\geq 1}$. Let $g_1, \ldots, g_\theta$ be distinct prime numbers. For $1 \leq i \leq \theta$, let $\hat{m}_i \in \mathbb{Z}^{\ell} \cap [0, g_i)^{\ell}$ be vectors such that $\hat{m}_i \not\equiv 0 \pmod{g_i}$. Let $\{u_j : 1 \leq j \leq \ell\}$ be a basis of the lattice of vectors $u \in \mathbb{Z}^{\ell}$ such that $\langle u, \hat{m}_i \rangle \equiv 0 \pmod{g_i}$ for all $1 \leq i \leq \theta$. Then, given $g = \prod_{i=1}^{\theta} g_i$ and the vectors $\{u_j : 1 \leq j \leq \ell\}$, one can compute in polynomial time a vector $\lambda \cdot \hat{m} \in \mathbb{Z}^{\ell} \cap [0, g)^{\ell}$ with $\gcd(\lambda, g) = 1$, such that $\hat{m} \equiv \hat{m}_i \pmod{g_i}$ for all $1 \leq i \leq \theta$.*

Proof. By the Chinese Remainder Theorem, there exists a unique vector $\hat{m} \in \mathbb{Z}^{\ell} \cap [0, g)^{\ell}$ such that $\hat{m} \equiv \hat{m}_i \pmod{g_i}$ for all $1 \leq i \leq \theta$. Consider the composition of maps $\mathbb{Z}^{\ell} \xrightarrow{\pi} (\mathbb{Z}/g\mathbb{Z})^{\ell} \xrightarrow{\phi} \mathbb{Z}/g\mathbb{Z}$, where π is reduction modulo g and ϕ sends u to $\langle u, \hat{m} \rangle$. By the Chinese Remainder Theorem, the map ϕ corresponds to a vector of maps $\phi = (\phi_1, \ldots, \phi_\theta) : (\prod_i \mathbb{Z}/g_i\mathbb{Z})^{\ell} \to \prod_i \mathbb{Z}/g_i\mathbb{Z}$

with components $\phi_i : (\mathbb{Z}/g_i\mathbb{Z})^\ell \to \mathbb{Z}/g_i\mathbb{Z}$ for $1 \le i \le \theta$. Let $1 \le i \le \theta$; since $\hat{m}_i \not\equiv 0 \pmod{g_i}$, the map ϕ_i is surjective with kernel $\ker(\phi_i) = \mathrm{im}(U_{g_i})$ where $U_{g_i} \equiv U \pmod{g_i}$ (with the congruence being understood entrywise). Since g_i is prime, $\ker(\phi_i) = \mathrm{im}(U_{g_i})$ is a $\mathbb{Z}/g_i\mathbb{Z}$-vector space of dimension $\ell - 1$. It follows that the kernel of U_{g_i} has dimension 1 over $\mathbb{Z}/g_i\mathbb{Z}$. This holds for all $1 \le i \le \theta$, so by the Chinese Remainder Theorem, the kernel of U_g, with U_g being the matrix U reduced modulo g, is a free $\mathbb{Z}/g\mathbb{Z}$-module of rank 1, generated by $\hat{m}$. In particular, there exists $k \in \mathbb{Z}^\ell$ such that $(\hat{m}, k)$ belongs to the $\mathbb{Z}$-kernel of the matrix $U(g) = [U|g \cdot 1_\ell]$. The integer kernel of this matrix can be computed in polynomial time from g and U and the left $\ell \times \ell$ submatrix of the Hermite normal form of the basis of the $\mathbb{Z}$-kernel gives in the first row a vector $\lambda\hat{m}$ with $\lambda \in (\mathbb{Z}/g\mathbb{Z})^\times$. □

Algorithm 5 Compute (without factoring) a multiple of a hidden plaintext vector with $\mathbb{Z}/g_i\mathbb{Z}$-residues the hidden CLT13 plaintexts

Parameters: Integer $\ell \ge 1$, $0 < \theta \le n$ and the CLT13 parameters

Input: $g = \prod_{i=1}^{\theta} g_i$ and a basis matrix $U \in \mathbb{Z}^{\ell \times \ell}$ of the lattice of vectors u with $\langle u, \hat{m}_i \rangle \equiv 0 \pmod{g_i}$ with $\hat{m}_i \not\equiv 0 \pmod{g_i}$, for all $1 \le i \le \theta$

Output: a vector $\lambda\hat{m} \in \mathbb{Z}^\ell$ with $\lambda \in \mathbb{Z}$ such that $\gcd(\lambda, g) = 1$ and $\hat{m} \equiv \hat{m}_i \pmod{g_i}$ for $1 \le i \le \theta$

1: Form the matrix $U(g) := [U \mid g \cdot 1_\ell] \in \mathbb{Z}^{\ell \times 2\ell}$
2: Compute the $\mathbb{Z}$-right kernel of $U(g)$ in Hermite normal form. Denote by $U^\perp(g) \in \mathbb{Z}^{\ell \times 2\ell}$ this matrix.
3: Return the first row of the $\ell \times \ell$ matrix $U^\perp(g)$

3.3.8 Concrete parameters and practical experiments

Concrete parameters. Below we provide concrete parameters for CLT13 multilinear maps with independent slots, for various values of the security parameter λ. We start from the same concrete parameters as provided in [CLT13]; we assume that the encoding noise is set so that the number of extracted bits is $\nu = 2\lambda + 12$ and we take $\alpha = \lambda$. We then provide the minimum value of θ that ensures the same level of security against lattice attacks; see Table 3.1. As in [CLT13], the goal is to ensure that the best attack takes at least 2^λ clock cycles.

While in Table 3.1 the number of independent slots $n_{\mathsf{slots}} = \lfloor n/\theta \rfloor$ appears to be relatively small, it is always possible to increase the number of independent slots by increasing the value of n.

Instantiation	λ	n	η	$\gamma = n \cdot \eta$	ν	θ	n_{slots}
Small	52	1080	1981	$2.1 \cdot 10^6$	116	540	2
Medium	62	2364	2055	$4.9 \cdot 10^6$	136	1182	2
Large	72	8250	2261	$18.7 \cdot 10^6$	156	1472	5
Extra	80	26115	2438	$63.7 \cdot 10^6$	172	1789	14

Table 3.1: Concrete parameters for CLT13 multilinear maps with independent slots

	θ	α	ν	$\ell = d$	dimension	running time
Basic attack	1	80	172	1	2	ε
Extended attack	2	80	172	2	4	ε
Extended attack	40	80	172	39	78	10 s
Extended attack	100	80	172	100	200	11 min
Extended attack	160	80	172	163	326	11 hours

Table 3.2: Running time of our LLL-based attack, as a function of the parameter θ, for the "Extra" parameters of CLT13

Practical experiments. We have run our extended attack from Section 3.3.6 with the "Extra" parameters of CLT13 from Table 3.1, for increasing values of θ. Note that for such parameters, the original attack from [GLW14] only applies for $\theta = 1$. To improve efficiency, we give as input to LLL a truncated matrix basis, where we keep only the ν most significant bits. Table 3.2 shows that our attack works in practice for much larger values of θ than the original attack from [GLW14], which can only work for $\theta = 1$. The line "Basic attack" corresponds to the attack from [GLW14], while the lines "Extended attack" refer to our attack from Section 3.3.6. The column "dimension" is the lattice dimension in which our algorithm is run, and we have chosen the symmetric case $\ell = d$ as suggested in Section 3.3.6. The quantity ε in the running time refers to a negligible time. The source code in SageMath [S$^+$17] is available in [CN19b].

3.3.9 Application to the Cheon et al. Attack

In this section, we consider the Cheon et al. attack against CLT13, recalled in Section 3.2.2.2. We show how to adapt this attack to the setting of CLT13 with independent slots when combined with our lattice-based attack: we assume that no encodings of zero are available to the attacker (otherwise the Cheon et al. attack would apply immediately), but as previously, the attacker can obtain low-level encodings where only $0 < \theta \leq n$ components of the plaintext are non-zero. In particular, this contributes to a cryptanalysis of CLT13 multilinear maps when no encodings of zero are available beforehand; for instance, this was considered as an open problem in [CFL$^+$16].

3.3.9.1 Adaptation of the Cheon et al. attack to our cryptanalysis

We now show how to adapt the Cheon et al. attack when no encodings of zero are available, but the attacker can obtain low-level encodings where only θ components of the underlying plaintexts are non-zero. The attack goes in two steps. From the given encodings, the attacker first generates encodings of zero by running the orthogonal lattice attack from Section 3.3.6. In second place, the attacker applies the original Cheon et al. attack on these newly created encodings of zero, to reveal the prime factors $\{p_i : 1 \leq i \leq n\}$ of x_0.

We work in the following setting, with $\kappa = 4$. Let $\ell \geq 1$ be an integer and consider a set

$$\mathcal{Y} = \{y_j : 1 \leq j \leq \ell\}$$

of level-one encodings of certain messages $m_1, \ldots, m_\ell \in \mathbb{Z}^n$, where only the first θ components of each m_j are non-zero. Moreover, we consider as in Section 3.2.2.2, three sets

$\mathcal{A} = \{\alpha_j : 1 \leq j \leq n\}$, $\mathcal{B} = \{\beta_1, \beta_2\}$ and $\mathcal{C} = \{\gamma_k : 1 \leq k \leq n\}$ of level-one encodings of non-zero messages.

First step: orthogonal lattice attack. The first step of the attack consists in generating an encoding of zero from the non-zero encodings in the public set $\mathcal{Y}$, which we achieve by our lattice attack. Denote by $y \in \mathbb{Z}^\ell$ the vector of encodings $(y_1, \ldots, y_\ell)$ constructed from $\mathcal{Y}$. We show that the orthogonal lattice attack from Section 3.3.5 can compute a short vector $u \in \mathbb{Z}^\ell$ such that $y' := \langle u, y \rangle = \sum_{j=1}^{\ell} u_j y_j$ is a level-one encoding of zero. For all $1 \leq j \leq \ell$, we write

$$y_j \equiv \frac{r_{ji} \cdot g_i + m_{ji}}{z} \pmod{p_i}, \quad 1 \leq i \leq n, \tag{3.27}$$

with the usual CLT13 notation from Section 3.2.1, and where $m_{ji} = 0$ for $\theta + 1 \leq i \leq n$. Note that our orthogonal lattice attack from Section 3.3.5 uses level-κ encodings; therefore it can be applied on level-κ encodings of the form $e_j = y_j \cdot \alpha_1 \cdot \beta_1 \cdot \gamma_1 \bmod x_0$ (remember that we use $\kappa = 4$) for level-one encodings $(\alpha_1, \beta_1, \gamma_1) \in \mathcal{A} \times \mathcal{B} \times \mathcal{C}$. By combining with Equation (3.27) this gives

$$e_j \equiv \frac{r'_{ji} \cdot g_i + m_{ji} \cdot x_i}{z^\kappa} \pmod{p_i}, \quad 1 \leq i \leq n,$$

for some integers $\{r'_{ji}\}_{i,j}$ and where x_i is the i-th component of the plaintext corresponding to the product encoding $\alpha_1 \cdot \beta_1 \cdot \gamma_1$. Clearly, since the messages $\{m_j : 1 \leq j \leq \ell\}$ have non-zero support of length θ, the messages $\{(m_{ji} \cdot x_i)_{1 \leq i \leq n} : 1 \leq j \leq \ell\}$ have non-zero support of length at most θ. Therefore, applying the orthogonal lattice attack from Section 3.3.5 on the encodings $\{e_j\}_j$ (or, more precisely, on the zero-tested vector $\omega = p_{zt} \cdot (e_j)_{1 \leq j \leq \ell} \bmod x_0$), we compute in polynomial time a vector $u \in \mathbb{Z}^\ell$ such that $\langle u, \hat{m}_i \cdot x_i \rangle \equiv 0 \pmod{g_i}$ for all $1 \leq i \leq \theta$, where $\{\hat{m}_i\}_i$ are the vectors $(m_{1i}, \ldots, m_{\ell i})$ for $1 \leq i \leq \theta$. Provided that $x_i \not\equiv 0 \pmod{g_i}$, this implies $\langle u, \hat{m}_i \rangle \equiv 0 \pmod{g_i}$ for all $1 \leq i \leq \theta$. Therefore, for all $1 \leq i \leq n$, we can write $\sum_{j=1}^{\ell} u_j m_{ji} = k_i g_i$ for integers $\{k_i\}$ with $k_i = 0$ for $\theta + 1 \leq i \leq n$. This implies:

$$y' = \sum_{j=1}^{\ell} u_j y_j \equiv g_i \left(\sum_{j=1}^{\ell} u_j r_{ji} + k_i \right) \cdot z^{-1} \pmod{p_i}, \quad 1 \leq i \leq n.$$

Thus we see that y' is a level-one encoding of zero. Moreover, since the vector u output by the orthogonal lattice attack is short, y' has small noise.

It is also worth noticing that this attack only requires a single vector u; therefore, the first step of the attack is proven by Proposition 3.3.4.

Second step: Cheon et al. attack. The second step of the attack consists in applying the algorithm of Cheon et al. described in Section 3.2.2.2 to the three sets

$$\mathcal{A}' = \{y' \cdot \alpha_j : 1 \leq j \leq n\}, \ \mathcal{B} = \{\beta_1, \beta_2\}, \ \mathcal{C} = \{\gamma_k : 1 \leq k \leq n\}$$

where $\mathcal{A}'$ is the set constructed from $\mathcal{A}$ by multiplying every encoding by the encoding y' generated in the first step. Since y' is an encoding of zero, all encodings in $\mathcal{A}'$ are encodings of zero at level 2. We can thus apply the Cheon et al. attack on these sets in order to reveal all prime factors $\{p_i\}_i$ of x_0. We refer to Section 3.2.2.2 for the details of the algorithm to factor x_0. Our complete algorithm can therefore be described as follows.

Algorithm 6 Cheon et al. attack against CLT13 with encodings of partial zero messages

Parameters: Integer $\ell \geq 1$, $0 < \theta \leq n$ and the CLT13 parameters with $\kappa = 4$

Input: Sets $\mathcal{A}, \mathcal{B}, \mathcal{C}$ of encodings at level one of non-zero messages, with $\#\mathcal{A} = \#\mathcal{C} = n$ and $\#B = 2$; a set $\mathcal{Y}$ of level-one encodings of messages with non-zero support of length θ with $\#\mathcal{Y} = \ell$

Output: The prime factors $\{p_i : 1 \leq i \leq n\}$ of x_0

1: Form the vector $y = (y_j)_j$ from $\mathcal{Y}$ and compute the product encoding $e = (e_j)_j$ with $e_j = y_j \alpha_1 \beta_1 \gamma_1$ for $1 \leq j \leq \ell$

2: Run the extended lattice attack (Algorithm 3) on the zero-tested vector $\omega = p_{zt} \cdot e \bmod x_0$. Let $u \in \mathbb{Z}^\ell$ be the shortest vector in the basis matrix U of the lattice of vectors orthogonal to $\{\hat{m}_i\}_i$ modulo $\{g_i\}_i$

3: Set $y' = \langle y, u \rangle \in \mathbb{Z}$ and construct the set $\mathcal{A}' = \{y'\alpha : \alpha \in \mathcal{A}\}$ of level-2 encodings of zero.

4: Run the Cheon et al. attack (Algorithm 2) on the sets $\mathcal{A}', \mathcal{B}, \mathcal{C}$.

Remark 3.3.6. This remark concerns a variant attack with $\kappa = 3$. Since the orthogonal lattice attack more generally provides a set of ℓ vectors $\{u_j\} \subseteq \mathbb{Z}^\ell$ (instead of a single u; and all satisfying $\langle u_j, \hat{m}_i \rangle \equiv 0 \pmod{g_i}$ for all i), a variant of the above attack with $\kappa = 3$ consists in starting from a set $\mathcal{A} = \{\alpha_j : 1 \leq j \leq n\}$ of $\ell = n$ encodings where only the first θ components of the underlying plaintexts are non-zero, and then, generating a set $\mathcal{A}' = \{\langle u_j, \alpha \rangle : 1 \leq j \leq n\}$ of encodings of zero, with the vector of encodings $\alpha = (\alpha_1, \ldots, \alpha_n)$. One can then apply the attack by Cheon et al. as previously on the three sets $\mathcal{A}'$, $\mathcal{B}$ and $\mathcal{C}$.

Note that in our description of the first step of the attack, that is, the generation of encodings of zero, we have relied on the orthogonal lattice attack from Section 3.3.5 with the bound $\alpha\theta < \nu$. However, the attack from Section 3.3.6 is easily adapted to reach the improved bound $\alpha\theta = O(\nu^2)$. In this case, following the notation of Section 3.3.6, the attacker can obtain $\ell \cdot d$ level-two encodings of zero given by $\{\langle u_j, c_k \rangle : 1 \leq j \leq \ell, 1 \leq k \leq d\}$, where c_k is the vector of encodings $(c_j \cdot c'_k)_{1 \leq j \leq \ell}$ with the encodings $\{c_j \cdot c'_k\}$ considered in Section 3.3.6.

3.3.10 Application to CLT13-based constructions with independent slots

In this section, we show that our orthogonal lattice attack from Section 3.3.4 can be applied to various constructions over CLT13 multilinear maps with independent slots.

3.3.10.1 The multilinear subgroup elimination assumption

The multilinear subgroup elimination assumption is used in [GLW14] for witness encryption and in [GLSW15] for constructing program obfuscation, based on a single assumption, independent of the particular circuit to be obfuscated. The multilinear subgroup elimination assumption is stated for a generic model of composite-order multilinear maps. Below, we show that our attacks break this assumption over CLT13 composite-order multilinear maps. We note that since the GLW14 scheme also includes encodings of zeroes, it could also be broken more directly by the Cheon et al. attack. We recall the definition from [GLSW15].

Definition 3.3.7 ((μ, ν)-multilinear subgroup elimination assumption [GLSW15])**.** *Let G be a group of order $N = a_1 \cdots a_\mu b_1 \cdots b_\nu c$ where $a_1, \ldots, a_\mu, b_1, \ldots, b_\nu, c$ are $\mu+\nu+1$ distinct primes. We*

give out generators $x_{a_1}, \ldots, x_{a_\mu}, x_{b_1}, \ldots, x_{b_\nu}$ for each prime order subgroup except for the subgroup of order c. For each $1 \leq i \leq \mu$, we also give out a group element h_i sampled uniformly at random from the subgroup of order $ca_1 \cdots a_{i-1}a_{i+1} \cdots a_\mu$. The challenge term is a group element $T \in G$ that is either sampled uniformly at random from the subgroup of order $ca_1 \cdots a_\mu$ or uniformly at random from the subgroup of order $a_1 \cdots a_\mu$. The task is to distinguish between these two distributions of T.

For simplicity, we consider the assumption with $\mu = 1$ and $\nu = 0$; the generalization of our attack to any (μ, ν) is straightforward. With this assumption, G is a group of order $a_1 c$. The challenge $T \in G$ is either generated at random from the subgroup of order $a_1 c$, or from the subgroup of order a_1. In the context of a CLT13 instantiation, we may assume that $a_1 = \prod_{i=1}^{\theta} g_i$ and $c = \prod_{i=\theta+1}^{n} g_i$. Note that in that case, a_1 and c are not prime numbers, but the assumption can still be considered for composite integers $\{a_i\}_i$, $\{b_i\}_i$ and c. The encoding T is then either generated from a random plaintext $m \in \bigoplus_{i=1}^{n} \mathbb{Z}/g_i\mathbb{Z}$, or from a random plaintext with only the θ first components non-zero, that is $m \equiv 0 \pmod{g_i}$ for $\theta + 1 \leq i \leq n$. It is easy to see that our attacks from Section 3.3.5 and Section 3.3.6 apply in this setting. Namely, when only the first θ components of the plaintext m corresponding to the challenge T are non-zero, our attacks recover the product $a_1 = \prod_{i=1}^{\theta} g_i$, whereas the attacks will fail when m is a random plaintext. Therefore the challenge T is easily distinguished unless θ is large enough; more precisely, θ must heuristically satisfy the bound given by Equation (3.26) in order to prevent the attack.

3.3.10.2 The Zimmerman circuit obfuscation scheme

At Eurocrypt 2015, Zimmerman described a technique to obfuscate programs without matrix branching programs, based on composite-order multilinear maps, see [Zim15]. A plaintext m belongs to $\mathbb{Z}/N\mathbb{Z}$ for a composite modulus $N = N_{\mathsf{ev}} \cdot N_{\mathsf{chk}}$, and the ring $\mathbb{Z}/N\mathbb{Z}$ is viewed as a direct product of an "evaluation" ring $\mathbb{Z}/N_{\mathsf{ev}}\mathbb{Z}$ to evaluate the circuit, and a "checksum" ring $\mathbb{Z}/N_{\mathsf{chk}}\mathbb{Z}$ to prevent the adversary from evaluating a different circuit; those two evaluations are performed in parallel. Using the CLT13 notations from Section 3.2.1, one can let $N_{\mathsf{ev}} = \prod_{i=1}^{\theta} g_i$ and $N_{\mathsf{chk}} = \prod_{i=\theta+1}^{n} g_i$. In that case, the parameter θ must satisfy the bound given by (3.26) to prevent our lattice attack.

3.3.10.3 The FRS17 construction for preventing input partitioning attacks

At Asiacrypt 2017, Fernando, Rasmussen and Sahai described constructions of "stamping functions" for preventing input-partitioning attacks on matrix branching programs [FRS17]. Their third construction is based on permutation hash functions and is instantiated over CLT13 multilinear maps with independent slots. More precisely, the permutation hash function is written as a matrix branching program, and multiple such permutation hash functions $\{h_i\}_i$ are evaluated in parallel along with the main matrix branching program; this is to ensure that only inputs of the form $x \| h(x)$ can be evaluated, where $h(x) = h_1(x) \| \cdots \| h_t(x)$, which prevents input partitioning attacks.

Matrix branching programs. Let us first recall the construction of [GGH$^+$13b] to obfuscate matrix branching programs. A matrix branching program BP of length n_p on ℓ-bit inputs

$x \in \{0,1\}^\ell$ is evaluated by computing:

$$C(x) = b_0 \cdot \prod_{i=1}^{n_p} B_{i,x_{\mathsf{inp}(i)}} \cdot b_{n_p+1} \,, \tag{3.28}$$

where $\{B_{i,b} : 1 \le i \le n_p, b \in \{0,1\}\}$ are $2n_p$ square matrices and b_0 and b_{n_p+1} are bookend vectors; then $\mathsf{BP}(x) = 0$ if $C(x) = 0$, and $\mathsf{BP}(x) = 1$ otherwise. The role of the integer $\mathsf{inp}(i) \in \{1, \ldots, \ell\}$ is to indicate which bit of x is read at step i of the product matrix computation. The matrices $B_{i,b}$ are first randomized by choosing $n_p + 1$ random invertible matrices $\{R_i : 0 \le i \le n_p\}$ and then letting $\tilde{B}_{i,b} = R_{i-1} B_{i,b} R_i^{-1}$ for $1 \le i \le n_p$, with also $\tilde{b}_0 = b_0 R_0^{-1}$ and $\tilde{b}_{n_p+1} = R_{n_p} b_{n_p+1}$. We obtain a randomized matrix branching program with the same result since the randomization matrices $\{R_i\}_i$ cancel each other:

$$C(x) = \tilde{b}_0 \cdot \prod_{i=1}^{n_p} \tilde{B}_{i,x_{\mathsf{inp}(i)}} \cdot \tilde{b}_{n_p+1}. \tag{3.29}$$

The entries of $\{\tilde{B}_{i,b}\}_{i,b}$ are then independently encoded, as well as the bookend vectors $\tilde{b}_0$ and $\tilde{b}_{n_p}$. We obtain the matrices and vectors $\hat{B}_{i,b} = \mathsf{Encode}_{\{i+1\}}(\tilde{B}_{i,b})$, $\hat{b}_0 = \mathsf{Encode}_{\{1\}}(\tilde{b}_0)$ and $\hat{b}_{n_p+1} = \mathsf{Encode}_{\{n_p+2\}}(\tilde{b}_{n_p+1})$. Here we have denoted by $\mathsf{Encode}_{\{i\}}(\cdot)$ an encoding relative to the singleton $\{i\}$. The matrix branching program from Equation (3.28) can then be evaluated over the encoded matrices:

$$\hat{C}(x) = \hat{b}_0 \cdot \prod_{i=1}^{n_p} \hat{B}_{i,x_{\mathsf{inp}(i)}} \cdot \hat{b}_{n_p+1} \tag{3.30}$$

The resulting $\hat{C}(x)$ is then a last-level encoding that can be zero-tested to check if $C(x) = 0$, which reveals the output of the branching program $\mathsf{BP}(x)$, without revealing the matrices $\{B_{i,b}\}_{i,b}$.

Application to the FRS17 construction. The [FRS17] scheme constructs a modified matrix branching program BP' that receives as input $u \| v_1 \ldots v_t$ and checks whether $v_i = h_i(u)$ for all $1 \le i \le t$, where $\{h_i\}_i$ are permutation hash functions; in that case, BP' returns $\mathsf{BP}(u)$ where BP is the original branching program; otherwise, it returns some non-zero value. As explained in [FRS17], multiple branching programs can be evaluated in parallel with composite order multilinear maps. With the countermeasure from [GLW14] over CLT13 (see Section 3.3.3.1), each branching program is then evaluated modulo a product of θ of the primes $\{g_i\}_i$, instead of a single g_i in [FRS17].

It is easy to generate an input $u \| v_1 \ldots v_t$ such that $\mathsf{BP}(u) = 0$ and $v_i = h_i(u)$ for all $1 \le i \le t$ except for some $i = i^\star$; in that case, only one of the $t+1$ parallel matrix branching program will evaluate to a non-zero value. Our lattice attack from Section 3.3.5 can therefore recover the secret plaintext ring $\bigoplus_{i=1}^{n} \mathbb{Z}/g_i\mathbb{Z}$ of CLT13, under the condition $\alpha\theta < \nu$. Alternatively, assuming that multiple zero-test parameters are available, our extended attack applies under the condition $\alpha\theta = O(\nu^2)$, as described at the end of Section 3.3.6.

We shall note however that in both cases, our attack against [FRS17] only limits itself to the recovery of the secret plaintext ring of CLT13, and not all secret parameters of CLT13. We leave the latter as an open problem.

3.3.11 Appendix: Proof of Proposition 3.3.4

We provide a proof of Proposition 3.3.4, based on Lemma 3.3.2.

Let $a = (\alpha_1, \ldots, \alpha_\theta, 1) \in \mathbb{Z}^{\theta+1}$. We let $C = 2^{\rho_R - \alpha + 1}$ and consider the lattice $\mathcal{A}^\perp$ of vectors $(Cx, y) \in \mathbb{Z}^\theta \times \mathbb{Z}$ such that (x, y) is orthogonal to a modulo x_0. Further, we let $B = \theta 2^{\rho_R+2}$ and let $\mathcal{L} \subseteq \mathbb{Z}^{\ell+1}$ denote the lattice of vectors $(Bu, v) \in \mathbb{Z}^\ell \times \mathbb{Z}$ such that the vector (u, v) is orthogonal to $(\omega, 1)$ modulo x_0. Let Λ be the lattice of vectors $u \in \mathbb{Z}^\ell$ such that $\langle u, \hat{m}_i \rangle \equiv 0 \pmod{g_i}$ for all $1 \leq i \leq \theta$. We denote by u_0 a shortest non-zero vector of Λ. We write $\langle u_0, \hat{m}_i \rangle = k_i g_i$ with $k_i \in \mathbb{Z}$. To u_0 we thus associate the vector $\Phi(u_0) = (Bu_0, -\sum_{i=1}^\theta k_i s_i - \langle u_0, R \rangle)$. From the definition of ω and the congruence relations $g_i \alpha_i \equiv s_i \pmod{x_0}$, we have that $(u_0, -\sum_{i=1}^\theta k_i s_i - \langle u_0, R \rangle)$ is orthogonal to $(\omega, 1)$ modulo x_0, and consequently, $\Phi(u_0) \in \mathcal{L}$.

Let $g = \prod_{i=1}^\theta g_i$. We shall now show that $\Phi(u_0)$ has square norm upper satisfying

$$\|\Phi(u_0)\|^2 \quad \leq \quad (\ell+1)B^2\|u_0\|^2 \leq \ell(\ell+1)B^2 g^{2/\ell}\,. \tag{3.31}$$

Indeed, we write $\|\Phi(u_0)\|^2 \leq B^2\|u_0\|^2 + (\sum_{i=1}^\theta |k_i s_i| + \|u_0\|\|R\|)^2$. From $\|\hat{m}_i\| \leq \sqrt{\ell} 2^\alpha$, we obtain $2^{\alpha-1}|k_i| \leq |k_i| g_i \leq \|u_0\|\|\hat{m}_i\| \leq \sqrt{\ell} 2^\alpha \|u_0\|$; *i.e.* $|k_i| \leq 2\sqrt{\ell}\|u_0\|$ for all i. Combined with $\|R\| \leq \sqrt{\ell}\|R\|_\infty \leq \sqrt{\ell} 2^{\rho_R}$, this gives

$$\sum_{i=1}^\theta |k_i s_i| + \|u_0\|\|R\| \leq \sqrt{\ell}\|u_0\| \cdot (2^{\rho_R+1}\theta + 2^{\rho_R}) \leq \sqrt{\ell}\|u_0\|(2 \cdot 2^{\rho_R+1}\theta) = \sqrt{\ell} B \|u_0\|$$

Therefore, $\|\Phi(u_0)\|^2 \leq B^2\|u_0\|^2 + \ell B^2\|u_0\|^2 = (\ell+1)B^2\|u_0\|^2$. Now, since u_0 has length $\lambda_1(\Lambda)$, it follows from Minkowski's Theorem that $\|u_0\| \leq \sqrt{\ell} g^{1/\ell}$ where $g = \det(\Lambda)$, and the estimate in (3.31) easily follows.

Let $x_1 = (Bu_1, v_1)$ be the first vector in a $(3/4)$-reduced basis of the lattice $\mathcal{L}$, obtained from LLL. By Equation (2.2), it satisfies $\|x_1\| \leq 2^{\ell/2}\|\Phi(u_0)\|$, that is, combined with (3.31), $\|x_1\| \leq 2^{\ell/2}\sqrt{\ell(\ell+1)} B g^{1/\ell}$. In particular, the following bounds hold:

$$\|u_1\| \leq 2^{\ell/2}\sqrt{\ell(\ell+1)} \cdot g^{1/\ell} \tag{3.32}$$

$$|v_1| \leq 2^{\ell/2} B \sqrt{\ell(\ell+1)} \cdot g^{1/\ell}\,. \tag{3.33}$$

For simplifying the notation, we put $K := 2^{\ell/2}\sqrt{\ell(\ell+1)} g^{1/\ell}$. Now, to the vector $x_1 \in L$, we associate, for C as above, the vector $\varphi(x_1) = (C\langle u_1, \hat{m}_1 \rangle, \ldots, C\langle u_1, \hat{m}_\theta \rangle, \langle u_1, R \rangle + v_1) \in A^\perp$. As $(Bu_1, v_1) \in \mathcal{L}$, it is a direct check that $\varphi(x_1) \in \mathcal{A}^\perp$, with square norm at most

$$\|\varphi(x_1)\|^2 \leq C^2 \sum_{i=1}^\theta \|u_1\|^2 \|\hat{m}_i\|^2 + (\|u_1\|\|R\| + v_1)^2.$$

Using once more that $\|\hat{m}_i\| \leq 2^\alpha\sqrt{\ell}$ and $\|R\| \leq 2^{\rho_R}\sqrt{\ell}$, and combining with Equation (3.32) and Equation (3.33), we then obtain

$$\begin{aligned} \|\varphi(x_1)\|^2 \quad &\leq \quad C^2K^2 \cdot \theta\ell 2^{2\alpha} + (K\sqrt{\ell}2^{\rho_R} + KB)^2 \leq C^2K^2 \cdot \theta\ell 2^{2\alpha} + (2K\sqrt{\ell}B)^2 \\ &= \quad K^2\ell(C^2\theta 2^{2\alpha} + 4B^2) \end{aligned}$$

so that, using $C^2\theta 2^{2\alpha} \leq B^2 = 16\theta^2 2^{2\rho_R}$, this gives

$$\|\varphi(x_1)\| \leq 4\sqrt{5} \cdot \sqrt{\ell} \cdot \theta \cdot K \cdot 2^{\rho_R} \,. \tag{3.34}$$

Consider the vectors $\{q_i : 1 \leq i \leq \theta\} \subseteq \mathbb{Z}^{\theta+1}$ defined by $q_i = (0, \ldots 0, Cg_i, 0, \ldots, 0, -s_i)$ with Cg_i at the ith position. These vectors are clearly linearly independent; moreover, from the congruence relations $g_i\alpha_i \equiv s_i \pmod{x_0}$ for $1 \leq i \leq \theta$ we deduce that for all i, $\langle q_i, a\rangle \equiv 0 \pmod{x_0}$; i.e. $q_i \in \mathcal{A}^\perp$. Further, as $|s_i| \leq 2^{\rho_R}$, their norm is upper bounded by $\|q_i\|^2 \leq C^2 g_i^2 + 2^{2\rho_R} \leq C^2 g_i^2 + Cg_i^2 \leq 2C^2 g_i^2$ because $Cg_i \geq 2^{\rho_R - \alpha + 1} \cdot 2^{\alpha - 1} = 2^{\rho_R}$. Consequently,

$$\prod_{i=1}^{\theta} \|q_i\| \leq 2^{\theta/2} C^\theta \prod_{i=1}^{\theta} g_i = 2^{\theta/2} C^\theta g \,. \tag{3.35}$$

Equation (3.20) together with $g \leq 2^{\alpha\theta}$ imply $(1+1/\ell)\log_2(g) + (\ell+\theta)/2 + \log_2(4\sqrt{5}\sqrt{\ell+1}\theta\ell) < \log_2(x_0) - \rho_R$ and, by raising to the power of 2, we obtain $g^{1+1/\ell} \cdot 2^{\ell/2} \cdot 2^{\theta/2} \cdot 4\sqrt{5}\sqrt{\ell+1}\theta\ell < x_0/2^{\rho_R}$. This is equivalent to

$$g^{1/\ell} \cdot 2^{\ell/2} \cdot 2^{\rho_R} \cdot 4\sqrt{5}\sqrt{\ell+1} \cdot \theta\ell < \frac{C^\theta x_0}{C^\theta 2^{\theta/2} g} \,. \tag{3.36}$$

The left-hand side is lower bounded by $\|\varphi(x_1)\|$ by Equation (3.34), and the right-hand side is upper bounded by $\det(\mathcal{A}^\perp)/\prod_{i=1}^{\theta}\|q_i\|$, by Equation (3.35) together with $\det(\mathcal{A}^\perp) = C^\theta x_0$. Therefore, Equation (3.36) implies $\|\varphi(x_1)\| < \mathrm{Vol}(\mathcal{A}^\perp)/\prod_{i=1}^{\theta}\|q_i\|$. It now follows from Lemma 3.3.2 that $\varphi(x_1)$ is in the linear span generated by the vectors $\{q_i : 1 \leq i \leq \theta\}$. Since $\{g_i\}_i$ are distinct prime numbers and $\gcd(s_i, g_i) = 1$ for $1 \leq i \leq \theta$, we conclude that $\varphi(x_1)$ is in the sublattice generated by the vectors $\{q_i : 1 \leq i \leq \theta\}$. Consequently, for all $1 \leq i \leq \theta$, one has $\langle u_1, \hat{m}_i\rangle \equiv 0 \pmod{g_i}$.

The rows $\{b_j : 1 \leq j \leq \ell+1\}$ of the matrix

$$\begin{bmatrix} B \cdot 1_\ell & -\omega^T \\ 0 & x_0 \end{bmatrix},$$

form a $\mathbb{Z}$-basis of $\mathcal{L}$. Hence, by running LLL on this matrix with $\delta = 3/4$, we compute a vector x_1 of which the first ℓ entries, divided by B, produce a vector $u = u_1$ satisfying $\langle u_1, \hat{m}_i\rangle \equiv 0 \pmod{g_i}$ for all i. The algorithm terminates in polynomial time because LLL does. □

CHAPTER 4

The Hidden Lattice Problem

In this chapter, we consider the problem of revealing a small hidden lattice from the knowledge of a low-rank sublattice modulo a given sufficiently large integer – the *Hidden Lattice Problem*.

A central motivation of study for this problem is the Hidden Subset Sum Problem, whose hardness is essentially determined by that of the hidden lattice problem. We describe and compare two algorithms for the hidden lattice problem: we first adapt the algorithm by Nguyen and Stern for the hidden subset sum problem, based on orthogonal lattices, and propose a new variant, which we explain to be related by duality in lattice theory. Following heuristic, rigorous and practical analyses, we find that our new algorithm brings some advantages as well as a competitive alternative for algorithms for problems with cryptographic interest, such as Approximate Common Divisor Problems, and the Hidden Subset Sum Problem. Finally, we study variations of the problem and highlight its relevance to cryptanalysis.

The content of this chapter is based on joint work with Gabor Wiese, and which has been submitted.

4.1 Introduction

The Hidden Subset Sum Problem asks to reveal a set of binary vectors from a given linear combination modulo a sufficiently large integer. At Crypto 1999, Nguyen and Stern have proposed an algorithm for this problem, based on lattices, [NS99]. Their solution crucially relies on revealing, in the first place, the "small" lattice generated by the binary vectors: this is the underlying *Hidden Lattice Problem* (HLP). The starting point of this work is to investigate the HLP independently. For this article, we define "small" lattices as follows.

Definition 4.1.1. *Let $0 < n \leq m$ be integers and let $\Lambda \subseteq \mathbb{Z}^m$ be a lattice of rank n equipped with the standard Euclidean norm. The* size *$\sigma(\mathfrak{B})$ of a basis $\mathfrak{B}$ of Λ is defined to be*

$$\sigma(\mathfrak{B}) = \sqrt{\frac{1}{n} \sum_{v \in \mathfrak{B}} \|v\|^2}.$$

*For a real number $\mu \geq 1$, we say that Λ is μ-*small *if Λ possesses a basis $\mathfrak{B}$ of size $\sigma(\mathfrak{B}) \leq \mu$.*

Small lattices naturally occur in computational problems in number theory and cryptography. For Λ as in Definition 4.1.1, we let $\Lambda_{\mathbb{Q}}$ (resp. $\Lambda_{\mathbb{R}}$) be the $\mathbb{Q}$-span (resp. $\mathbb{R}$-span) of Λ in $\mathbb{R}^m$. The completion $\overline{\Lambda}$ of Λ is $\Lambda_{\mathbb{Q}} \cap \mathbb{Z}^m = \Lambda_{\mathbb{R}} \cap \mathbb{Z}^m$ and we say that Λ is complete if $\overline{\Lambda} = \Lambda$. As is customary in many computational problems we also work modulo $N \in \mathbb{Z}$ and write $v \in \Lambda \pmod N$ if there exists $w \in \Lambda$ such that $v - w \in (N\mathbb{Z})^m$. If $\Lambda' \subseteq \mathbb{Z}^m$, then $\Lambda' \subseteq \Lambda \pmod N$ shall mean $v \in \Lambda \pmod N$ for all $v \in \Lambda'$. We then define the Hidden Lattice Problem as follows.

Definition 4.1.2 (Hidden Lattice Problem (HLP))**.** *Let $\mu \in \mathbb{R}_{\geq 1}$, integers $1 \leq r \leq n \leq m$ and $N \in \mathbb{Z}$. Let $\mathcal{L} \subseteq \mathbb{Z}^m$ be a μ-small lattice of rank n. Further, let $\mathcal{M} \subseteq \mathbb{Z}^m$ be a lattice of rank r such that $\mathcal{M} \subseteq \mathcal{L} \pmod N$.*

The Hidden Lattice Problem (HLP) *is the task to compute from the knowledge of n, N and a basis of $\mathcal{M}$, a basis of the completion of any μ-small lattice Λ of rank n such that $\mathcal{M} \subseteq \Lambda \pmod N$.*

Note that $\mathcal{M}$ is defined modulo N and we may thus view $\mathcal{M} \subseteq (\mathbb{Z}/N\mathbb{Z})^m$. We analyse for which values of $\mu \in \mathbb{R}_{\geq 1}$ a generic HLP can be expected to be solvable. Random choices of $\mathcal{M}$ are likely to uniquely define the lattice $\mathcal{L}$, thus $\Lambda = \mathcal{L}$. We will see that $\overline{\mathcal{L}}$ is very often equal to $\mathcal{L}$: it is the *hidden lattice* to be uncovered (note that the completion makes the lattice only smaller). Our definition is more general than the framework in [NS99] and deviates in two ways: first, we do not require $\mathcal{L}$ to possess a basis of *binary* vectors as in [NS99], but instead control the size of $\mathcal{L}$ by μ. Also, instead of assuming a unique vector to be public ($r = 1$), we assume a (basis of a) sublattice $\mathcal{M}$ of arbitrary rank r to be public.

4.2 Our contributions

Our principal aims are to describe algorithms for the HLP, analyse them theoretically, heuristically and practically, and give applications.

Algorithms for the HLP. We describe two algorithms for the HLP. First, we adapt the orthogonal lattice algorithm of Nguyen and Stern [NS99], based on the (public) lattice $\mathcal{M}^{\perp_N}$ of vectors orthogonal to $\mathcal{M}$ modulo N. It naturally contains the relatively small lattice $\mathcal{L}^{\perp}$, which we can identify by lattice reduction, provided that the parameters satisfy certain conditions. Our major contribution is to propose a new two-step alternative algorithm, based on the (public) lattice $\mathcal{M}_N$ of vectors that lie in $\mathcal{M}$ modulo N. In this case, we first explain how to recognize vectors lying directly in a relatively small sublattice of the completion of the hidden lattice $\mathcal{L}$ and compute them by lattice reduction. We explain that the second step of our new algorithm can be designed to only perform linear algebra over finite fields, which is generally very fast. Therefore, our second algorithm is often faster than the orthogonal lattice algorithm. As we can directly compute short vectors in $\mathcal{L}$ instead of $\mathcal{L}^{\perp}$ (which avoids the computation of orthogonal complements), it is also conceptually easier than the orthogonal lattice algorithm. We finally justify that both algorithms are related by duality. Using celebrated transference results for the successive minima of dual lattices, we explain how to bridge both algorithms theoretically. Throughout this paper, we refer to both algorithms as Algorithm I and Algorithm II, respectively. In cryptanalysis, the orthogonal lattice has been used extensively, since its introduction in [NS97]. Lattice duality has been used for example in the context of the LWE Problem, see e.g. [Alb17]. We refer to Section 4.4 for a full description of our algorithms.

Analysis of our algorithms. We provide a heuristic analysis of our algorithms based on the Gaussian Heuristic for "random lattices". For Algorithm I, we follow the intuition of [NS99]: short enough vectors $u \in \mathcal{M}^{\perp_N}$ (which we compute by lattice reduction) must lie in $\mathcal{L}^{\perp}$. Since $\mathcal{L}^{\perp}$ has rank $m-n$, we expect to find $m-n$ such vectors. For Algorithm II, we derive an explicit lower bound on the norm of the vectors lying in $\mathcal{M}_N$ but outside $\mathcal{L}_{\mathbb{Q}}$, which gives us a criterion for establishing an explicit parameter selection. In both cases, it turns out that the HLP is solvable when the N is sufficiently large with respect to μ. Quantifying this dependence theoretically and practically is a natural question. For example, both algorithms detect hidden lattices of size $\mu = O(N^{\frac{r(m-n)}{nm}})$ up to some terms which differ according to the algorithm; in the balanced case $m = 2n = 4r$, this gives $\mu = O(N^{1/4})$. To quantify the difference between N and μ in a compact formula, we propose a definition for an arithmetic invariant attached to the HLP, which we justify to behave like an inverse-density, a handy and well-studied invariant for knapsack-type problems (see e.g. [LO85, NS99]).

We next establish proven results, not conditioned on the Gaussian Heuristic. Such formal statements are not included in [NS99]. For our proofs, we rely on a discrete counting technique. For a fixed μ-small basis $\mathfrak{B}$ of $\mathcal{L}$ (sampled from some set of collections of vectors) and a given integer N, we denote by $\mathcal{H}(\mathfrak{B})$ a finite sample set of vectors constructed from $\mathfrak{B}$ and N. To an element of $\mathcal{H}(\mathfrak{B})$, we naturally associate a HLP with hidden lattice $\mathcal{L}$. On each of these problems, we "run" either Algorithm I or Algorithm II, and "count" how often our algorithm successfully computes a basis of $\overline{\mathcal{L}}$ by using LLL, [LLL82]. Our proven bounds asymptotically confirm our heuristic prediction.

We have implemented our algorithms in SageMath [S$^+$20]. Our practical results confirm our theoretical findings quite accurately. Moreover, we see that both algorithms practically perform equivalently well, which is heuristically understandable from the duality between them. In some cases, Algorithm II outperforms Algorithm I: the second step of Algorithm II is computationally simpler than for Algorithm I, leading to strongly improved running times.

At informal level, we can state the following simplified lower bounds (see Table 4.1) for $\log(N)$ in our heuristic and proven analyses. In the proven case (for $r = 1$), the lower bound stands for $\log(N\varepsilon)$, where $\varepsilon \in (0,1)$ is fixed such that the success rate of the algorithms is $1-\varepsilon$. Here ι denotes the root Hermite factor depending on the chosen lattice reduction algorithm (which is LLL in our proven analysis).

We refer to Section 4.5 and 4.6 for our theoretical analyses and to Section 4.9 for our practical experiments.

Variations and applications. Some variations of Def. 4.1.2 are of interest to us. First, we study the case where given vectors lie in a small lattice modulo N only up to unknown short "noise" vectors; we call this the *noisy hidden lattice problem* (NHLP). We notice that we can cancel the effect of the noise, by reducing the NHLP to a HLP with a "larger" (in the sense of size and dimension) hidden lattice, and apply our previous algorithms without changes. We also consider a *decisional* version (DHLP) of the hidden lattice problem, asking about the existence of a μ-small lattice $\mathcal{L}$ containing $\mathcal{M}$ modulo N. This problem, although not asking for the computation of $\overline{\mathcal{L}}$ lies at the heart of many cryptanalytic settings, and may thus be of interest to cryptanalysts. We recognize that the existence of such $\mathcal{L}$ strongly impacts the

	Algorithm	Lower bound for $\log(N)$
Heuristic	I	$\frac{mn}{r}\log(\iota) + \frac{mn}{r(m-n)}\log(\mu) + \frac{n}{2r}\log(m-n)$
	II	$\frac{m}{m-n}\frac{mn}{r}\log(\iota) + \frac{mn}{r(m-n)}\log(\mu) + \frac{mn}{2r(m-n)}\log\left(\frac{n}{m}\right)$
Proven	I	$mn\log(\iota) + n(n+1)\log(\mu) + \frac{n(m-n)}{2}\log(\frac{2(m-n)}{3}) + n\log(3\sqrt{n}) + 1$
	II	$mn\log(\iota) + n(n+2)\log(\mu) + n\log(3n^2) + 1$

Table 4.1: Lower bounds for $\log(N)$ as functions of n, m, r, μ

geometry of $\mathcal{M}^{\perp_N}$ (or $\mathcal{M}_N$) and, consequently, our algorithms solve the decisional version heuristically.

Finally, we describe applications of the HLP together with some improvements implied by our Algorithm II. Our applications show that the HLP appears somewhat naturally in many different frameworks. We mostly refer to the works [CP19, CG20, CN19a, CNT10, BNNT11]. We refer to Section 4.7 and 4.8.

4.3 Background and notation on lattices

4.3.1 Lattices

Throughout this section we fix a lattice $\Lambda \subseteq \mathbb{Z}^m$ of positive rank n. For more background, we refer to Chapter 2.

Definition 4.3.1. *For $N \in \mathbb{Z}$ we define the N-*orthogonal lattice of Λ *by*

$$\Lambda^{\perp_N} = \{v \in \mathbb{Z}^m \mid \forall\, w \in \Lambda \pmod N : \langle v, w\rangle \equiv 0 \pmod N\}$$

*and the N-*congruence lattice of Λ *by*

$$\Lambda_N = \{v \in \mathbb{Z}^m \mid v \in \Lambda \pmod N\}\,.$$

These only depend on Λ modulo N.

Note that $\Lambda^{\perp_0} = \Lambda^\perp$ is the usual orthogonal lattice, and $\Lambda_0 = \Lambda$. Assume now $N \neq 0$. The map $\Lambda^{\perp_N} \to \mathrm{Hom}_{\mathbb{Z}}(\Lambda_N, N\mathbb{Z}) \simeq N \cdot \mathrm{Hom}_{\mathbb{Z}}(\Lambda_N, \mathbb{Z}) \simeq N \cdot (\Lambda_N)^\vee$ sending w to $(v \mapsto \langle v, w\rangle)$ is an isomorphism. Moreover, for a basis matrix $A \in \mathbb{Z}^{m\times n}$ of Λ, $\Lambda^{\perp_N}$ is precisely the kernel of $\mathbb{Z}^m \to (\mathbb{Z}/N\mathbb{Z})^n, v \mapsto A^T v$, and thus has volume dividing N^n. Since the product of the volumes of dual lattices is 1, we conclude that N^{m-n} divides $\mathrm{Vol}(\Lambda_N)$. If the Gram matrix $A^T A$ of A is invertible over $\mathbb{Z}/N\mathbb{Z}$, then these inequalities are equalities, i.e. $\mathrm{Vol}(\Lambda^{\perp_N}) = N^n$ and $\mathrm{Vol}(\Lambda_N) = N^{m-n}$.

Lemma 4.3.2. *Let Λ be a lattice of positive rank n.*

(i) The completion of Λ *satisfies* $(\Lambda^\perp)^\perp = \overline{\Lambda}$.

(ii) If $\Lambda' \subseteq \Lambda^\perp$ *is a sublattice of the same rank as* $\Lambda^\perp$*, then* $\overline{\Lambda} = \Lambda'^\perp$.

(iii) (Hadamard's inequality) $\mathrm{Vol}(\Lambda) \leq \prod_{v\in\mathfrak{B}} \|v\| \leq \sigma(\mathfrak{B})^n$ *for any basis* $\mathfrak{B}$ *of* Λ.

(iv) $\mathrm{Vol}(\Lambda^\perp) = \mathrm{Vol}(\overline{\Lambda}) \leq \mathrm{Vol}(\Lambda)$

Proof. For *(i)* and *(iv)*, see Section 2 and Corollary 2 in [NS97]. The inequality in *(iv)* follows because $\Lambda \subseteq \overline{\Lambda}$. Statement *(ii)* follows from *(i)* because $\overline{\Lambda} = (\Lambda^\perp)^\perp \subseteq \Lambda'^\perp$ are of the same rank with $\overline{\Lambda}$ complete. The first inequality in *(iii)* is Hadamard's Inequality (Proposition 2.2.1) and the last inequality in *(iii)* follows from the arithmetic-geometric mean inequality. □

4.3.2 Lattice reduction

We rely on lattice reduction. We refer to Chapter 2 for more details. We assume that given as input a basis of a lattice Λ, a lattice reduction algorithm outputs a basis $\{b'_i\}_i$ of Λ with

$$\|b'_i\| \leq \iota^n \lambda_i(\Lambda)\ , 1 \leq i \leq n\ ,$$

as mentioned in Equation 2.4. In practice, one uses LLL [LLL82] or BKZ [HPS11a, CN12]. Throughout this chapter we mainly rely on the LLL algorithm (see Theorem 2.2.5 in Section 2.2.2) to obtain proved upper bounds on the length of short vectors. For the complexity of lattice reduction we make use of the formulae for the L^2-algorithm and the heuristic complexity for the BKZ algorithm (see Section 2.2.2).

4.4 Algorithms for the HLP

We provide and compare two algorithms for the HLP. Our first algorithm follows the orthogonal lattice algorithm proposed by Nguyen and Stern, [NS99]. We then propose a variant based on the N-congruence lattice, which has some advantages over the first algorithm. We justify that both algorithms are related via the duality (up to scalar) of the lattices in Definition 4.3.1. Finally, we present some practical observations for both algorithms.

Let us introduce some notation. For a basis $\mathfrak{B} = \{v_1, \ldots, v_n\}$ of $\mathcal{L}$, consider the coordinate isomorphism $c_\mathfrak{B} : \mathcal{L} \to \mathbb{Z}^n$, sending $\sum_{i=1}^n a_i v_i$ to $(a_1, \ldots, a_n) \in \mathbb{Z}^n$. Let π_N be the natural projection $\mathbb{Z}^n \to (\mathbb{Z}/N\mathbb{Z})^n$, and denote by $c_{\mathfrak{B},N} : \mathcal{L} \to (\mathbb{Z}/N\mathbb{Z})^n$ the composition $\pi_N c_\mathfrak{B}$. We now assume that $\mathcal{L}$ is complete, that is, $\mathcal{L} = \overline{\mathcal{L}}$. Then $N\mathcal{L} = \mathcal{L} \cap N\mathbb{Z}^m$, and thus we can extend $c_{\mathfrak{B},N}$ to $\mathcal{L}_N \to (\mathbb{Z}/N\mathbb{Z})^n$, by setting $c_{\mathfrak{B},N}(\ell + Nt) = c_{\mathfrak{B},N}(\ell)$ for every $\ell \in \mathcal{L}$ and $t \in \mathbb{Z}^m$. For $\mathcal{M} \subseteq \mathcal{L} \pmod N$, that is, $\mathcal{M} \subseteq \mathcal{L}_N$, define

$$\mathcal{M}_\mathfrak{B} = c_{\mathfrak{B},N}(\mathcal{M}) \subseteq (\mathbb{Z}/N\mathbb{Z})^n\ ,$$

the image of $\mathcal{M}$ under $c_{\mathfrak{B},N}$. For our algorithms, we consider the lattices $(\mathcal{M}_\mathfrak{B})^{\perp_N} \subseteq \mathbb{Z}^n$ and $(\mathcal{M}_\mathfrak{B})_N = \pi_N^{-1}(\mathcal{M}_\mathfrak{B}) \subseteq \mathbb{Z}^n$ of rank n, respectively.

4.4.1 The orthogonal lattice algorithm for the HLP

We adapt the algorithm from [NS99] based on the orthogonal lattice. Given an instance of the HLP with notation as in Definition 4.1.2, we have $\mathcal{L}^\perp \subseteq \mathcal{M}^\perp \subseteq \mathcal{M}^{\perp_N}$. In imprecise terms, the smallness of $\mathcal{L}$ implies the smallness of $\mathcal{L}^\perp$. We argue below that in a sufficiently generic case, the lattice $\mathcal{M}^{\perp_N}$ contains a sublattice $\mathcal{N}_\mathrm{I}$ of $\mathcal{L}^\perp$ of the same rank $m-n$. By Lemma 4.3.2, $\overline{\mathcal{L}} = \mathcal{N}_\mathrm{I}^\perp$ is a solution to the given HLP. Under the assumption $\mathcal{N}_\mathrm{I} \subseteq \mathcal{L}^\perp$, one has $\mathcal{N}_\mathrm{I}^\perp = \overline{\mathcal{L}}$, and so a solution to the HLP. The *orthogonal lattice algorithm* is as follows; we refer to it as Algorithm I.

Algorithm 7 Solve the HLP using the orthogonal lattice (Algorithm I)

Parameters: The HLP parameters n, m, r, μ, N from Definition 4.1.2

Input: A valid input for the HLP: a basis of a lattice $\mathcal{M} \subseteq \mathbb{Z}^m$ of rank r such that $\mathcal{M} \subseteq \mathcal{L} \pmod N$ where $\mathcal{L}$ is a μ-small lattice of rank n in $\mathbb{Z}^m$

Output: A basis of the lattice $\overline{\mathcal{L}}$ (under suitable parameter choice)

1: Compute a basis matrix $B(\mathcal{M}, N)$ of $\mathcal{M}^{\perp_N}$
2: Run a lattice reduction algorithm on the basis $B(\mathcal{M}, N)$ to compute a reduced basis $u_1, \ldots, u_\ell$ of $\mathcal{M}^{\perp_N}$, where $\ell = m - r$ if $N = 0$ and $\ell = m$ otherwise; order the vectors $\{u_i\}_i$ by increasing norm
3: Construct the lattice $\mathcal{N}_\mathrm{I} = \bigoplus_{i=1}^{m-n} \mathbb{Z} u_i$
4: Compute and return a basis of $\mathcal{N}_\mathrm{I}^\perp$ (see Section 4.4.4)

Identifying $\mathcal{L}^\perp$ by its smallness. The decisive point in this algorithm is that $\mathcal{N}_\mathrm{I}$ is expected to lie in $\mathcal{L}^\perp$ due to the smallness of the latter for the following heuristic argumentation. A more precise discussion follows below. Recall that $\mathcal{L}^\perp$ is a "small" sublattice of $\mathcal{M}^{\perp_N}$ of rank $m-n$. One hence expects that lattice reduction algorithms identify $m-n$ linearly independent "short" vectors in $\mathcal{M}^{\perp_N}$. Indeed, in practice one sees a significant jump in the size of the basis vectors after the first $m-n$ vectors, i.e. $\mathcal{N}_\mathrm{I}$ is the unique "small" sublattice of $\mathcal{L}^\perp$ of rank $m-n$. In Section 4.7.2 we will formulate the decisional hidden lattice problem (DHLP), asking for the existence of $\mathcal{L}$. This size jump is exactly what is detected by the algorithm for the decisional version. Let us note that heuristically a "short" vector that is orthogonal to $\mathcal{M}$ modulo N is genuinely orthogonal over $\mathbb{Z}$. Consequently, if $m - n' > m - n$ "short" vectors were found by the lattice reduction algorithm in the first step, then $\mathcal{M}$ would already lie in a small lattice of rank $n' < n$, which is heuristically not the case.

Proposition 4.4.1 makes the smallness of $\mathcal{L}^\perp$ precise by giving a lower bound for vectors lying outside $\mathcal{L}^\perp$. For a basis $\mathfrak{B}$ of $\mathcal{L}$, consider the coordinate morphism $c_\mathfrak{B} : \mathcal{L}_N \to (\mathbb{Z}/N\mathbb{Z})^n$ sending a vector to its coordinates with respect to $\mathfrak{B}$ modulo N. For $\mathcal{M} \subseteq \mathcal{L} \pmod N$, denote by $\mathcal{M}_\mathfrak{B} := c_\mathfrak{B}(\mathcal{M}) \subseteq (\mathbb{Z}/N\mathbb{Z})^n$ the image of $\mathcal{M}$ under $c_\mathfrak{B}$. For $\mathfrak{B} = \{v_1, \ldots, v_n\}$, we also define the group homomorphism

$$\Phi_\mathfrak{B} : \mathcal{M}^{\perp_N} \to \mathcal{M}_\mathfrak{B}^{\perp_N} \ , \ u \mapsto (\langle u, v_i \rangle)_{i=1,\ldots,n} \ . \tag{4.1}$$

Indeed, by the linearity of the inner product, it is easy to see that for vectors u in the N-orthogonal lattice of $\mathcal{M}$, the vector $\Phi_\mathfrak{B}(u)$ lies in the N-orthogonal lattice of $\mathcal{M}_\mathfrak{B}$. Note that the kernel of $\Phi_\mathfrak{B}$ equals $\mathcal{L}^\perp$, and is hence independent of the choice of basis $\mathfrak{B}$. Following

[NS99], the idea is as follows: for short enough vectors in $\mathcal{M}^{\perp_N}$, their image under $\Phi_{\mathfrak{B}}$ inside $\mathcal{M}_{\mathfrak{B}}^{\perp_N}$ will not become significantly longer (since the vectors $\{v_i\}_i$ generate a small basis); if these vectors have Euclidean norm less than the first minimum of $\mathcal{M}_{\mathfrak{B}}^{\perp_N}$, they must be zero in $\mathcal{M}_{\mathfrak{B}}^{\perp_N}$, and therefore $u \in \mathcal{L}^{\perp}$. More precisely, we have the following proposition.

Proposition 4.4.1. *If $u \in \mathcal{M}^{\perp_N} \setminus \mathcal{L}^{\perp}$, then $\|u\| \geq \lambda_1(\mathcal{M}_{\mathfrak{B}}^{\perp_N})/(\sqrt{n} \cdot \mu)$ for any basis $\mathfrak{B}$ of $\mathcal{L}$.*

Proof. The kernel of $\Phi_{\mathfrak{B}}$ is $\mathcal{L}^{\perp}$ and is independent of the choice of $\mathfrak{B}$. So, for $u \in \mathcal{M}^{\perp_N} \setminus \mathcal{L}^{\perp}$, we have $\lambda_1(\mathcal{M}_{\mathfrak{B}}^{\perp_N}) \leq \|\Phi_{\mathfrak{B}}(u)\| \leq \sqrt{\sum_{i=1}^{n} \|u\|^2 \|v_i\|^2} = \|u\|\sqrt{n}\sigma(\mathfrak{B})$ by the Cauchy-Schwarz inequality. We conclude using $\sigma(\mathfrak{B}) \leq \mu$, as $\mathcal{L}$ is μ-small. □

4.4.2 An alternative algorithm for the HLP

We describe a variant of Algorithm I based on the (public) lattice $\mathcal{M}_N$ for $N \neq 0$. For an instance of the HLP as in Definition 4.1.2, we have $\mathcal{M}_N \subseteq \mathcal{L}_N$. We argue below that in a sufficiently generic case, the lattice $\mathcal{M}_N$ contains a sublattice $\mathcal{N}_{\mathrm{II}}$ of $\overline{\mathcal{L}}$ of rank n. Then $\overline{\mathcal{N}_{\mathrm{II}}} = \overline{\mathcal{L}}$ solves the HLP. We summarize the algorithm as follows and refer to it as Algorithm II.

Algorithm 8 Solve the HLP using the congruence lattice (Algorithm II)

Parameters: The HLP parameters n, m, r, μ, N from Definition 4.1.2
Input: A valid input for the HLP: a basis of a lattice $\mathcal{M} \subseteq \mathbb{Z}^m$ of rank r such that $\mathcal{M} \subseteq \mathcal{L} \pmod N$ where $\mathcal{L}$ is a μ-small lattice of rank n in $\mathbb{Z}^m$
Output: A basis of the lattice $\overline{\mathcal{L}}$ (under suitable parameter choice)
1: Compute a basis matrix $B'(\mathcal{M}, N)$ of $\mathcal{M}_N$
2: Run a lattice reduction algorithm on the basis $B'(\mathcal{M}, N)$ to compute a reduced basis $u_1, \ldots, u_m$ of $\mathcal{M}_N$; order the vectors $\{u_i\}_i$ by increasing norm
3: Construct the lattice $\mathcal{N}_{\mathrm{II}} = \bigoplus_{i=1}^{n} \mathbb{Z}u_i$
4: Compute and return a basis of $\overline{\mathcal{N}_{\mathrm{II}}}$ (see Section 4.4.4)

Identifying $\overline{\mathcal{L}}$ by its smallness. The key point of the algorithm is the existence of a somewhat small sublattice $Q \subseteq \mathcal{M}_N$ of rank n. Its existence implies lattice reduction applied on $\mathcal{M}_N$ to output n short vectors. To define Q, we use the same notation as above. Let

$$Q := Q_{\mathfrak{B},N} := c_{\mathfrak{B}}^{-1}((\mathcal{M}_{\mathfrak{B}})_N) \cap \mathcal{L} = {c_{\mathfrak{B},N}}^{-1}(\mathcal{M}_{\mathfrak{B}}) \cap \mathcal{L} \,. \tag{4.2}$$

Note that $Q \simeq (\mathcal{M}_{\mathfrak{B}})_N$ via the isomorphism $c_{\mathfrak{B}}$. The following lemma describes some properties of the lattice Q.

Lemma 4.4.2. *(a) The lattice Q is equal to $\mathcal{M}_N \cap \mathcal{L}$.*
(b) The index $(\mathcal{L} : Q)$ is a multiple of N^{n-r}, and equal to N^{n-r} if $\mathrm{Vol}((\mathcal{M}_{\mathfrak{B}})_N) = N^{n-r}$

Proof. (a) To see that $Q \subseteq \mathcal{M}_N \cap \mathcal{L}$, it suffices to show that $Q \subseteq \mathcal{M}_N$. For $q \in Q$, we have by definition, $c_{\mathfrak{B},N}(q) \in \mathcal{M}_{\mathfrak{B}}$, so $c_{\mathfrak{B},N}(q) = c_{\mathfrak{B},N}(x)$ for some $x \in \mathcal{M}$. Therfore $c_{\mathfrak{B},N}(x - q) = 0$, i.e., $q \in x + N\mathbb{Z}^m \subseteq \mathcal{M}_N$.

To see that $\mathcal{M}_N \cap \mathcal{L} \subseteq Q$, let $\ell = x + Nt \in \mathcal{M}_N \cap \mathcal{L}$ with $\ell \in \mathcal{L}, x \in \mathcal{M}$ and $t \in \mathbb{Z}^m$. This gives $c_{\mathfrak{B},N}(\ell) = c_{\mathfrak{B},N}(x) \in \mathcal{M}_{\mathfrak{B}}$, so $\ell \in c_{\mathfrak{B},N}^{-1}(\mathcal{M}_{\mathfrak{B}}) \cap \mathcal{L} = Q$.
(b) The isomorphism $Q \simeq (\mathcal{M}_{\mathfrak{B}})_N$ implies $(\mathcal{L} : Q) = (\mathbb{Z}^n : (\mathcal{M}_{\mathfrak{B}})_N) = \mathrm{Vol}((\mathcal{M}_{\mathfrak{B}})_N)$, which is a multiple of N^{n-r}. If $\mathrm{Vol}((\mathcal{M}_{\mathfrak{B}})_N) = N^{n-r}$, then $(\mathcal{L} : Q) = N^{n-r}$. □

The key point is the following general lemma (Lemma 4.4.3). When applied to $\Lambda' = \mathcal{L} \subseteq \mathcal{L}_N = \Lambda$, it gives Proposition 4.4.4.

Lemma 4.4.3. *Let $\Lambda \subseteq \mathbb{Z}^m$ be a lattice of rank m. Let $\Lambda' \subseteq \Lambda$ be a sublattice of rank $1 \leq n \leq m$. For every basis $\mathfrak{B}'$ of Λ' and every $u \in \Lambda$ with $u \notin \Lambda'_{\mathbb{Q}}$, we have*

$$\|u\| \geq \frac{\mathrm{Vol}(\Lambda)}{\prod_{v \in \mathfrak{B}'} \|v\| \cdot \prod_{i=n+2}^{m} \lambda_i(\Lambda)}$$

Proof. Since $u \in \Lambda$ and $u \notin \Lambda'_{\mathbb{Q}}$, $\Lambda' \oplus \mathbb{Z}u$ is a sublattice of Λ of rank $n+1$. There are linearly independent (ordered) vectors $t_1, \ldots, t_m \in \Lambda$ with $\|t_j\| = \lambda_j(\Lambda)$ for all j. Since $\{t_j\}_j$ are linearly independent, we can choose $m-n-1$ vectors $t'_1, \ldots, t'_{m-n-1}$ among $\{t_j\}_j$ such that $\Omega = \Lambda' \oplus (\mathbb{Z}u) \oplus (\bigoplus_{j=1}^{m-n-1} \mathbb{Z}t'_j)$ is a sublattice of Λ of finite index. In particular, $\mathrm{Vol}(\Lambda) \leq \mathrm{Vol}(\Omega)$. Since $\prod_{j=1}^{m-n-1} \|t'_j\| \leq \prod_{i=n+2}^{m} \lambda_i(\Lambda)$, we obtain by Hadamard's Inequality (Proposition 2.2.1) that the volume of Ω is upper bounded by

$$\left(\prod_{v \in \mathfrak{B}'} \|v\| \right) \cdot \|u\| \cdot \prod_{j=1}^{m-n-1} \|t'_j\| \leq \left(\prod_{v \in \mathfrak{B}'} \|v\| \right) \cdot \|u\| \cdot \prod_{i=n+2}^{m} \lambda_i(\Lambda) \,.$$

□

Proposition 4.4.4. *Let $u \in \mathcal{M}_N \setminus \mathcal{L}_{\mathbb{Q}}$. Then we have*

$$\|u\| \geq \frac{1}{\mu^n} \cdot \frac{\mathrm{Vol}(\mathcal{L}_N)}{\prod_{i=n+2}^{m} \lambda_i(\mathcal{M}_N)}$$

Proof. By Lemma 4.4.3 with $\Lambda' = \mathcal{L} \subseteq \mathcal{L}_N = \Lambda$ we obtain together with Lemma 4.3.2 *(iii)* that $\|u\| \geq (1/\mu^n) \cdot (\mathrm{Vol}(\mathcal{L}_N) / \prod_{i=n+2}^{m} \lambda_i(\mathcal{L}_N))$ and the claimed bound follows using the inclusion $\mathcal{M}_N \subseteq \mathcal{L}_N$. □

As a consequence, short enough vectors in $\mathcal{M}_N$, which we seek by lattice reduction, must eventually lie in $\mathcal{L}_{\mathbb{Q}}$, and as they are integral, also in $\overline{\mathcal{L}}$.

4.4.3 Relation between the algorithms

Algorithms I and II are related by the duality relations of $\mathcal{M}^{\perp_N}$ and $\mathcal{M}_N$. We refer to Section 2.2.1 for more generalities on dual lattices. We argue that it is equivalent to obtain short vectors in both lattices $\mathcal{M}_N$ and $\mathcal{M}^{\perp_N}$, by relying on Banaszczyk's Transference Theorem (see Theorem 2.2.3).

Proposition 4.4.5. *For every lattice $\mathcal{M} \subseteq \mathbb{Z}^m$, the following hold:*

(i) $\prod_{j=1}^{m-n} \lambda_j(\mathcal{M}^{\perp_N}) \leq \gamma_m^{m/2} \frac{\mathrm{Vol}(\mathcal{M}^{\perp_N})}{N^n} \prod_{j=1}^{n} \lambda_j(\mathcal{M}_N)$

(ii) $\prod_{j=1}^{n} \lambda_j(\mathcal{M}_N) \leq \gamma_m^{m/2} \frac{\mathrm{Vol}(\mathcal{M}_N)}{N^{m-n}} \prod_{j=1}^{m-n} \lambda_j(\mathcal{M}^{\perp_N})$

Proof. (i) Minkowski's Second Theorem (Theorem 2.2.2) gives

$$\prod_{j=1}^{m-n} \lambda_j(\mathcal{M}^{\perp_N}) \leq \frac{\gamma_m^{m/2} \cdot \mathrm{Vol}(\mathcal{M}^{\perp_N})}{\prod_{j=m-n+1}^{m} \lambda_j(\mathcal{M}^{\perp_N})} \tag{4.3}$$

and we find a lower bound for $\prod_{j=m-n+1}^{m} \lambda_j(\mathcal{M}^{\perp_N})$. Theorem 2.2.3 with $\Lambda = \mathcal{M}^{\perp_N}$ and $\Lambda^\vee = N^{-1}\mathcal{M}_N$ gives $\lambda_j(\mathcal{M}^{\perp_N})\lambda_{m-j+1}(\mathcal{M}_N) \in [N, mN]$ for all $1 \leq j \leq m$. Taking the product over $j = m-n+1, \ldots, m$ yields

$$\prod_{j=m-n+1}^{m} \lambda_j(\mathcal{M}^{\perp_N}) \prod_{j=1}^{n} \lambda_j(\mathcal{M}_N) \in [N^n, (mN)^n]$$

and we conclude with Equation (4.3). To establish *(b)*, we proceed similarly. □

Therefore, an upper bound on the first $m-n$ successive minima of $\mathcal{M}^{\perp_N}$ implies an upper bound on the first n successive minima of $\mathcal{M}_N$, and vice-versa.

4.4.4 Practical discussion on Algorithm I and II

Algorithm I reveals $\overline{\mathcal{L}}$ by means of orthogonal lattices. On the other side, Algorithm II is conceptually easier than Algorithm I, in the sense that it recovers $\overline{\mathcal{L}}$ much more directly. In fact, as explained in Section 4.4.2, Algorithm II solves a "hidden sublattice problem" in the first place, by recovering the lattice $\mathcal{N}_{\mathrm{II}} \subseteq \overline{\mathcal{L}}$. We now detail the different steps of the algorithms with a practical focus.

Bases for $\mathcal{M}^{\perp_N}$ and $\mathcal{M}_N$. Given N and an input basis matrix for $\mathcal{M}$, bases for $\mathcal{M}^{\perp_N}$ and $\mathcal{M}_N$ are easily computed. For example, to compute a basis matrix for $\mathcal{M}^{\perp_N}$, given a basis matrix $M \in \mathbb{Z}^{r\times m}$ for $\mathcal{M}$, where the rows of M are a basis of $\mathcal{M}$, one may proceed as follows: write $M = [M_1|M_2]$ with $M_1 \in \mathbb{Z}^{r\times m-r}$ and $M_2 \in \mathbb{Z}^{r\times r}$. Let M_1' and M_2' be the reductions of M_1 and M_2 modulo N. If $M_2' \in \mathrm{GL}(r, \mathbb{Z}/N\mathbb{Z})$, let $M_2'^{-1}$ be its inverse and put $\tilde{M} := (-M_2'^{-1}M_1')^T$. Then the rows of the block matrix

$$B(\mathcal{M}, N) = \begin{bmatrix} 1_{m-r} & \tilde{M} \\ 0_{r\times m-r} & N \cdot 1_r \end{bmatrix} \tag{4.4}$$

is a basis matrix for $\mathcal{M}^{\perp_N}$. This is the matrix $B(\mathcal{M}, N)$ computed in the first step of Algorithm I (see Algorithm 7). Indeed, $u \in \mathbb{Z}^{1\times m}$ lies in $\mathcal{M}^{\perp_N}$ if and only if u is orthogonal modulo N to the rows of M, i.e. $Mu^T \equiv 0_{r\times 1} \pmod N$. Putting $u = (u_1, u_2) \in \mathbb{Z}^{1\times m-r} \times \mathbb{Z}^{1\times r}$, this gives $M_1u_1^T + M_2u_2^T \equiv 0_{r\times 1} \pmod N$, or equivalently, $M_2^{-1}M_1u_1^T + u_2^T \equiv 0_{r\times 1} \pmod N$, which over the integers reads as $u_2^T = \tilde{M}u_1^T + N \cdot 1_r l^T$ for some $l^T \in \mathbb{Z}^{r\times 1}$. Thus $u = (u_1, u_2) = (u_1, l) \cdot B(\mathcal{M}, N)$ is the image of (u_1, l) under $B(\mathcal{M}, N)$.

A basis matrix for the lattice $\mathcal{M}_N$ for the first step of Algorithm II (see Algorithm 8) is constructed similarly, or, one may directly use the duality: if B is a basis matrix for a lattice $\Lambda \subseteq \mathbb{Q}^m$ of full rank m, then a basis matrix for $\Lambda^\vee$ is $B^\vee := (B^T)^{-1}$, where the inverse is taken over $\mathrm{GL}(m, \mathbb{Q})$. In view of the relation $\mathcal{M}_N = N(\mathcal{M}^{\perp_N})^\vee$, a basis matrix $B'(\mathcal{M}, N)$ for $\mathcal{M}_N$ is

therefore $NA^\vee$, with $A = B(\mathcal{M}, N)$ as in Equation (4.4).

The first steps of Algorithm I and Algorithm II rely on running lattice reduction on these basis matrices, respectively. Subsequently the lattices $\mathcal{N}_{\mathrm{I}}$ and $\mathcal{N}_{\mathrm{II}}$ are constructed from those reduced bases as indicated.

The second steps of our algorithms differ more substantially. Algorithm I computes a basis of the orthogonal complement (over $\mathbb{Z}$) of $\mathcal{N}_{\mathrm{I}}$, while Algorithm II computes a basis of the completion of $\mathcal{N}_{\mathrm{II}}$. We detail these algorithms below.

Orthogonal of $\mathcal{N}_{\mathrm{I}}$. In Algorithm I, once a basis for $\mathcal{N}_{\mathrm{I}}$ is constructed, the second part computes a basis for $\mathcal{N}_{\mathrm{I}}^\perp$. This can be done using the LLL algorithm following [NS97, Theorem 4 and Algorithm 5]; also see [CSV18, Proposition 4.1]. Generally, for a lattice $\Lambda \subseteq \mathbb{Z}^m$ of rank $0 < n < m$ with basis matrix $B \in \mathbb{Z}^{n\times m}$ (with basis vectors in rows), the technique relies on LLL-reducing the rows of the matrix $\left[K_B \cdot B^T \mid 1_m\right] \in \mathbb{Z}^{(m+n)\times m}$ for a well-chosen sufficiently large constant $K_B \in \mathbb{N}$ depending on B, and then, projecting the first $m-n$ vectors of the resulting reduced basis on their last m components. For the computation of $\mathcal{N}_{\mathrm{I}}^\perp$, following [NS97, Algorithm 5], it suffices to choose the constant $K_U = \lceil 2^\ell \prod_{i=1}^{m-n} \|u_i\| \rceil$ with $\ell = (m-1)/2 + n(n-1)/4$ and where U is a basis matrix of $\mathcal{N}_{\mathrm{I}}$ with row vectors $\{u_i : 1 \le i \le m-n\}$, computed in the first part.

Completion of $\mathcal{N}_{\mathrm{II}}$. Recall that the completion of $\Lambda \subseteq \mathbb{Z}^m$ is the lattice $\overline{\Lambda} = \Lambda_{\mathbb{Q}} \cap \mathbb{Z}^m$, which is $\{v \in \mathbb{Z}^m \mid dv \in \Lambda, \text{for some } d \in \mathbb{Z}\setminus\{0\}\}$. In Algorithm II, once a basis for $\mathcal{N}_{\mathrm{II}}$ is constructed, we compute a basis for $\overline{\mathcal{N}_{\mathrm{II}}}$ in the second part. One may compute $\overline{\mathcal{N}_{\mathrm{II}}}$ as $(\mathcal{N}_{\mathrm{II}}^\perp)^\perp$ by using LLL twice, as in [NS97, Theorem 4 and Algorithm 5], and the output is then LLL-reduced.

We describe an alternative method, which in practice works well (see Section 4.9). As predicted by Lemma 4.4.2, the index of Q in $\mathcal{L}$ is N^{n-r} in most of the cases. In practical experiments with a solvable hidden lattice problem, we observe that $\mathcal{N}_{\mathrm{II}}$ is exactly $\mathcal{M}_N \cap \mathcal{L} = Q$, thus $(\mathcal{L} : \mathcal{N}_{\mathrm{II}}) = N^{n-r}$. Therefore, more directly, we can complete $\mathcal{N}_{\mathrm{II}}$ locally at primes p dividing N. For a prime p, define the *p-completion* of $\Lambda \subseteq \mathbb{Z}^m$ by

$$\Lambda^{p^\infty} := \{v \in \mathbb{Z}^m \mid p^k v \in \Lambda, \text{for some } k \in \mathbb{N}\} .$$

Let $B \in \mathbb{Z}^{n\times m}$ be a basis matrix (with rows $\{b_i\}_i$) of some lattice $\Lambda \subseteq \mathbb{Z}^m$ of rank n; assume that p divides the index $(\overline{\Lambda} : \Lambda)$. We proceed as follows to compute a basis of Λ^{p^∞}. Let $\overline{B} \in \mathbb{F}_p^{n\times m}$ be reduction of B modulo p; let $\overline{\alpha} \in \mathbb{F}_p^n$ be an element of $\ker(\overline{B})$, i.e. $\overline{\alpha}\overline{B} \equiv 0 \pmod p$. We represent $\alpha = (\alpha_i)_i \in \mathbb{Z}^n$ by choosing the entries of $\overline{\alpha}$ by their unique representatives in $\mathbb{Z} \cap [-p/2, p/2)$. We may assume that one of the coefficients of $\bar{\alpha}$ equals 1, say it is the ith coefficient. Let $x \in \mathbb{Z}^m$ such that $\alpha B = px$. Let $\Lambda' \subseteq \mathbb{Z}^m$ be the lattice generated by $B' \in \mathbb{Z}^{n\times m}$ where B' is the matrix obtained from B after replacing the ith row of B by x; then $\Lambda \subseteq \Lambda'$ and $\Lambda_{\mathbb{Q}} = \Lambda'_{\mathbb{Q}}$. By the choice of x, the rank of B' over $\mathbb{F}_\ell$ for every prime number $\ell \neq p$, does not decrease. We repeat this for every basis vector in the $\mathbb{F}_p$-kernel of $\overline{B}$ and update B' accordingly.

In Section 4.9, we report that the second step for Algorithm II can, in general, be carried out much more rapidly than the second step of Algorithm I. This also gives an improved total running time for Algorithm II against Algorithm I.

4.5 Heuristic analysis of the algorithms

We provide a heuristic analysis and comparison of Algorithms I and II for $N > 0$. For $N < 0$, it suffices to replace N by $-N$ throughout the analysis. We write $\log$ for the logarithm in base 2. Proposition 4.4.1 and 4.4.4 are the keys in our analysis. We rely on the Gaussian Heuristic (GH) for the successive minima for "random"[1] lattices (see Section 2.2.2). Accordingly, we heuristically approximate $\lambda_1(\Lambda)$ by $\sqrt{\gamma_n} \cdot \mathrm{Vol}(\Lambda)^{1/n}$. Additionnally, we heuristically assume all the minima to be approximately equal, that is,

$$\lambda_k(\Lambda) \approx \sqrt{\frac{n}{2\pi e}} \cdot \mathrm{Vol}(\Lambda)^{1/n} \ , \ 1 \leq k \leq n \ . \tag{4.5}$$

Since $\mathcal{L}, \mathcal{L}^\perp, (\mathcal{M}_\mathfrak{B})^{\perp_N}, (\mathcal{M}_\mathfrak{B})_N$ (contrary to $\mathcal{M}^{\perp_N}$ and $\mathcal{M}_N$) do not possess "small" sublattices, it is reasonable to follow this heuristic for these lattices. As $n \to \infty$, we will use the approximation $\gamma_n \approx n/(2\pi e)$.

4.5.1 Analysis of Algorithm I

Lattice reduction computes short vectors in $\mathcal{M}^{\perp_N}$; let $u_1, \ldots, u_{m-n}$ be the first $m-n$ vectors in a basis of $\mathcal{M}^{\perp_N}$ output by a lattice reduction algorithm. Since $\mathcal{M}^{\perp_N}$ contains $\mathcal{L}^\perp$, one has $\|u_{m-n}\| \leq \iota^m \lambda_{m-n}(\mathcal{L}^\perp)$ for some $\iota > 1$ depending on the lattice reduction algorithm. By Proposition 4.4.1, if

$$\iota^m \lambda_{m-n}(\mathcal{L}^\perp) < \frac{\lambda_1(\mathcal{M}_\mathfrak{B}^{\perp_N})}{\sqrt{n} \cdot \mu} \tag{4.6}$$

then $u_{m-n} \in \mathcal{L}^\perp$ and since the vectors $\{u_i\}_i$ are ordered by size, we obtain a sublattice $\mathcal{N}_\mathrm{I} = \bigoplus_{i=1}^{m-n} \mathbb{Z} u_i$ of $\mathcal{L}^\perp$ of the same rank. The orthogonal complement $\mathcal{N}_\mathrm{I}^\perp$ is then the completion of $\mathcal{L}$, by Lemma 4.3.2.

We rely on the Gaussian Heuristic to estimate $\lambda_{m-n}(\mathcal{L}^\perp)$ and $\lambda_1(\mathcal{M}_\mathfrak{B}^{\perp_N})$. Using $\mathrm{Vol}(\mathcal{L}^\perp) \leq \mathrm{Vol}(\mathcal{L})$ and Hadamard's Inequality from Lemma 4.3.2, we have $\lambda_{m-n}(\mathcal{L}^\perp) \lesssim \sqrt{(m-n)/(2\pi e)} \cdot \mu^{n/(m-n)}$. Assuming $\mathrm{Vol}(\mathcal{M}_\mathfrak{B}^{\perp_N}) = N^r$, we obtain under (GH): $\lambda_1(\mathcal{M}_\mathfrak{B}^{\perp_N}) \approx \sqrt{n/(2\pi e)} \cdot N^{r/n}$. Putting the bounds together gives

$$N^{r/n} > \iota^m \cdot (m-n)^{1/2} \cdot \mu^{\frac{m}{m-n}} \ . \tag{4.7}$$

There are more ways to read such an inequality: since our investigation is on the hidden lattice, we could either bound μ, the size of the small basis of the hidden lattice $\mathcal{L}$, as a function of the other parameters, or else, consider μ as fixed and bound the modulus N in terms of the remaining parameters. Following this latter approach, by taking logarithms, Equation (4.6) imposes the following heuristic condition on the parameters:

$$\log(N) \quad > \quad \frac{mn}{r(m-n)} \log(\mu) + \frac{mn}{r} \log(\iota) + \frac{n}{2r} \log(m-n) \tag{4.8}$$

Equation (4.8) is a heuristic sufficient condition that the chosen lattice reduction algorithm outputs $m-n$ vectors $u_1, \ldots, u_{m-n} \in \mathcal{L}^\perp$. Equation (4.8) is a heuristic sufficient condition that the chosen lattice reduction algorithm outputs $m-n$ vectors $u_1, \ldots, u_{m-n} \in \mathcal{L}^\perp$.

[1] see e.g. [Ajt06] for a precise setting; here we shall mean "generic" lattices, i.e. lattices with no extra assumptions, such as the existence of particularly small sublattices

4.5.2 Analysis of Algorithm II

We present two alternative analyses: a "direct analysis" without relying on Proposition 4.4.4, and one using Proposition 4.4.4.

Direct analysis

We run lattice reduction on $\mathcal{M}_N$; let $u_1, \ldots, u_m$ be the first n vectors of a reduced basis of $\mathcal{M}_N$. The existence of the hidden lattice $\mathcal{L}$ implies the existence of the sublattice $Q = \mathcal{M}_N \cap \mathcal{L}$ of $\mathcal{M}_N$ (defined in Equation (4.2)), which impacts the geometry of $\mathcal{M}_N$ in the following way: the first n minima of $\mathcal{M}_N$ are heuristically of the same size than the first n minima of Q, and the remaining $m - n$ minima are much larger. In particular, the first n minima of $\mathcal{M}_N$ are expected to be significantly smaller than the quantity predicted by Equation (4.5):

$$\sqrt{\gamma_m} \cdot \mathrm{Vol}(\mathcal{M}_N)^{1/m} \approx \sqrt{\gamma_m} \cdot N^{1-r/m} ,$$

which would heuristically be a valid approximation if $\mathcal{M}_N$ were a "generic" lattice (i.e. without the existence of Q). To measure this gap, we introduce a threshold constant $\theta \geq 1$. We heuristically expect to have $u_1, \ldots, u_n \in \mathcal{L}_{\mathbb{Q}}$ under the condition

$$\theta \cdot \|u_n\| < \sqrt{\gamma_m} \cdot N^{1-r/m} . \tag{4.9}$$

Since $\mathcal{M}_N$ contains Q, we have $\|u_n\| \leq \iota^m \lambda_n(Q)$ for some $\iota > 1$ depending on the lattice reduction algorithm. We assume $(\mathcal{L} : Q) = N^{n-r}$ by Lemma 4.4.2 *(b)*. Then $\mathrm{Vol}(Q) = N^{n-r}\mathrm{Vol}(\mathcal{L})$. Since $\prod_{i=1}^{n} \lambda_i(Q) \leq \gamma_n^{n/2}\mathrm{Vol}(Q)$, this gives with $\mathrm{Vol}(\mathcal{L}) \leq \mu^n$, the approximation

$$\prod_{i=1}^{n} \lambda_i(Q) \lesssim \gamma_n^{n/2} \mu^n N^{n-r} . \tag{4.10}$$

With the heuristic assumption that the successive minima of Q are roughly of equal size, this implies the heuristic upper bound

$$\lambda_i(Q) \lesssim \sqrt{\gamma_n} \mu N^{1-r/n} , \ 1 \leq i \leq n .$$

It follows that

$$\|u_n\| \lesssim \iota^m \sqrt{\gamma_n} \mu N^{1-r/n} . \tag{4.11}$$

Consequently, from Equation (4.9), we expect to have $u_1, \ldots, u_n \in \mathcal{L}_{\mathbb{Q}}$ as soon as

$$\theta \iota^m \sqrt{\gamma_n} \mu N^{1-r/n} < \sqrt{\gamma_m} N^{1-r/m} .$$

Taking logarithms, this gives the condition

$$\log(N) > \frac{mn}{r(m-n)} \log(\mu) + \frac{m}{m-n} \frac{mn}{r} \log(\iota) + \frac{mn}{r(m-n)} \log\left(\theta \frac{\sqrt{n}}{\sqrt{m}}\right) . \tag{4.12}$$

Analysis using Proposition 4.4.4.

Following Proposition 4.4.4, we compute a heuristic upper bound for $\|u_n\|$ and a lower bound for $N^{m-n}/(\mu^n \prod_{i=n+2}^{m} \lambda_i(\mathcal{M}_N))$, as $\mathrm{Vol}(\mathcal{L}_N) \geq N^{m-n}$.

The heuristic upper bound for $\|u_n\|$ is given by Equation (4.11). To establish a heuristic lower bound for $N^{m-n}/(\mu^n \prod_{i=n+2}^{m} \lambda_i(\mathcal{M}_N))$, we establish a heuristic upper bound for $\prod_{i=n+2}^{m} \lambda_i(\mathcal{M}_N)$. By Minkowski's Second Theorem,

$$\prod_{i=n+1}^{m} \lambda_i(\mathcal{M}_N) \leq \gamma_m^{m/2} N^{m-r} / \prod_{i=1}^{n} \lambda_i(\mathcal{M}_N) \ ,$$

where we have assumed that $\mathrm{Vol}(\mathcal{M}_N) = N^{m-r}$ (see Section 4.3.1), which is the generic case and heuristically (almost) always true. The first n minima of $\mathcal{M}_N$ are heuristically equal to the n minima of Q, because Q is heuristically the only relatively small sublattice of $\mathcal{M}_N$. We can heuristically consider the upper bound provided in (4.10) as a lower bound, too. Indeed, since $\mathrm{Vol}(Q) \leq \prod_{i=1}^{n} \lambda_i(Q) \leq \gamma_n^{n/2} \mathrm{Vol}(Q)$, Minkowski's bound in (4.10) is loose by a factor at most $\gamma_n^{n/2}$. Note also that assuming equality in (4.10) is compatible with Equation (4.5) for the lattice Q.

Therefore, we obtain $\prod_{i=1}^{n} \lambda_i(\mathcal{M}_N) \approx \prod_{i=1}^{n} \lambda_i(Q) \approx \gamma_n^{n/2} \mu^n N^{n-r}$. This implies that

$$\prod_{i=n+1}^{m} \lambda_i(\mathcal{M}_N) \lesssim \frac{\gamma_m^{m/2} N^{m-r}}{\gamma_n^{n/2} \mu^n N^{n-r}} \approx \frac{(2\pi e)^{n/2}}{(2\pi e)^{m/2}} \frac{m^{m/2}}{n^{n/2}} \frac{N^{m-n}}{\mu^n} =: K(m,n,N,\mu) =: K \ .$$

Since we expect the minima $\lambda_i(\mathcal{M}_N)$ for $n+1 \leq i \leq m$ to be roughly equal, we obtain $\prod_{i=n+2}^{m} \lambda_i(\mathcal{M}_N) \lesssim K^{(m-n-1)/(m-n)}$. Consequently, we derive the heuristic lower bound

$$\frac{N^{m-n}}{\mu^n \prod_{i=n+2}^{m} \lambda_i(\mathcal{M}_N)} \gtrsim \frac{N^{m-n}}{\mu^n K^{\frac{m-n-1}{m-n}}} \ ,$$

which gives,

$$\frac{N^{m-n}}{\mu^n \prod_{i=n+2}^{m} \lambda_i(\mathcal{M}_N)} \gtrsim \frac{N}{\mu^{\frac{n}{m-n}}} \cdot \left(\frac{n^{n/2}}{m^{m/2}} \right)^{\frac{m-n-1}{m-n}} \cdot \sqrt{2\pi e}^{m-n-1} \ .$$

Combined with Equation (4.11), Proposition 4.4.4 says that if

$$\iota^m \sqrt{\gamma_n} \mu N^{1-r/n} < \frac{N}{\mu^{\frac{n}{m-n}}} \cdot \left(\frac{n^{n/2}}{m^{m/2}} \right)^{\frac{m-n-1}{m-n}} \cdot \sqrt{2\pi e}^{m-n-1} \ ,$$

then $u_n \in \mathcal{L}_{\mathbb{Q}}$ (and thus $\overline{\mathcal{L}}$). Since $\{u_i\}_i$ are ordered by size, $\mathcal{N}_{\mathrm{II}} = \bigoplus_{i=1}^{n} \mathbb{Z} u_i$ is a sublattice of $\overline{\mathcal{L}}$ of rank n. Thus, the completion of $\mathcal{N}_{\mathrm{II}}$ is the completion of $\mathcal{L}$. Simplifying and taking logarithms, gives the asymptotic condition

$$\begin{aligned} \log(N) \ > \ & \frac{mn}{r(m-n)} \log(\mu) + \frac{mn}{r} \log(\iota) + \frac{n}{2r} \log(n) \\ & + \frac{n}{2r} \log\left(\frac{m^m}{n^n} \right) - \frac{n(m-n)}{2r} \log(2\pi e) \ , \end{aligned} \tag{4.13}$$

where we have used the mild approximation $m-n-1 \approx m-n$. Equation (4.13) is a heuristic sufficient condition that the chosen lattice reduction algorithm outputs n vectors in $\mathcal{L}_{\mathbb{Q}} \cap \mathbb{Z}^m = \overline{\mathcal{L}}$.

In Section 4.5.3, we will see that, asymptotically, the heuristic bounds for Algorithms I and II perform very similarly.

4.5.3 Parameter comparison of Algorithms I and II

In light of Equations (4.8) and (4.13) we deduce that $\log(N) > \frac{mn}{r(m-n)} \log(\mu)$, and therefore we heuristically expect both algorithms to detect μ-small lattices of size μ approximately

$$\mu = O(N^{\frac{r(m-n)}{nm}}) \,, \tag{4.14}$$

when r, n, m are fixed and N tends to infinity. Since $r < n$ and $m - n < m$, the exponent is strictly less than 1. For example, in the balanced case $m = 2n = 4r$, this gives $\mu = O(N^{1/4})$. Further, we see that larger values of r make the hidden lattice problem easier (as is expected, since more information is public) as it can be solved with a modulus of r times smaller bitsize.

We now turn to a more detailed comparison of Equation (4.8) and Equation (4.12). For fixed m, n, r, μ, a sufficiently large value of N satisfies (4.8), resp. (4.12). When m, n, r are considered as constants, then the right-hand sides of (4.8) and (4.12) differ only by a constant.

To study the value of N asymptotically as $n \to \infty$, we consider the rank r as constant, and view m as a function of n. The term $\log(\iota)$ is constant and relatively small; for example, in practice one achieves a root Hermite factor ι approximately 1.021 for LLL, so $\log(\iota) \approx 0.03$ is of impact only in large dimensions. Table 4.2 shows three cases: when $m - n = O(1)$ is bounded absolutely (independently of n), when $m = O(n)$, and last, and when $m = O(n^\ell)$ for $\ell > 1$.

	$\log(N)$	
m	Algorithm I	Algorithm II
$n + O(1)$	$O(\frac{n^2}{r} \log(\mu))$	$O(\frac{n^2}{r} \max(\log(\mu), n))$
$O(n)$	$O(\frac{n}{r} \max(\log(\mu), n))$	$O(\frac{n}{r} \max(\log(\mu), n))$
$O(n^\ell), \ell \in \mathbb{R}_{>1}$	$O(\frac{n}{r} \max(\log(\mu), n^\ell))$	$O(\frac{n}{r} \max(\log(\mu), n^\ell))$

Table 4.2: Asymptotic lower bounds for $\log(N)$ as functions of n, r, μ

When $m - n = O(1)$, our algorithms heuristically require larger (asymptotically equal) values of N. The last line of Table 4.2 remains meaningful for $\ell = 1$ and recovers the case $m = O(n)$; we have separated both for better readability. Alternatively, we may rewrite (4.8) and (4.13) as

$$\Delta := \log \left(\frac{N^{r/n}}{\mu^{m/(m-n)}} \right) > \Delta_*(n, m, \iota) \quad , \tag{4.15}$$

where $\Delta_*(n, m, \iota)$ with $* \in \{\mathrm{I}, \mathrm{II}\}$ (depending on whether the bound stands for Algorithm I

or II) are the functions depending on n, m and ι, defined by:

$$\Delta_{\mathrm{I}}(n,m,\iota) = m\log(\iota) + \frac{1}{2}\log(m-n) \tag{4.16}$$

$$\Delta_{\mathrm{II}}(n,m,\iota) = \frac{m^2}{m-n}\log(\iota) + \frac{m}{m-n}\log\left(\theta\frac{\sqrt{n}}{\sqrt{m}}\right) \tag{4.17}$$

We consider ι as a constant once the lattice reduction algorithm is chosen, and treat $m = m(n)$ as a function of n. Thus we just write $\Delta_*(n)$ as function of n only. The number Δ is regarded as an arithmetic invariant for the (geometric) hidden lattice problem, depending on all the parameters of the problem.

Remark 4.5.1. In the language of knapsack-type problems, Δ^{-1} can be regarded as a density for the HLP. Namely, one commonly attributes a density to knapsack-type problems as a measure of their hardness. For the classical "binary" subset sum problem [LO85], asking to reveal $x_1, \ldots, x_n \in \{0,1\}$ from a given sum $\alpha = \sum_{i=1}^n \alpha_i x_i$ with given $\alpha_1, \ldots, \alpha_n \in \mathbb{Z}$, the density is defined as $n/\log(\max_i \alpha_i)$. When the $\{x_i\}_i$ are not binary, [NS05] argues that this definition is not "complete" enough, and introduces the notion of "pseudo-density", as $(\sum_i x_i^2)\cdot n/\log(\max_i \alpha_i)$, taking into account the weights $\{x_i\}_i$. In [PZ11], the authors study higher-dimensional subset sums where $k \geq 1$ equations are given instead of only one; thereby the density is generalized as $(1/k)\cdot n/\log(\max_i \alpha_i)$. For the hidden subset sum problem [NS99] (see also Section 2.2.3.1 of Chapter 2 and 4.8.2 of this chapter), asking to reveal vectors $x_1, \ldots, x_n \in \{0,1\}^m$ and weights $\alpha_1, \ldots, \alpha_n$ from a given vector $v \equiv \sum_{i=1}^n \alpha_i x_i \pmod N$, the density has been defined as $n/\log(N)$, which, however, is independent of the dimension m. In light of this discussion, we believe that the definition of Δ^{-1} is a more complete definition for a density of the HLP. In fact, for large enough m (say $m \to \infty$) together with $r = 1$, our definition (4.15) roughly recovers that of [NS99] since $\Delta^{-1} \to 1/\log(N^{1/n}/\mu) = n/\log(N/\mu^n)$, as $m \to \infty$. Our bounds show that heuristically our algorithms are more likely to succeed for larger values of Δ (i.e. larger gaps between N and μ).

Proposition 4.5.2. *(a) Let $m = \ell n$ for $\ell > 1$. Then $\Delta_{\mathrm{I}}(n) = O(n)$ and $\Delta_{\mathrm{II}}(n) = O(n)$.*
(b) Let $m = n^\ell$ for $\ell > 1$. Then $\Delta_{\mathrm{I}}(n) = O(n^\ell)$ and $\Delta_{\mathrm{II}}(n) = O(n^\ell)$.

The proof is immediate from growth comparisons in the formulae (4.16) and (4.17).

4.5.4 Complexity of lattice reduction

The computations of $\mathcal{N}_{\mathrm{I}}$ and $\mathcal{N}_{\mathrm{II}}$ are carried out by lattice reduction. We describe their complexity by the LLL and BKZ algorithm. We see that the LLL reduction (L^2-reduction) step in Algorithm II is faster than in Algorithm I when $r \geq m/2$.

When applying the L^2-algorithm with naive integer multiplication in Algorithm I, it is run on a basis of $\mathcal{M}^{\perp_N}$, given by the rows of the matrix B in Equation (4.4). The top right block $\tilde{M}$ in B has size $(m-r)\times r$ and entries of size at most N, so, every row in B has Euclidean norm at most $\max((rN^2+1)^{1/2}, N) = (rN^2+1)^{1/2}$. This gives, as stated in Section 2.2.2, Equation (2.3), a complexity $O(m^6\log((rN^2+1)^{1/2}) + m^5\log^2((rN^2+1)^{1/2}))$ which approximately is

$$O(m^6\log(r^{1/2}N) + m^5\log^2(r^{1/2}N))\,,$$

for computing $\mathcal{N}_{\mathrm{I}}$ by the L^2-algorithm. For Algorithm II, the L^2-algorithm is run on the basis matrix $NB^\vee$ of $\mathcal{M}_N$, with $N \cdot 1_{m-r}$ in the top left corner. The rows have Euclidean norm at most $\max(N, ((m-r)N^2+1)^{1/2}) = ((m-r)N^2+1)^{1/2}$. This gives a time complexity $O(m^6 \log(((m-r)N^2+1))^{1/2} + m^5 \log^2(((m-r)N^2+1))^{1/2})$, which approximately is

$$O(m^6 \log((m-r)^{1/2}N) + m^5 \log^2((m-r)^{1/2}N)) \ ,$$

for computing $\mathcal{N}_{\mathrm{II}}$ by the L^2-algorithm. In particular, this complexity is lower than that for computing $\mathcal{N}_{\mathrm{I}}$ when $r \geq m/2$. In the common case that $r = 1$, the computation of $\mathcal{N}_{\mathrm{I}}$ is thus faster than that of $\mathcal{N}_{\mathrm{II}}$, which we confirm practically in Section 4.9.

When the prime factorization of N is known and p denotes the smallest prime factor of N, then the complexity can be reduced by replacing N by p in the aforementioned formulae, provided that $\log(p)$ satisfies the (heuristic) bounds (4.8) and (4.12), respectively, and by performing the first steps of the algorithms over $\mathbb{Z}/p\mathbb{Z}$ instead of $\mathbb{Z}/N\mathbb{Z}$. Namely, in that case $\mathcal{M} \subseteq \mathcal{L} \pmod N$ implies also $\mathcal{M} \subseteq \mathcal{L} \pmod p$, which in the first step of Algorithm I and II, leads to consider the lattices $\mathcal{M}^{\perp_p}$ and $\mathcal{M}_p$, respectively.

When using BKZ lattice reduction, we rely on our heuristic analyses to obtain a lower bound on the complexity for computing $\mathcal{N}_{\mathrm{I}}$ and $\mathcal{N}_{\mathrm{II}}$. A root Hermite factor ι is achieved within time at least $2^{\Theta(1/\log(\iota))}$ by using BKZ with block-size $\Theta(1/\log(\iota))$ ([HPS11a], see also Section 2.2.2). For both algorithms, $\log(\iota) < \frac{r}{mn}\log(N)$ gives a heuristic time complexity $2^{\Theta(mn/(r\log(N)))}$ to compute $\mathcal{N}_{\mathrm{I}}$, resp. $\mathcal{N}_{\mathrm{II}}$, with BKZ.

4.6 Theoretical analysis by counting

In this section, we restrict to the most basic case of $r = 1$, i.e. the lattice $\mathcal{M}$ is generated by a single vector, lying in a hidden μ-small lattice $\mathcal{L}$ modulo N. As we will see in Section 4.8, this is the case in many concrete applications of the hidden lattice problem. We obtain two results (see Theorems 4.6.1 and 4.6.2 below) concerning the success rate of Algorithms I and II, respectively, on a large class of hidden lattice problems, under a well-chosen set of parameters. Our proof technique is a discrete counting argument. We then compare our results with the heuristic analysis from Section 4.5.

4.6.1 Notation and main results

We fix $n, m \in \mathbb{Z}_{\geq 2}$ with $m > n$ and $\mu \in \mathbb{R}_{\geq 1}$, $N \in \mathbb{Z}_{>0}$. Let $\Omega := \Omega(n, m, \mu)$ be the set of collections $\mathfrak{B} = \{v_i\}_i$ of n $\mathbb{Z}/N\mathbb{Z}$-linearly independent vectors in $\mathbb{Z}^m$ satisfying $\sigma(\mathfrak{B}) := (\frac{1}{n}\sum_i \|v_i\|^2)^{1/2} \leq \mu$. For $\mathfrak{B} \in \Omega$, let $\mathcal{L}(\mathfrak{B})$ be the μ-small lattice generated by $\mathfrak{B}$; this is the 'hidden lattice'. Consider the homomorphism

$$F_{\mathfrak{B}} : (\mathbb{Z}/N\mathbb{Z})^n \to (\mathbb{Z}/N\mathbb{Z})^m \ , \ a = (a_i)_i \mapsto F_{\mathfrak{B}}(a) := \sum_i a_i \pi_N(v_i) \ ,$$

where $\pi_N : \mathbb{Z}^m \to (\mathbb{Z}/N\mathbb{Z})^m$ is reduction modulo N. Let $\mathcal{M}(a)$ be the lattice $\mathbb{Z}F_{\mathfrak{B}}(a)$ of rank 1 with generator $F_{\mathfrak{B}}(a)$. By construction, $\mathcal{M}(a) \subseteq \mathcal{L}(\mathfrak{B}) \pmod N$ defines a Hidden Lattice

Problem, asking to compute a basis of $\overline{\mathcal{L}(\mathfrak{B})}$ on input $\mathcal{M}(a)$ and N (and n). We identify $F_{\mathfrak{B}}(a)$ with this problem and our sample space for the hidden lattice problems is

$$\mathcal{H}(\mathfrak{B}) = \{F_{\mathfrak{B}}(a) \mid a \in (\mathbb{Z}/N\mathbb{Z})^n\} .$$

Clearly, $\#\mathcal{H}(\mathfrak{B}) = N^n$. For $\delta \in (1/4, 1]$, denote by $\mathcal{H}_{\delta,\mathrm{I}}(\mathfrak{B}) \subseteq \mathcal{H}(\mathfrak{B})$ (resp. $\mathcal{H}_{\delta,\mathrm{II}}(\mathfrak{B}) \subseteq \mathcal{H}(\mathfrak{B})$), the subset of $\mathcal{H}(\mathfrak{B})$ for which Algorithm I (resp. Algorithm II) succeeds by using δ-LLL in the first step. We shall prove the following results, where Corollary 4.6.3 is a direct consequence of the theorems.

Theorem 4.6.1. *Let $\mu \in \mathbb{R}_{\geq 1}$ and $m > n \geq 3$ and $N > 0$ be integers. Let $\delta \in (1/4, 1)$, $c = (\delta - 1/4)^{-1}$ and $\varepsilon \in (0, 1)$ such that*

$$\begin{aligned} \log(N\varepsilon) \quad > \quad & \frac{mn}{2}\log(c) + n(n+1)\log(\mu) \\ & + \frac{n(m-n)}{2}\log((2/3)(m-n)) + n\log(3\sqrt{n}) + 1 \end{aligned} \tag{4.18}$$

For every $\mathfrak{B} \in \Omega$, at least $(1-\varepsilon)\#\mathcal{H}(\mathfrak{B})$ of the hidden lattice problems from $\mathcal{H}(\mathfrak{B})$ are solvable by Algorithm I *with δ-*LLL*; i.e.*

$$\frac{\#\mathcal{H}_{\delta,\mathrm{I}}(\mathfrak{B})}{\#\mathcal{H}(\mathfrak{B})} \geq 1 - \varepsilon .$$

Theorem 4.6.2. *Let $\mu \in \mathbb{R}_{\geq 1}$ and $m > n \geq 2$ and $N > 0$ be integers. Let $\delta \in (1/4, 1)$, $c = (\delta - 1/4)^{-1}$ and $\varepsilon \in (0, 1)$ such that*

$$\begin{aligned} \log(N\varepsilon) \quad > \quad & \frac{mn}{2}\log(c) + n(n+2)\log(\mu) \\ & + n\log(3n^2) + 1 \end{aligned} \tag{4.19}$$

For every $\mathfrak{B} \in \Omega$, at least $(1-\varepsilon)\#\mathcal{H}(\mathfrak{B})$ of the hidden lattice problems constructed from $\mathcal{H}(\mathfrak{B})$ are solvable by Algorithm II *with δ-*LLL*; i.e.*

$$\frac{\#\mathcal{H}_{\delta,\mathrm{II}}(\mathfrak{B})}{\#\mathcal{H}(\mathfrak{B})} \geq 1 - \varepsilon .$$

Corollary 4.6.3. *Let $m > n \geq 3$. For every $\delta \in (1/4, 1)$ and $\varepsilon \in (0, 1)$, there exist positive real numbers $N_{\mathrm{I}}^{\dagger} = N_{\delta,\mu,n,m}(\varepsilon)$ and $N_{\mathrm{II}}^{\dagger} = N_{\delta,\mu,n,m}(\varepsilon)$ depending on $n, m, \mu, \delta, \varepsilon$, such that for all integers $N > \min(N_{\mathrm{I}}^{\dagger}, N_{\mathrm{II}}^{\dagger})$ and all $\mathfrak{B} \in \Omega$, at least $(1-\varepsilon)\#\mathcal{H}(\mathfrak{B})$ of the hidden lattice problems from $\mathcal{H}(\mathfrak{B})$ are solvable (by Algorithm* I *if $\min(N_{\mathrm{I}}^{\dagger}, N_{\mathrm{II}}^{\dagger}) = N_{\mathrm{I}}^{\dagger}$ and Algorithm* II *otherwise) using δ-*LLL*.*

4.6.2 Proof of Theorem 4.6.1

Fix integers[2] $m > n \geq 3$, $N > 0$ and $\mu \in \mathbb{R}_{\geq 1}$. It is enough to show that under the assumption in Equation (4.18), we can compute a sublattice $\mathcal{N}_{\mathrm{I}}$ of $\mathcal{L}(\mathfrak{B})^{\perp}$ of rank $m - n$. A basis for the orthogonal complement of $\mathcal{N}_1$ then gives a basis of $\mathcal{L}(\mathfrak{B})$. To prove Theorem 4.6.1 we proceed in three steps, described in subsequent subsections. Given $a \in (\mathbb{Z}/N\mathbb{Z})^n$, we establish a lower bound for $\lambda_1((\mathbb{Z}a)^{\perp_N})$ (see Step 1) and then an upper bound for $\|u_{m-n}\|$ where $\{u_i\}_i$ is a δ-LLL reduced basis of $(\mathbb{Z}F_{\mathfrak{B}}(a))^{\perp_N}$ (see Step 2) We conclude the proof by combining with Proposition 4.4.1 (see Step 3).

[2] the condition $n \geq 3$ is used for Lemma 4.6.5.

4.6.2.1 Step 1

For a lower bound for $\lambda_1((\mathbb{Z}a)^{\perp_N})$, we use a counting argument. For an element $t = (t_1, \ldots, t_n)$ of $\mathbb{Z}^n$, let $\gcd(t, N) := \gcd(t_1, \ldots, t_n, N)$.

Lemma 4.6.4. *For every non-zero vector $t \in \mathbb{Z}^n$ with $d = \gcd(t, N)$, one has*

$$\#\{a \in (\mathbb{Z}/N\mathbb{Z})^n \mid \langle a, t\rangle \equiv 0 \pmod N\} = dN^{n-1} .$$

Proof. If $d = 1$, then the set in the statement is the kernel of the surjective (as $\gcd(t, N) = 1$) homomorphism $\varphi_t : (\mathbb{Z}/N\mathbb{Z})^n \to \mathbb{Z}/N\mathbb{Z}, a \mapsto \langle a, t\rangle$ with $\#\ker(\varphi_t) = N^{n-1}$. If $d > 1$, let $t' = (1/d)t$. Then $\langle a, t\rangle \equiv 0 \pmod N$ if and only if $\langle a, t'\rangle \equiv 0 \pmod{N/d}$, and we represent a as $a_1 + (N/d)a_2$ with $a_1 \in (\mathbb{Z}/(N/d)\mathbb{Z})^n$ and $a_2 \in (\mathbb{Z}/d\mathbb{Z})^n$. The number of such a with $\langle a_1, t'\rangle \equiv 0 \pmod{N/d}$ is $(N/d)^{n-1} \cdot d^n$. □

For $R > 0$, let $B_n(R)$ be the n-dimensional closed ball of radius R centered at the origin and let $S_n(R) = \#\{x \in \mathbb{Z}^n \mid \|x\| \le R\}$ the number of integral points in $B_n(R)$. We rely on the simple upper bound $S_n(R) \le (2R + 1)^n \le (3R)^n$ if $R \ge 1$.

Lemma 4.6.5. *For $\varepsilon \in (0, 1)$, let $k_\varepsilon := k_\varepsilon(n, N) = \frac{1}{3}(\frac{6\varepsilon}{\pi^2})^{1/n} N^{1/n}$. Then*

$$\frac{1}{N^n} \cdot \#\{a \in (\mathbb{Z}/N\mathbb{Z})^n \mid \lambda_1((\mathbb{Z}a)^{\perp_N}) > k_\varepsilon\} \ge 1 - \varepsilon$$

Proof. For $R > 0$, let $\alpha_n(R) = N^{-n} \cdot \#\{a \in (\mathbb{Z}/N\mathbb{Z})^n \mid \lambda_1((\mathbb{Z}a)^{\perp_N}) \le R\}$; we prove $\alpha_n(k_\varepsilon) \le \varepsilon$. Without loss of generality, we let $1 \le R < N$ as the vectors $\{Ne_i\}_i$ (for the canonical basis $\{e_i\}_i$) of norm N lie in $(\mathbb{Z}a)^{\perp_N}$, and so $\alpha_n(R) = 1$ for $R \ge N$. Then $N^n\alpha_n(R) = \#\{a \in (\mathbb{Z}/N\mathbb{Z})^n \mid \exists\, t \in B_n(0, R) \cap \mathbb{Z}^n \setminus \{0\}, \langle a, t\rangle \equiv 0 \pmod N\}$ is upper bounded by $\sum_t \#\{a \mid \langle a, t\rangle \equiv 0 \pmod N\}$, which we write as

$$\sum_{d|N, d\neq N} \left(\sum_{t, \gcd(t,N)=d} \#\{a \mid \langle a, t\rangle \equiv 0 \pmod N\} \right) \tag{4.20}$$

where t runs over $B_n(0, R) \cap \mathbb{Z}^n \setminus \{0\}$. Note that in the outer sum we omit $d = N$ as $\|t\| \le R < N$ and therefore every entry of t is less than N. We estimate the number of terms in the inner sum for a given divisor d of N. By dividing every entry of t by d we have $\#\{t \in \mathbb{Z}^n \setminus \{0\} : \|t\| \le R, \gcd(t, N) = d\} \le S_n(0, R/d) \le 3^n(R/d)^n$, if $R \ge d$. Otherwise, the same bound still holds, because we we count non-zero points. Using Lemma 4.6.4, one has $\#\{a \mid \langle a, t\rangle \equiv 0 \pmod N\} = dN^{n-1}$ for vectors t with $\gcd(t, N) = d$. Finally, the sum in Equation (4.20) is at most

$$\begin{aligned} 3^n \sum_d (R/d)^n (dN^{n-1}) &= 3^n R^n N^{n-1} \sum_d d^{1-n} \\ &\le 3^n R^n N^{n-1} \sum_{d \ge 1} d^{-2} \\ &= 3^n R^n N^{n-1} \pi^2/6 , \end{aligned}$$

because $n \ge 3$. Hence $\alpha_n(R) \le 3^n R^n \pi^2/(6N)$. Choosing R to be equal to $R_\varepsilon := \frac{1}{3}(6N\varepsilon/\pi^2)^{1/n}$ gives $\alpha_n(R_\varepsilon) \le \varepsilon$. In conclusion, letting $k_\varepsilon = \min(N, R_\varepsilon) = R_\varepsilon$, gives the result. □

4.6.2.2 Step 2

To $(\mathfrak{B}, a) \in \Omega \times (\mathbb{Z}/N\mathbb{Z})^n$, we associate the vector $F_{\mathfrak{B}}(a)$, which we identify with an HLP. The first step of Algorithm I computes a reduced basis of $(\mathcal{M}(a))^{\perp_N}$. For $\delta \in (1/4, 1]$, we consider a δ-LLL reduced basis $\{u_i^{(\mathfrak{B},a,\delta)}\}_i$ of $\mathcal{M}(a)^{\perp_N}$. We establish an upper bound for $\|u_{m-n}^{(\mathfrak{B},a,\delta)}\|$, by Minkowski's Second Theorem and a counting argument similar to Step 1. Using $\mathcal{L}(\mathfrak{B})^{\perp} \subseteq (\mathcal{M}(a))^{\perp_N}$, δ-LLL (Theorem 2.2.5) outputs vectors $\{u_i^{(\mathfrak{B},a,\delta)}\}_i$ such that

$$\|u_{m-n}^{(\mathfrak{B},a,\delta)}\| \leq c^{(m-1)/2} \lambda_{m-n}(\mathcal{L}(\mathfrak{B})^{\perp}) \tag{4.21}$$

where $c = (\delta - 1/4)^{-1}$. We obtain an upper bound for $\lambda_{m-n}(\mathcal{L}(\mathfrak{B})^{\perp})$ by Minkowski's Second Theorem (Theorem 2.2.2):

$$\lambda_{m-n}(\mathcal{L}(\mathfrak{B})^{\perp}) \leq \prod_{i=1}^{m-n} \lambda_i(\mathcal{L}(\mathfrak{B})^{\perp}) \leq ((2/3)(m-n))^{(m-n)/2} \mathrm{Vol}(\mathcal{L}(\mathfrak{B})^{\perp}) \,, \tag{4.22}$$

which gives $\lambda_{m-n}(\mathcal{L}(\mathfrak{B})^{\perp}) \leq ((2/3)(m-n))^{(m-n)/2} \mu^n$, since $\mathrm{Vol}(\mathcal{L}(\mathfrak{B})^{\perp}) \leq \mathrm{Vol}(\mathcal{L}(\mathfrak{B})) \leq \mu^n$ (Lemma 4.3.2). This gives

$$\|u_{m-n}^{(\mathfrak{B},a,\delta)}\| \leq c^{(m-1)/2} ((2/3)(m-n))^{(m-n)/2} \mu^n \tag{4.23}$$

for every $a \in (\mathbb{Z}/N\mathbb{Z})^n$.

4.6.2.3 Step 3: Proof of Theorem 4.6.1

Let $\mathfrak{B} \in \Omega$ and $\varepsilon \in (0, 1)$. We continue to use the notation k_ε introduced above. Equation (4.18) implies that $\log(N\varepsilon)$ is strictly larger than

$$\frac{n(m-1)}{2} \log(c) + n(n+1)\log(\mu) + \frac{n(m-n)}{2} \log((2/3)(m-n)) + n\log(3\sqrt{n}) + \log(\pi^2/6) \,;$$

and it is a direct computation to see that this is equivalent to

$$c^{(m-1)/2} ((2/3)(m-n))^{(m-n)/2} \mu^n < k_\varepsilon / (\sqrt{n}\mu) \,. \tag{4.24}$$

By Lemma 4.6.5, $k_\varepsilon < \lambda_1((\mathbb{Z}a)^{\perp_N})$ for at least $(1-\varepsilon)N^n$ choices of $a \in (\mathbb{Z}/N\mathbb{Z})^n$. By Equation (4.23), $c^{(m-1)/2}((2/3)(m-n))^{(m-n)/2}\mu^n$ is an upper bound for $\|u_{m-n}^{(\mathfrak{B},a,\delta)}\|$ where $\{u_i^{(\mathfrak{B},a,\delta)}\}_i$ is a δ-LLL reduced basis of $\mathcal{M}(a)^{\perp_N}$ for every a. Hence, for at least $(1-\varepsilon)N^n$ choices of $a \in (\mathbb{Z}/N\mathbb{Z})^n$, Equation (4.24) implies $\|u_{m-n}^{(\mathfrak{B},a,\delta)}\| < \lambda_1((\mathbb{Z}a)^{\perp_N})/(\sqrt{n}\mu)$. Proposition 4.4.1 gives $u_i^{(\mathfrak{B},a,\delta)} \in \mathcal{L}(\mathfrak{B})^{\perp}$ for all $1 \leq i \leq m-n$. This terminates the proof.

4.6.3 Proof of Theorem 4.6.2

Fix integers $m > n \geq 3$, $N > 0$ and $\mu \in \mathbb{R}_{\geq 1}$. It is enough to show that under the assumption in (4.19), we can compute a sublattice $\mathcal{N}_{\mathrm{II}}$ of $\overline{\mathcal{L}(\mathfrak{B})}$ of rank n. A basis for $\overline{\mathcal{N}_{\mathrm{II}}}$ then gives a basis of $\overline{\mathcal{L}(\mathfrak{B})}$. To prove Theorem 4.6.2, we again proceed in three steps, similarly to the proof of Thm. 4.6.1. Given $a \in (\mathbb{Z}/N\mathbb{Z})^n$, we first establish an upper bound for $\|u_n\|$, where $\{u_i\}_i$ is a δ-LLL reduced basis of $\mathcal{M}(a)_N$. We conclude the proof using Proposition 4.4.4.

4.6.3.1 Step 1

For a given $a \in (\mathbb{Z}/N\mathbb{Z})^n$, we consider a δ-LLL reduced basis $\{u_i^{(\mathfrak{B},a,\delta)}\}_i$ of $\mathcal{M}(a)$. Note that by construction $(\mathcal{M}(a)_{\mathfrak{B}})_N = (\mathbb{Z}a)_N$. The lattice $Q(a) = \mathcal{M}(a)_N \cap \mathcal{L}(\mathfrak{B})$ is defined as in Section 4.4 (see Lemma 4.4.2). The following lemma gives an upper bound for $\|u_n^{(\mathfrak{B},a,\delta)}\|$ for almost all $a \in (\mathbb{Z}/N\mathbb{Z})^n$.

Lemma 4.6.6. *For $\varepsilon \in (0,1)$, let $\ell_\varepsilon := \ell_\varepsilon(n,N) = 3n(\pi^2/(6\varepsilon))^{1/n} N^{1-1/n}$. Then*

$$\frac{1}{N^n} \cdot \#\{a \in (\mathbb{Z}/N\mathbb{Z})^n \mid \|u_n^{(\mathfrak{B},a,\delta)}\| < c^{(m-1)/2} n\mu^2 \ell_\varepsilon\} \geq 1 - \varepsilon \,.$$

Proof. Let $a \in (\mathbb{Z}/N\mathbb{Z})^n$. As $Q(a) \subseteq \mathcal{M}(a)_N$, we have

$$\|u_n^{(\mathfrak{B},a,\delta)}\| \leq c^{(m-1)/2} \lambda_n(Q(a)) \,. \tag{4.25}$$

The lattice $Q(a)$ contains the n "short" vectors $q_1 = c_{\mathfrak{B},N}^{-1}(x^{(1)}), \ldots, q_n = c_{\mathfrak{B},N}^{-1}(x^{(n)})$ with $\|x^{(j)}\| = \lambda_j((\mathbb{Z}a)_N)$ for $1 \leq j \leq n$. With $\mathfrak{B} = \{v_1, \ldots, v_n\}$, we can write, for every $1 \leq j \leq n$, $q_j = \sum_{i=1}^n x_i^{(j)} v_i$ with $x^{(j)} = (x_1^{(j)}, \ldots, x_n^{(j)}) \in (\mathbb{Z}a)_N$. Therefore, for all $1 \leq j \leq n$,

$$\|q_j\| \leq \sum_{i=1}^n |x_i^{(j)}| \|v_i\| \leq \sum_{i=1}^n \lambda_j((\mathbb{Z}a)_N)\|v_i\| \leq \lambda_n((\mathbb{Z}a)_N) \sum_{i=1}^n \|v_i\|^2 \,. \tag{4.26}$$

This implies, since $\mathfrak{B}$ is μ-small,

$$\lambda_n(Q(a)) \leq \max_{1 \leq j \leq n} \|q_j\| \leq \lambda_n((\mathbb{Z}a)_N) n\mu^2 \,. \tag{4.27}$$

Theorem 2.2.3 applied with $\Lambda = (\mathbb{Z}a)_N$ and $\Lambda^\vee = N^{-1}(\mathbb{Z}a)^{\perp_N}$ implies that

$$\lambda_n((\mathbb{Z}a)_N) \leq \frac{nN}{\lambda_1((\mathbb{Z}a)^{\perp_N})} \,.$$

By Lemma 4.6.5, $\lambda_1((\mathbb{Z}a)^{\perp_N}) > k_\varepsilon = \frac{1}{3}(6\varepsilon/\pi^2)^{1/n} N^{1/n}$ for at least $(1-\varepsilon)N^n$ choices of $a \in (\mathbb{Z}/N\mathbb{Z})^n$. Therefore, $\lambda_n((\mathbb{Z}a)_N) < nN/k_\varepsilon = 3n(\pi^2/(6\varepsilon))^{1/n} N^{1-1/n} = \ell_\varepsilon$ for at least $(1-\varepsilon)N^n$ choices of $a \in (\mathbb{Z}/N\mathbb{Z})^n$. The bound for $\|u_n^{(\mathfrak{B},a,\delta)}\|$ then follows by combining Equations (4.25) and (4.27). □

4.6.3.2 Step 2

We now compute a lower bound for the right-hand side of the formula in Proposition 4.4.4, for every $a \in (\mathbb{Z}/N\mathbb{Z})^n$. We clearly have:

$$\frac{1}{\mu^n} \cdot \frac{\mathrm{Vol}(\mathcal{L}(\mathfrak{B})_N)}{\prod_{i=n+2}^m \lambda_i(\mathcal{M}(a)_N)} \geq \frac{1}{\mu^n} \cdot \frac{N^{m-n}}{\prod_{i=n+2}^m \lambda_i(\mathcal{M}(a)_N)} \,. \tag{4.28}$$

Since $N\mathbb{Z}^m \subseteq \mathcal{M}(a)_N$, we have $\lambda_i(\mathcal{M}(a)_N) \leq N$ for every $1 \leq i \leq m$. Thereby

$$\prod_{i=n+2}^m \lambda_i(\mathcal{M}(a)_N) \leq N^{m-n-1} \,,$$

which in Equation (4.28) gives, for every $a \in (\mathbb{Z}/N\mathbb{Z})^n$:

$$\frac{1}{\mu^n} \cdot \frac{\mathrm{Vol}(\mathcal{L}(\mathfrak{B})_N)}{\prod_{i=n+2}^m \lambda_i(\mathcal{M}(a)_N)} \geq \frac{N}{\mu^n} \,. \tag{4.29}$$

4.6.3.3 Proof of Theorem 4.6.2.

Let $\mathfrak{B} \in \Omega$ and $\varepsilon \in (0,1)$. The assumption in Equation (4.19) implies that

$$\log(N\varepsilon) > \frac{(m-1)n}{2}\log(c) + n(n+2)\log(\mu) + n\log(3n^2) + \log(\pi^2/6)\ ,$$

which by a direct computation is equivalent to

$$c^{(m-1)/2} n\mu^2 \ell_\varepsilon < \frac{N}{\mu^n}\ , \tag{4.30}$$

where $\ell_\varepsilon = 3n(\pi^2/(6\varepsilon))^{1/n} N^{1-1/n}$ is as in Lemma 4.6.6. By Lemma 4.6.6, the left-hand side is an upper bound for $\|u_n^{(\mathfrak{B},a,\delta)}\|$ for at least $(1-\varepsilon)N^n$ of the choices of a, where $\{u_i^{(\mathfrak{B},a,\delta)}\}_i$ is a δ-LLL reduced basis of $\mathcal{M}(a)_N$. By Equation (4.29) the right-hand side is a lower bound for $\frac{1}{\mu^n} \cdot \frac{\mathrm{Vol}(\mathcal{L}(\mathfrak{B})_N)}{\prod_{i=n+2}^{m} \lambda_i(\mathcal{M}(a)_N)}$, for every $a \in (\mathbb{Z}/N\mathbb{Z})^n$. Hence, for at least $(1-\varepsilon)N^n$ of $a \in (\mathbb{Z}/N\mathbb{Z})^n$, Equation (4.19) and Proposition 4.4.4 give $u_i^{(\mathfrak{B},a,\delta)} \in \overline{\mathcal{L}(\mathfrak{B})}$ for all $1 \le i \le n$. This terminates the proof.

4.6.4 Comparison

We first compare Theorem 4.6.1 with Theorem 4.6.2. Table 4.3 summarizes the asymptotic lower bounds for $\log(N)$. It appears that Algorithm II achieves slightly better asymptotic bounds.

	$\log(N)$	
m	Algorithm I (Thm. 4.6.1)	Algorithm II (Thm. 4.6.2)
$n + O(1)$	$O(n^2 \log(\mu))$	$O(n^2 \log(\mu))$
$O(n)$	$O(n^2 \max(\log(\mu), \log(n)))$	$O(n^2 \log(\mu))$
$O(n^\ell), \ell \in \mathbb{R}_{>1}$	$O(n^2 \max(\log(\mu), n^{\ell-1}\log(n)))$	$O(n^2 \max(\log(\mu), n^{\ell-1}))$

Table 4.3: Asymptotic lower bounds for $\log(N)$ as functions of n, μ

We compare Theorem 4.6.1 and Theorem 4.6.2 with the heuristic estimates in Section 4.5 (with $r = 1$). The terms in $\log(c)$ are to be compared with those in $\log(\iota)$ in the heuristic analysis. As our proofs build upon non-tight upper bounds (e.g. Minkowski bounds, or the number of integral points in spheres), our proven formulae are expectedly weaker. The main difference between our heuristic and theoretical lower bounds for $\log(N)$ occurs in the term containing $\log(\mu)$. In the case of Algorithm I, this difference comes from our upper bound for the last minimum of $\mathcal{L}^\perp$ by Minkowski's Second Theorem in Equation (4.22). In the case of Algorithm II, this difference comes from our upper bound for $\prod_{i=n+2}^{m} \lambda_i(\mathcal{M}(a)_N)$ in Equation (4.29).

Remark 4.6.7. Using our theoretical bounds, we can derive proven estimates depending on ε, for the invariant Δ, defined in Equation 4.15.

4.7 Variations of the HLP

In this section, we consider variations of the hidden lattice problem from Definition 4.1.2. We first consider the problem with noise, as is customarily done in several applications. Next, we consider a decisional version.

4.7.1 HLP with noise

We can consider a noisy version of the HLP as follows.

Definition 4.7.1. *Let $\mu, \rho \in \mathbb{R}_{>0}$, integers $1 \leq r \leq n \leq m$ and N. Let $\mathcal{L} \subseteq \mathbb{Z}^m$ be a μ-small lattice of rank n and $\{w_j\}_{j=1,\ldots,r}$ be linearly independent vectors in $\mathbb{Z}^m$ such that there exist linearly independent vectors $\{x_j\}_{j=1\ldots,r}$ in $\mathbb{Z}^m$ satisfying $w_j - x_j \in \mathcal{L} \pmod N$ and $\|x_j\| \leq \rho$ for all j. The* Noisy Hidden Lattice Problem (NHLP) *is the task to compute from the knowledge of n, N and the vectors $\{w_j\}_j$, a basis of the completion of any lattice Λ satisfying the properties of $\mathcal{L}$.*

We solve the NHLP by reducing it in the first place to a HLP. Let $\mathcal{X}$ be the rank-r lattice generated by $\mathfrak{X} = \{x_j\}_j$; let $\mathfrak{B}$ be a μ-small basis of $\mathcal{L}$. We assume that $\mathcal{L} \cap \mathcal{X} = \{0\}$, so that $\mathcal{L} \oplus \mathcal{X}$ has rank $n+r$. Further, by assumption, $\mathcal{L} \oplus \mathcal{X}$ has size $\sigma(\mathfrak{B} \cup \mathfrak{X}) \leq \sigma(\mathfrak{B}) + \sigma(\mathfrak{X}) \leq \mu + \rho$ and contains $\{w_j\}_j$ modulo N. Therefore, the vectors $\{w_j\}_j$ are an instance of a hidden lattice problem with hidden lattice $\mathcal{L} \oplus \mathcal{X}$.

We first treat the special case when ρ is larger than μ. The application of either Algorithm I or Algorithm II to $\{w_j\}_j$, reveals a reduced basis of $\overline{\mathcal{L} \oplus \mathcal{X}}$ if the parameters are suitable. If ρ is larger than μ, then one can distinguish, in a reduced basis of $\overline{\mathcal{L} \oplus \mathcal{X}}$, the vectors belonging to $\overline{\mathcal{L}}$ from those belonging to $\overline{\mathcal{X}}$.

We now consider the general case and work without the assumption $\rho > \mu$, in which case, we do not expect a significant gap between vectors of $\overline{\mathcal{L}}$ and vectors of $\overline{\mathcal{X}}$ in a reduced basis of $\overline{\mathcal{L} \oplus \mathcal{X}}$, and consequently, cannot directly identify vectors in $\overline{\mathcal{L}'}$. We will overcome this identification problem via an embedding in larger dimension and the resolution of a system of linear equations. More precisely, let $\mathcal{L}' \in \mathbb{Z}^{m+r}$ be the lattice embedded in $\mathbb{Z}^{m+r}$ as $(\mathcal{L}, 0)$, that is, the set of vectors $(v, 0) \in \mathcal{L} \times \{0\}^r$. For $1 \leq j \leq r$, let $w'_j = (w_j, e_j) \in \mathbb{Z}^{m+r}$ and $x'_j = (x_j, e_j) \in \mathbb{Z}^{m+r}$, where $e_j \in \mathbb{Z}^r$ is the jth standard unit vector; let $\mathcal{M}' \subseteq \mathbb{Z}^{m+r}$ be the rank-r lattice generated by $\{w'_j\}_j$, and $\mathcal{X}' \subseteq \mathbb{Z}^{m+r}$ be the rank-r lattice generated by $\{x'_j\}_j$. Clearly, $\mathcal{M}' \subseteq \mathcal{L}' \oplus \mathcal{X}' \pmod N$, and $\mathcal{L}' \oplus \mathcal{X}'$ is a small hidden lattice of rank $n+r$ in dimension $m+r$. We proceed as follows to compute $\overline{\mathcal{L}}$. Let $\pi : \mathbb{Z}^{m+r} \to \mathbb{Z}^r$ be the projection onto the last r coordinates. We will distinguish between vectors in $\overline{\mathcal{L}}$ and $\overline{\mathcal{X}}$ by the fact that for every $v \in \mathcal{L}'$ it holds $\pi(v) = 0$, and $\mathcal{X}' \cap \{v \in \mathbb{Z}^{m+r} : \pi(v) = 0\} = \{0\}$. We therefore recover (basis) vectors v of $\overline{\mathcal{L}}$ from vectors in $\overline{\mathcal{L}'}$ by solving a system of linear equations, imposing that $\pi(v) = 0$.

Practically, let B be a reduced basis matrix of $\overline{\mathcal{L}' \oplus \mathcal{X}'}$, written as $B = [V|U] \in \mathbb{Z}^{(n+r)\times(m+r)}$, where $V \in \mathbb{Z}^{(n+r)\times m}$ and $U \in \mathbb{Z}^{(n+r)\times r}$, and computed by either Algorithm I or Algorithm II, on input $\mathcal{M}'$. By our analyses in Section 4.5, we expect to compute such B successfully under the heuristic conditions (4.8) and (4.12), with, essentially, n replaced by $n+r$, m replaced by $m+r$ and μ replaced by $\mu + \rho$. We next compute the left-kernel of U, that is, $K \in \mathbb{Z}^{n\times(n+r)}$ such that $KU = 0_{n,r}$. This implies $KB = [KV|0_{n,r}]$. Heuristically, the rows in KB must be

vectors in $\overline{\mathcal{L}'}$, as the last r components are zero. Therefore, the rows of KV are heuristically a basis for $\overline{\mathcal{L}}$. Namely, we heuristically expect to uniquely recover $\overline{\mathcal{L}'}$, as it is unlikely in the "generic" case, that there exists a small lattice $\Lambda \neq \mathcal{L}$ of rank n in $\mathbb{Z}^m$ such that $\Lambda \oplus \mathcal{X}$ contains $\mathcal{M}$ modulo N.

Algorithm 9 Solve the NHLP in general

Parameters: The HLP parameters n, m, r, μ, ρ, N from Definition 4.7.1
Input: A valid input for the NHLP
Output: A basis of the lattice $\overline{\mathcal{L}}$ (under suitable parameter choice)
(1) Run Algorithm I or Algorithm II on the lattice $\mathcal{M}' \subseteq \mathbb{Z}^{m+r}$ generated by $\{w'_j\}_j$; write the output basis vectors into the rows of a matrix $B \in \mathbb{Z}^{(n+r)\times(m+r)}$
(2) Write $B = [V|U]$ with $V \in \mathbb{Z}^{(n+r)\times m}$ and $U \in \mathbb{Z}^{(n+r)\times r}$. Compute $K \in \mathbb{Z}^{n\times(n+r)}$ such that $KU = 0_{n\times r}$
(3) Return the basis given by the rows of KV

4.7.2 Decisional HLP

We propose the following decisional version of the HLP.

Definition 4.7.2. *Let $\mu \in \mathbb{R}_{\geq 1}$, integers $1 \leq r \leq m$ and $N \in \mathbb{Z}$. Let $\mathcal{M} \subseteq \mathbb{Z}^m$ be a lattice of rank r. The* Decisional Hidden Lattice Problem (DHLP) *is the task to decide from the knowledge of μ, N and a basis of $\mathcal{M}$, whether there exists a μ-small lattice $\mathcal{L} \subseteq \mathbb{Z}^m$ of rank $1 \leq n \leq m$ such that $\mathcal{M} \subseteq \mathcal{L} \pmod N$.*

The rank of $\mathcal{L}$ is not given as input, and our algorithm is able to detect it. Note that when $\mathcal{L}$ exists, then there exist many small lattices of lower ranks (e.g. the sublattices of $\mathcal{L}$). Therefore, we would like $\mathcal{L}$ to be of maximal rank.

4.7.2.1 A geometric approach

To solve the DHLP for $\mathcal{M}$, we consider the successive minima of $\mathcal{M}^{\perp_N}$ (or $\mathcal{M}_N$) and show that the existence of a small lattice $\mathcal{L}$ impacts the geometry of $\mathcal{M}^{\perp_N}$ (or $\mathcal{M}_N$). In particular, we show that we can solve the DHLP via recognizing significant gaps in short bases of $\mathcal{M}^{\perp_N}$ (or $\mathcal{M}_N$). Lattices with gaps in their successive minima and their impact on known cryptosystems are for example studied in [LWXZ11].

Lemma 4.7.3. *For every lattice $\Lambda \subseteq \mathbb{Z}^m$ of rank m and every sublattice $\Lambda' \subseteq \Lambda$ of rank $0 < m' < m$, one has*

$$\frac{\prod_{k=m'+1}^{m} \lambda_k(\Lambda)}{\prod_{k=1}^{m'} \lambda_k(\Lambda)} \geq \gamma_{m'}^{-m'} \frac{\mathrm{Vol}(\Lambda)}{\mathrm{Vol}(\Lambda')^2} .$$

Proof. The quotient can be written as $\prod_{k=1}^{m} \lambda_k(\Lambda)/(\prod_{k=1}^{m'} \lambda_k(\Lambda))^2$. Since vectors $\{z_k\}_k$ in Λ with $\|z_k\| = \lambda_k(\Lambda)$ for every $1 \leq k \leq m$ form a sublattice of Λ of rank m, one has $\prod_{k=1}^{m} \lambda_k(\Lambda) \geq \mathrm{Vol}(\Lambda)$. The denominator is upper bounded by $(\prod_{k=1}^{m'} \lambda_k(\Lambda'))^2$ as Λ' is a sublattice of Λ. Finally, Minkowski's Second Theorem (Theorem 2.2.2) entails the claimed bound. □

Corollary 4.7.4. *Let $\mathcal{M} \subseteq \mathbb{Z}^m$ be a lattice of rank r and let $N > 0$ be an integer. Assume that $\mathrm{Vol}(\mathcal{M}^{\perp_N}) = N^r$. If there exists a μ-small lattice $\mathcal{L} \subseteq \mathbb{Z}^m$ of rank n such that $\mathcal{M} \subseteq \mathcal{L} \pmod N$ then*

$$\frac{\prod_{k=m-n+1}^{m} \lambda_k(\mathcal{M}^{\perp_N})}{\prod_{k=1}^{m-n} \lambda_k(\mathcal{M}^{\perp_N})} \geq \gamma_{m-n}^{-(m-n)} \frac{N^r}{\mu^{2n}} \tag{4.31}$$

Proof. Lemma 4.7.3 with $\Lambda = \mathcal{M}^{\perp_N}$ and $\Lambda' = \mathcal{L}^{\perp}$ of rank $m-n$ gives the lower bound $\gamma_{m-n}^{-(m-n)}\mathrm{Vol}(\mathcal{M}^{\perp_N})/\mathrm{Vol}(\mathcal{L}^{\perp})^2$ on the considered ratio. We conclude using $\mathrm{Vol}(\mathcal{M}^{\perp_N}) = N^r$, and $\mathrm{Vol}(\mathcal{L}^{\perp}) \leq \mathrm{Vol}(\mathcal{L}) \leq \mu^n$, by Lemma 4.3.2. □

We observe that this ratio grows as N gets larger. A similar lower bound can be obtained for a similar ratio of successive minima of $\mathcal{M}_N$ by either using Lemma 4.7.3 with $\Lambda' = \mathcal{N}_{\mathrm{II}} \subseteq \overline{\mathcal{L}}$ or invoking Banaszczyk's Theorem (Theorem 2.2.3). Since $\mathcal{N}_{\mathrm{II}}$ has rank n, the gap is visible between the nth minimum and the $(n+1)$th minimum.

Non-HLP instances. We compare with lattices $\mathcal{M}$ not lying in a μ-small lattice $\mathcal{L}$ modulo N (we call this a Non-HLP instance). Expectedly, this is the case for random lattices $\mathcal{M}$, when r basis vectors are uniformly chosen from $(\mathbb{Z}/N\mathbb{Z})^m$. We rely on the Gaussian Heuristic (4.5) to estimate the ratio of successive minima considered in Corollary 4.7.4. Assuming all the minima to be balanced, they are approximately equal to $\sqrt{m/(2\pi e)}N^{r/m}$, if $\mathrm{Vol}(\mathcal{M}^{\perp_N}) = N^r$ (as assumed in Corollary 4.7.4), which is likely for a random choice of $\mathcal{M}$. Therefore, for any $1 \leq n \leq m-1$:

$$\frac{\prod_{k=m-n+1}^{m} \lambda_k(\mathcal{M}^{\perp_N})}{\prod_{k=1}^{m-n} \lambda_k(\mathcal{M}^{\perp_N})} \gtrapprox \frac{(\sqrt{m/(2\pi e)}N^{\frac{r}{m}})^n}{(\sqrt{m/(2\pi e)}N^{\frac{r}{m}})^{m-n}} = \sqrt{\frac{m}{2\pi e}}^{2n-m} N^{\frac{r(2n-m)}{m}} \tag{4.32}$$

For $m = 2n$, this approximation is 1, and much larger if $2n > m$. In particular, we observe that (4.32) is in general much smaller than (4.31), as can be seen when choosing $m > 2n$ and relatively small values of μ.

4.7.2.2 Heuristic Algorithm for DHLP

Since we cannot compute the successive minima efficiently, the ratio in Corollary 4.7.4 is not practical. Instead, we approximate the minima by the norms of the vectors in an LLL-reduced basis. Using the proven bounds for LLL (Theorem 2.2.5), it is immediate to establish a similar lower bound for the ratio

$$g_{m-n}(\mathcal{M}^{\perp_N}) := \frac{\prod_{k=m-n+1}^{m} \|u_k\|}{\prod_{k=1}^{m-n} \|u_k\|}$$

where $\{u_k\}_k$ is an LLL-reduced basis of $\mathcal{M}^{\perp_N}$. Such a lower bound gives a necessary condition for the existence of a μ-small lattice $\mathcal{L}$ such that $\mathcal{M} \subseteq \mathcal{L} \pmod N$. Equation (4.31) shows an explicit dependence on n, the rank of $\mathcal{L}$. Since n is unknown, one first detects $m-n$ (the rank of $\mathcal{L}^{\perp}$) by computing the successive ratios $\{g_{m-\ell}(\mathcal{M}^{\perp_N})\}_\ell$ defined by $g_{m-\ell}(\mathcal{M}^{\perp_N}) = \prod_{k=m-\ell+1}^{m} \|u_k\| / \prod_{k=1}^{m-\ell} \|u_k\|$ for $\ell = 1, \ldots, m-1$ and $\{u_k\}_k$ a reduced basis of $\mathcal{M}^{\perp_N}$; one has $g_1(\mathcal{M}^{\perp_N}) \geq g_2(\mathcal{M}^{\perp_N}) \geq \ldots \geq g_{m-1}(\mathcal{M}^{\perp_N})$. One then identifies the smallest index $m - \ell_0$ such that $g_{m-\ell_0}(\mathcal{M}^{\perp_N})$ is significantly larger than $g_{m-\ell}(\mathcal{M}^{\perp_N})$ for all $\ell < \ell_0$. In that case, we expect the existence of a hidden small lattice of rank $n = \ell_0$. Again, this is easily adapted for Algorithm II when considering $\mathcal{M}_N$ instead of $\mathcal{M}^{\perp_N}$. Although this approach only

solves DHLP in one direction, we heuristically expect the converse to be true: if these gaps are sufficiently large, then there exists a small lattice $\mathcal{L}$ containing $\mathcal{M}$ modulo N.

4.8 Applications and Impacts on Cryptographic Problems

In this section we address applications of the hidden lattice problem in cryptography and discuss the impact of our algorithms. In the literature, these problems are typically solved by means of Algorithm I. Our Algorithm II provides a competitive alternative for solving these problems.

4.8.1 CRT-Approximate Common Divisor Problem

We first consider the CRT-ACD Problem from Definition 2.2.7 and the two-step algorithm by Coron and Pereira [CP19] which we have recalled in Section 2.2.3.2 (see Algorithm 1). Recall that [CP19] works when $\#\mathcal{S} = O(n)$. Following [CNW20a], we will in Chapter 5 describe an improved algorithm (see Algorithm 12) when $\#\mathcal{S} = O(\sqrt{n})$ only. This improvement is mainly obtained by modifying the second step of the algorithm from [CP19]. However, both these algorithms essentially agree on their first step, which is based, as we explain below, on solving a hidden lattice problem with $r = 1$. Therefore the presentation in the current section affects both algorithms (Algorithm 1 and Algorithm 12).

HLP and the CRT-ACD problem. As recalled in Section 2.2.3.2, the first step of the algorithm in [CP19] works as follows. Let $N = \prod_{i=1}^{n} p_i$ be the public integer which is to factor, and consider $\mathcal{S} = \{x_1, \ldots, x_n, y\}$ and $x = (x_1, \ldots, x_n) \in \mathcal{S}^n$, with $\#\mathcal{S} = n + 1$. Then, the vector $b = (x, y \cdot x) \in \mathbb{Z}^{2n}$ is public, and by the Chinese Remainder Theorem, letting $x \equiv x^{(i)} \pmod{p_i}$ and $y \equiv y^{(i)} \pmod{p_i}$ for all $1 \leq i \leq n$, one has

$$b \equiv \sum_{i=1}^{n} c_i (x^{(i)}, y^{(i)} x^{(i)}) =: \sum_{i=1}^{n} c_i b^{(i)} \pmod{N},$$

for some integers $c_1, \ldots, c_n$. If the vectors $\{x^{(i)}\}_i$ are $\mathbb{R}$-linearly independent, then so are the vectors $\{b^{(i)}\}_i$ and generate a $2n$-dimensional lattice $\mathcal{L}$ of rank n. Importantly, by Definition 2.2.7, $\{b^{(i)}\}_i$ are reasonably short vectors with entries bounded by $2^{2\rho}$, approximately. The basis $\{b^{(i)}\}_i$ of $\mathcal{L}$ has size $\mu := \sigma(\{b^{(i)}\}_i) = (n^{-1} \sum_{i=1}^{n} \|b^{(i)}\|^2)^{1/2} \lessapprox \sqrt{2n} \cdot 2^{2\rho}$, i.e. $\mu = O(n^{1/2} 2^{2\rho})$. Since the basis $\{b^{(i)}\}_i$ of $\mathcal{L}$ is secret, and $b \in \mathcal{L} \pmod{N}$, we view the vector b (or rather, the rank-one lattice $\mathcal{M} = \mathbb{Z}b$) as an instance of a HLP of rank $r = 1$, with hidden lattice $\mathcal{L}$ of size $O(n^{1/2} 2^{2\rho})$. Based on this observation, Algorithm 1 uses the orthogonal lattice attack to compute a basis of the completion of $\mathcal{L}$, and succeeds when η is sufficiently larger than ρ (i.e. when N is sufficiently larger than μ). The first step is to run lattice reduction on $\mathcal{M}^{\perp_N}$. Upon recovery of a basis of $\overline{\mathcal{L}}$, the authors proceed with an "algebraic attack" based on computing the eigenvalues of a well-chosen (public) matrix and then revealing the prime numbers $\{p_i\}_i$ by a gcd-computation, as discussed in Section 2.2.3.2.

4.8.2 The Hidden Subset Sum Problem

Our second application concerns the hidden subset sum problem, stated in Definition 2.2.6 following [NS99], and also considered in [CG20].

HLP and the Hidden Subset Sum Problem. Consider an instance of the hidden subset sum problem, that is, a public vector $v \in \mathbb{Z}^m$ and an integer N such that $v \equiv \sum_{i=1}^{n} \alpha_i x_i \pmod N$, where the vectors $\{x_i\}_i$ have binary entries. In [NS99], Nguyen and Stern propose an algorithm in two steps. The first step actually relies on solving a hidden lattice problem with $r = 1$: namely, the lattice

$$\mathcal{L} := \bigoplus_{i=1}^{n} \mathbb{Z} x_i$$

is μ-small with $\mu = \sigma(\{x_i\}_i) \leq \sqrt{m}$, and $v \in \mathcal{L} \pmod N$. Nguyen and Stern therefore compute in the first place a basis of $\overline{\mathcal{L}}$, by the orthogonal lattice algorithm (Algorithm I). The second step reveals $\{x_i\}_i$ and $\{\alpha_i\}_i$ from such a basis. Recently, Coron and Gini [CG20] argued that, due to the need of running BKZ with an increasingly large block-size to compute a short lattice basis, the second step of the algorithm in [NS99] has exponential (in n) complexity, and is therefore practical only in low dimension. Moreover, [CG20] gives an alternative second step, based on solving a system of multivariate equations, which works in polynomial-time, at the cost of having a larger dimension $m = O(n^2)$. The first step, i.e. the resolution of the associated HLP, remains unchanged and is solved by Algorithm I in [CG20]. In dimension $m = O(n^2)$, lattice reduction in the first step becomes unpractical. Therefore the authors employ a technique to compute $\overline{\mathcal{L}}$ when m is a lot larger than n (see [CG20, Section 4.1, Section 5]), based on computing a reduced basis for $(\mathbb{Z}v)^{\perp_N} \subseteq \mathbb{Z}^m$ by parallelizing lattice reduction over several components of v, of smaller dimension, say $2n$. The idea is to reduce the HLP given by $\mathcal{M} \subseteq \mathcal{L} \pmod N$ in dimension m, to multiple HLP's $\pi_j(\mathcal{M}) \subseteq \pi_j(\mathcal{L}) \pmod N$, where $\{\pi_j\}_j$ are projections $\mathbb{Z}^m \to \mathbb{Z}^{m'}$ onto block-components, where $m' < n$ (e.g. when $m = O(n^2)$, one can let $m' = O(n)$ and $O(n)$ projections $\{\pi_j\}_j$). One then solves the HLP's $\pi_j(\mathcal{M}) \subseteq \pi_j(\mathcal{L}) \pmod N$ for every j (this can be done in parallel) and reconstructs a full basis of $\overline{\mathcal{L}}$. We note that it is immediate to adapt this method to our Algorithm II.

In [NS99, CG20], the density for the hidden subset sum problem is defined as $n/\log(N)$, as analogy to the density for the classical subset sum problem [LO85] (we refer to Remark 4.5.1 for a deeper discussion on the density). By giving a heuristic lower bound for $\log(N)$, Coron and Gini justify that the orthogonal lattice attack solves the hidden lattice problem in the first step with density $O(1/n)$, when $m = 2n$.

Our Algorithm II can in turn be used to solve the HLP in the first step of the algorithms of [NS99, CG20]. Note that when $m = 2n$ (and $\mu = O(\sqrt{n})$), the density is heuristically at most $O(1/n)$ and proven $O(1/(n \log(\sqrt{n})))$ according to our Table 4.3. This gives a factor 2 improvement compared to [CG20].

4.8.3 More applications related to Cryptography

Cryptanalysis of CLT13 with independent slots

In Chapter 3, we have studied the security of CLT13 Multilinear Maps [CLT13] when instantiated with independent slots, and proposed an improved attack compared to [GLW14]. In fact, this cryptanalysis (see Chapter 3, Section 3.3.6) can be interpreted as an example of our NHLP from Definition 4.7.1. We refer to Chapter 3 for full details. The expression $w \equiv \sum_{i=1}^{\theta} \alpha_i \hat{m}_i + R \pmod{x_0}$ considered in Equation (3.14), where w is a zero-tested vector encoding, $\{\alpha_i\}_i$ are certain unknown integers, $\{\hat{m}_i\}_i \subseteq \mathbb{Z}^\ell$ are short vectors describing the

non-zero components of the plaintexts, and R is an unknown "noise" vector of small Euclidean norm, exactly translates into a hidden lattice problem with noise vector R and $r = 1$. Namely, $w - R \in \mathcal{L} := \bigoplus_{i=1}^{\theta} \mathbb{Z}\hat{m}_i \pmod{x_0}$, and the basis $\{\hat{m}_i\}_i$ of $\mathcal{L}$ consists of short vectors (compared to x_0), making $\mathcal{L}$ a small lattice. Similarly, Equation (3.22) corresponds to an instance of the NHLP with $r = d$. As noticed in Chapter 3, the integers $\{\alpha_i\}_i$ carry a particular structure, which makes our attack actually work more directly than what is presented in Section 4.7.1. Namely, by relying on the first step of Algorithm I for the NHLP (Algorithm 9), we have revealed a basis of the lattice $\Lambda \subseteq \mathbb{Z}^\ell$ of vectors orthogonal to $\{\hat{m}_i\}_i$ modulo the prime numbers $\{g_i\}_i$ defining the plaintext space. One may therefore rather view Λ as hidden a lattice, instead of $\mathcal{L}$. Upon the computation of a basis of Λ, our cryptanalysis proceeds by computing the (secret) volume $\prod_{i=1}^{\theta} g_i$ of Λ.

Fault attacks on RSA Signatures

In [CNT10], Coron et al. describe a cryptanalysis on a signature scheme based on RSA, based on an attack similar to [NS98]. In [BNNT11], a very similar attack is described against RSA-CRT signatures. Following the notation in [CNT10, Section 3], by considering ℓ faulty signatures together with a public modulus $N = pq$ (as in RSA), one derives an equation of the form $a_i + x_i + cy_i \equiv 0 \pmod{p}$ for $1 \leq i \leq \ell$, where $\{a_i\}_i$ are known integers, $\{x_i\}_i, \{y_i\}_i, c$ are unknown. Letting $a = (a_i)_i$, we derive that $a \in \mathcal{L} \pmod{p}$, where $\mathcal{L}$ is the rank-2 lattice $\mathbb{Z}x \oplus \mathbb{Z}y$ generated by $x = (x_i)_i$ and $y = (y_i)_i$ in $\mathbb{Z}^\ell$. If $\{x_i\}_i$ and $\{y_i\}_i$ are sufficiently small, then $\mathcal{L}$ is a suitably small lattice, describing a HLP of rank 2 in dimension ℓ. The authors follow Algorithm I to compute a basis $\{x', y'\}$ of $\overline{\mathcal{L}}$. Upon recovery of x', y', the attack proceeds by simple linear algebra and a gcd computation to reveal p.

In this case, the hidden lattice has rank only 2. Therefore Algorithm II is much more direct in the second step. While ℓ is not very large in [CNT10, BNNT11], we note that, in general, computing the completion of the rank-2 lattice $\mathcal{N}_{\mathrm{II}}$, is much faster than computing the orthogonal complement of the rank-$(\ell - 2)$ lattice $\mathcal{N}_{\mathrm{I}}$, as in [CNT10, BNNT11].

4.9 Practical aspects of our algorithms

We provide practical results for the HLP obtained in SageMath [S+20]. Our experiments are done on a standard 3.3 GHz Intel Core i7 processor. The source code is available in [NW21]. For $a \in \mathbb{Z}_{\geq 2}$, let $\mathfrak{p}(a)$ denote the smallest prime number larger than 2^a.

Instance generation. We generate random instances of the HLP and test Algorithms I and II (Algorithm 7 and Algorithm 8). Given fixed integers r, n, m, N as in Definition 4.1.2, we uniformly at random generate a basis $\mathfrak{B}$ for a hidden lattice $\mathcal{L}$, where the absolute values of the entries of each vector are bounded by some positive integer α, i.e. every vector has infinity norm at most α. We let $\mu := \sigma(\mathfrak{B})$, and by construction, $\mathcal{L}$ is μ-small. To generate a lattice $\mathcal{M} \subseteq \mathcal{L} \pmod{N}$ of rank r, we generate r random (uniformly) linear combinations modulo N of the basis vectors in $\mathfrak{B}$. For large n, m, μ, the lattice $\mathcal{L}$ is likely to be complete.

Running Times. In Table 4.6 we compare the running times for our algorithms. Here $N = \mathfrak{p}(a)$ where a is indicated in the column "$\log(N)$". For Algorithm I, "Step 1" runs LLL on $\mathcal{M}^{\perp_N}$

and computes $\mathcal{N}_{\mathrm{I}}$; "Step 2" computes $\mathcal{N}_{\mathrm{I}}^{\perp}$ following Section 4.4.4. For Algorithm II, "Step 1" runs LLL on $\mathcal{M}_N$ and computes the generating set of $\mathcal{N}_{\mathrm{II}}$, while "Step 2" computes $\overline{\mathcal{N}_{\mathrm{II}}}$. For the latter, we compute $\overline{\mathcal{N}_{\mathrm{II}}}^{N^{\infty}}$; namely, in these cases we have $(\overline{\mathcal{L}} : \mathcal{N}_{\mathrm{II}}) = N^{n-r}$. For this step, we compare the running time with Magma [BCP97], which seems to perform the finite field linear algebra much faster. The total running times for Algorithm II are therefore very competitive and constitute a major strength of Algorithm II against Algorithm I. Our observation from Section 4.4.4 is confirmed: the running time for Algorithm II is largely reduced for larger values of r (e.g. $\geq m/2$). In these cases, the running time for Algorithm II outperforms Algorithm I.

Modulus size for Algorithms I and II. We fix n, m and and bound $\|v\|_{\infty}$ for basis vectors v of $\mathcal{L}$ by some α. We observe that for both algorithms there is a critical value $N^{\dagger}$ such that in 100 random instances of the HLP, the algorithm always succeeds (i.e. always outputs a basis of $\overline{\mathcal{L}}$) for N a bit larger than $N^{\dagger}$, never succeeds for N a bit smaller than $N^{\dagger}$ and partially succeeds in between. We denote these values by $N_{\mathrm{I}}^{\dagger}$ and $N_{\mathrm{II}}^{\dagger}$ for Algorithms I and II, respectively, as in Corollary 4.6.3.

For example, for $n = 35$, $m = 52$ and $\alpha = 2^{10}$, we find that for Algorithm I with $r = 1$, we must have $N > N_1^{\dagger} = 2^{1260}$ and for $r = 2$ we must have $N > N_2^{\dagger} = 2^{632}$. Note that $N_2^{\dagger} \approx N_1^{\dagger}/2$, as predicted by Equation (4.8). Similarly, we run the extended Algorithm I for NHLP (Algorithm 9, Section 4.7.1) with $r = 1$ and a noise vector x such that $\|x\| \leq \rho = 2^{75}$. We observe $N > N_3^{\dagger} := 2^{1525} > N_1^{\dagger}$. We refer to Figure 4.1 and compare with Corollary 4.6.3.

In Table 4.7, we fix m, r and μ and find, for increasing values of n, the smallest value for $\log(N)$ such that a randomly generated HLP with parameters n, m, r, μ, N is solvable by our algorithms. We compare the practical values of N with the heuristic values derived in Section 4.5. The columns "theoretical" stand for the lower bounds in (4.8), resp. (4.12), according to the algorithm. Practically, we observe that the condition in Equation (4.9) is already satisfied for $\theta = 1$, thus we may neglect the last term in Equation (4.12), which becomes negative. We run LLL, so we set $\log(\iota) = 0.03$. We study two series (Series 1 and 2) according to m, r. Conjecturally, we see that the practical bound for $\log(N)$ is the same for both algorithms; this is to be expected from the duality relation (see Section 4.4.2). An interesting question is to find a theoretical optimal bound for $\log(N)$ fitting best with the practical behaviour of our algorithms.

Output quality of basis. In random generations, $\mathcal{L}$ is complete with high probability, and we compare μ to the size of the basis output by Algorithms I and II. We observe that our algorithms compute much smaller (LLL-reduced) bases of $\mathcal{L}$, and in fact sometimes recover the basis uniquely (up to sign). In particular, they sometimes solve the stronger version of Definition 4.1.2, that of computing a *μ-small* basis instead of *any*. Table 4.4 is obtained for $m = 2n = 4r$ for increasing values of m; note that in this case μ is approximately $N^{1/4}$, as predicted theoretically by Equation (4.14).

Decisional HLP. We test the decisional version of the HLP of Section 4.7.2. Table 4.5 shows different values for $g_{m-n}(\mathcal{M}^{\perp_N})$ if $\mathcal{M}$ lies in a μ-small lattice $\mathcal{L}$ modulo N (HLP instance), and if $\mathcal{M}$ is randomly sampled (random instance). For the latter, we compare with the heuristic bound (4.32) and report this value in the column "heuristic". We fix $n = 25$ and consider

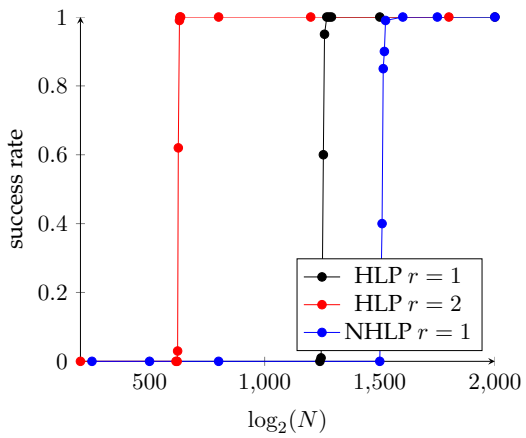

Figure 4.1: Success rate of Algorithm I as a function of N for HLP and NHLP

			Size of output basis	
m	$\log(N)$	$\log(\mu)$	Algorithm I	Algorithm II
76	175	42.335	42.335	42.335
160	325	77.874	79.927	79.814
320	80	13.36	18.132	18.182

Table 4.4: Sizes of output bases for Algorithms I and II

increasing values of $r < 35$.

Overall, we conclude that our practical experiments confirm the theoretical aspects studied in Section 4.5, 4.6, 4.7.

4.10 Appendix: Solving the HLP in large dimensions

Consider an instance of a hidden lattice problem $\mathcal{M} \subseteq \mathcal{L} \pmod N$, where the dimension m is large. The computation of $\mathcal{N}_{\mathrm{I}}$ and $\mathcal{N}_{\mathrm{II}}$ in the first step of our algorithms relies on lattice reduction, which in such large dimensions quickly becomes unpractical. Following the idea of [CG20, Section 4.1] in the context of the hidden subset sum problem, we describe a better

			HLP instance		random instance	
r	m	$\log(N)$	μ	$g_{m-n}(\mathcal{M}^{\perp_N})$	$g_{m-n}(\mathcal{M}^{\perp_N})$	heuristic
1	45	350	581.73	$3.52 \cdot 10^{21}$	$2.74 \cdot 10^{12}$	$5.73 \cdot 10^{12}$
5	70	100	580.11	$7.65 \cdot 10^{49}$	$3.39 \cdot 10^{-56}$	$7.4 \cdot 10^{-50}$
10	50	100	2037.15	$2.95 \cdot 10^{188}$	1.27	1
20	80	85	2949.66	$6.06 \cdot 10^{305}$	$3.16 \cdot 10^{-191}$	$4.09 \cdot 10^{-180}$

Table 4.5: Gaps in LLL-reduced bases of $\mathcal{M}^{\perp_N}$

				Running Time				
				Algorithm I		Algorithm II		
r	n	m	$\log(N)$	Step 1	Step 2	Step 1	Step 2 (Sage)	Step 2 (Magma)
60	150	200	140	7 min 13 s	1 min 20 s	10 min 2 s	3 min 4 s	0.37 s
110	150	200	90	6 min 20 s	1 min 29 s	4 min	1 min 33 s	0.24 s
175	180	200	140	6 min 56 s	1 min 24 s	1 min 39 s	20 s	0.19 s
80	100	300	75	3 min 51 s	30 min 17 s	22 min 51 s	30 s	0.12 s
150	200	300	75	145 min 29 s	22 min 23 s	116 min 14 s	6 min 19 s	0.56 s
75	150	400	80	75 min 16 s	326 min 44 s	414 min 51 s	5 min 13 s	0.61 s
235	275	400	80	527 min 43 s	117 min 2 s	304 min 10 s	15 min 39 s	0.95 s

Table 4.6: Running times for Algorithms I and II; the entries of a small basis of $\mathcal{L}$ are bounded by 2^{10}, which gives $\log(\mu) \approx 13$ in all instances

		$\log(N)$			
		Algorithm I		Algorithm II	
	n	heuristic	practice	heuristic	practice
Series 1 ($m = 100, r = 5$)	10	52	41	46	41
	20	113	92	103	92
	40	282	241	274	241
	80	1486	1384	1643	1384
	90	3240	3075	3695	3075
Series 2 ($m = 250, r = 30$)	50	57	48	54	48
	100	139	121	143	121
	160	327	296	381	296
	200	676	629	857	629

Table 4.7: Minimal values for $\log(N)$ as a function of the other parameters; the entries of a small basis of $\mathcal{L}$ lie in $(-2^{15}, 2^{15}) \cap \mathbb{Z}$, which gives $\log(\mu) \approx 18$ in all instances

algorithm for computing $\mathcal{N}_{\mathrm{I}}$ and $\mathcal{N}_{\mathrm{II}}$, based on parallel lattice reductions over several lower-dimensional lattices.

Assume for simplicity of exposition that $m = n^a$ with $a \geq 2$, i.e. m is at least quadratic in n. In such large dimensions, lattice reduction quickly becomes unpractical. Moreover, the size of N is required to be large, as can be seen from our heuristic analysis in Section 4.5. The method described below allows to run lattice reduction in dimension $O(n)$, in parallel on several bases; this is much faster in practice.

Let $b, \ell \in \mathbb{Z}_{\geq 1}$ such that $m = n^a = (b+1)\ell n$, that is, $b = n^{a-1}/\ell - 1$. Let $M \in \mathbb{Z}^{r \times m}$ be a basis matrix for $\mathcal{M}$, with basis vectors written in rows. We write $M = [M_0|M_1|\ldots|M_b]$ with $M_j \in \mathbb{Z}^{r \times \ell n}$ for $1 \leq j \leq b$. Then $\mathcal{M} \subseteq \mathcal{L} \pmod N$ gives $\pi_j(\mathcal{M}) \subseteq \pi_j(\mathcal{L}) \pmod N$ where $\pi_j : \mathbb{Z}^m \to \mathbb{Z}^{2\ell n}$ sends $(a_1, \ldots, a_m)$ to $(a_1, \ldots, a_{\ell n}, a_{j\ell n+1}, \ldots, a_{(j+1)\ell n})$, the first block-component of length ℓn of $(a_1, \ldots, a_m)$ followed by the jth block-component of length ℓn of $(a_1, \ldots, a_m)$. The latter is a hidden lattice problem in dimension $m' = 2\ell n$, *linear* in n. By running either Algorithm I or II in parallel on the basis matrices $[M_0|M_j] \in \mathbb{Z}^{r \times 2\ell n}$ of $\pi_j(\mathcal{M})$, we obtain (under a suitable choice of parameters) bases $V_j := [C_{0,j}|C_j] \in \mathbb{Z}^{n \times 2\ell n}$ of the hidden (completed) lattices $\overline{\pi_j(\mathcal{L})}$ for $1 \leq j \leq b$. By this dimension-lowering, those lattice reductions eventually work for an even smaller modulus N, as would originally be needed when performing lattice reduction directly with $m = n^a$. So to make the bases $\{V_j\}_j$ all start by $C_{0,1}$ in the first block, one base changes $\{C_{0,j}\}_j$ to $\{C_{0,1}\}$, assuming that $\{C_{0,j}\}_j$ have full rank n. By standard linear algebra, one computes matrices $\{Q_j\}$ such that $Q_j C_{0,j} = C_{0,1}$ for all $1 \leq j \leq b$. This gives bases $V_1 = [C_{0,1}|C_1]$ and $V_j' := Q_j V_j = [Q_j C_{0,j}|Q_j C_j] = [C_{0,1}|Q_j C_j]$. Finally, a basis of $\mathcal{L}$ is given by combining the projected block-components back into the full dimension m, as follows:

$$[C_{0,1} \mid C_1 \mid Q_2 C_2 \mid \ \ldots \ \mid Q_b C_b] \in \mathbb{Z}^{n \times m} \ .$$

This algorithm requires applying lattice reduction $b = n^{a-1}/\ell - 1$ times, which can be performed in parallel.

Remark 4.10.1. The exposition above is a little bit more general than [CG20], where the authors require a dimension roughly $m = n^2$ (i.e. $a = 2$) and fix $\ell = 1$, so that lattice reduction is performed in dimension $2n$. By the above, this is generalized to higher values of r (the rank of the public lattice $\mathcal{M}$), higher dimensions m ($m \geq n^2$), and controlled by a varying parameter ℓ. The utility of ℓ is to lower the value of b, the number of parallel computations. Namely, when the dimension m is much bigger, a larger value of ℓ compensates the growth of b. As mentioned, this technique straightforwardly also applies to Algorithm II. An implementation is provided in [NW21].

CHAPTER 5

Simultaneous Diagonalization of Incomplete Matrices

In this chapter, we consider the problem of recovering the entries of diagonal matrices $\{U_a\}_a$ for $a = 1, \ldots, t$ from multiple "incomplete" samples $\{W_a\}_a$ of the form

$$W_a = PU_aQ \ ,$$

where P and Q are unknown matrices of low rank. We devise practical algorithms and explicit parameters for this problem depending on the ranks of P and Q and t, the number of input samples. We justify that this problem finds its motivation in cryptanalysis: we show how to significantly improve previous algorithms for solving the approximate common divisor problem based on the Chinese Remainder Theorem and the Cheon et al. attack against the CLT13 cryptographic multilinear map scheme. Our improvement lies in reducing the input size while maintaining the same output size.

The content of this chapter is based on joint work [CNW20a] with Jean-Sébastien Coron and Gabor Wiese, which has been published in the proceedings of Fourteenth Algorithmic Number Theory Symposium 2020. We closely follow the exposition of [CNW20a].

5.1 Introduction

In this chapter, we consider the following computational problem from linear algebra.

Definition 5.1.1 (Problems $\mathbb{A}, \mathbb{B}, \mathbb{C}, \mathbb{D}$)**.** *Let $n \geq 2, t \geq 2$ and $2 \leq p, q \leq n$ be integers. Let $\{U_a : 1 \leq a \leq t\}$ be diagonal matrices in $\mathbb{Q}^{n \times n}$. Let $\{W_a : 1 \leq a \leq t\}$ be matrices in $\mathbb{Q}^{p \times q}$ and $W_0 \in \mathbb{Q}^{p \times q}$ such that W_0 has full rank and there exist matrices $P \in \mathbb{Q}^{p \times n}$ of full rank p and $Q \in \mathbb{Q}^{n \times q}$ of full rank q, such that $W_0 = P \cdot Q$ and $W_a = P \cdot U_a \cdot Q$ for $1 \leq a \leq t$. We distinguish the following cases:*

$$\begin{array}{ll} (\mathbb{A}) \quad p = n \text{ and } q = n & (\mathbb{B}) \quad p = n \text{ and } q < n \\ (\mathbb{C}) \quad p < n \text{ and } q = n & (\mathbb{D}) \quad p < n \text{ and } q = p \end{array}$$

In each of the four cases, the problem states as follows:

(1) Given the matrices $\{W_a : 0 \leq a \leq t\}$, compute $\{(u_{1,i}, \ldots, u_{t,i}) : 1 \leq i \leq n\}$, where for $1 \leq a \leq t$, $u_{a,1}, \ldots, u_{a,n} \in \mathbb{Q}$ are the diagonal entries of matrices $\{U_a : 1 \leq a \leq t\}$ as above.

(2) Determine whether the solution is unique.

Motivation from Cryptanalysis. As already mentioned in Section 1.2.3.1, Definition 5.1.1 mainly finds it roots in cryptanalysis. For example, this problem arises naturally when trying to lower the number of encodings in the attack of Cheon et al. [CHL+15] against the CLT13 Scheme [CLT13]. In fact, we argue that the attack by Cheon et al. corresponds to solving Problem $\mathbb{A}$. Lowering the number of encodings then gives rise to Problem $\mathbb{C}$ or $\mathbb{D}$, for which we shall design algorithms.

5.2 Our Contributions

We describe efficient algorithms for Problems $\mathbb{C}$ and $\mathbb{D}$ of Definition 5.1.1, and show how to minimize the parameters p and t with respect to n. We further rely on our algorithms to improve on two concrete applications with interest in cryptography. We believe that our algorithms are of independent interest and hope that more applications are to be found. More concretely, our contributions can be summarized as follows.

Algorithms for Problems $\mathbb{C}$ and $\mathbb{D}$. Our first contribution is to propose practical algorithms for both Problems $\mathbb{C}$ and $\mathbb{D}$. Our approach to solve Problem $\mathbb{C}$ is to use the invertibility of the matrix Q and write

$$W_a = PU_aQ = PQQ^{-1}U_aQ = W_0Z_a$$

with $Z_a = Q^{-1}U_aQ$, for every $1 \leq a \leq t$. As W_0 is not invertible, we cannot recover Z_a directly. However we interpret this as a system of linear equations to solve for the matrices $\{Z_a\}_a$. This system is, in general, underdetermined and does not yield the matrices $\{Z_a\}_a$ uniquely. However, exploiting the special feature that $\{Z_a\}_a$ commute among each other leads to additional linear equations. This enables to recover $\{Z_a\}_a$ uniquely, and simultaneous diagonalization eventually yields the diagonal entries of $\{U_a\}_a$. We then determine exact bounds on the parameters to ensure that the system have at least as many linear equations as variables, and we obtain that p and t can be set as $O(\sqrt{n})$. Our algorithm is heuristic only, but performs well in practice.

We next provide an algorithm for Problem $\mathbb{D}$, where we can clearly no longer invert the matrix Q. In rather imprecise words, we therefore reduce Problem $\mathbb{D}$ to Problem $\mathbb{C}$ by "augmenting" the matrix Q with extra columns so that it becomes invertible, and preserving the solution to Problem $\mathbb{D}$. Upon making this augmentation process, we can then use our algorithm for Problem $\mathbb{C}$ to solve Problem $\mathbb{D}$, and we show that p can be close to $2n/3$. We refer to Section 5.4 and Section 5.5 for a complete description of our algorithms and provide the results of practical experiments in Section 5.7.

Improved algorithm for the CRT-approximate common divisor problem. Our second contribution consists in improving the two-step algorithm by [CP19] for the multiprime approximate common divisor problem based on Chinese Remainders, introduced in Definition 2.2.7. Namely, we remark that [CP19] relies on solving a certain instance of Problem $\mathbb{A}$. By solving an appropriate instance of Problem $\mathbb{C}$ instead, we obtain a quadratic improvement in the number of input samples. More precisely, letting n be the number of secret primes in the public modulus N, we can factor N completely given only $O(\sqrt{n})$ input samples, whereas [CP19] uses $O(n)$ samples. We therefore achieve complete factorization of the public modulus N

while limiting the input size drastically. We confirm our results with practical experiments in Section 5.7.

Improved cryptanalysis of CLT13 Multilinear Maps. Our third contribution in this chapter is an improvement of the cryptanalysis of Cheon et al. [CHL+15] against the CLT13 multilinear map, when fewer encodings are available to the attacker. Namely, the attack of Cheon et al., recalled in Section 3.2.2.2, relies on solving some instance of Problem $\mathbb{A}$. By solving instances of Problems $\mathbb{C}$ or $\mathbb{D}$ instead, we can lower the number of public encodings required for the cryptanalysis. Specifically, for a composite modulus x_0 of n prime numbers, we obtain improved algorithms using only $O(\sqrt{n})$ encodings of zero, compared to n in [CHL+15]. Further, we also limit the number of total encodings to roughly $4n/3$, compared to $2n + 2$ required in [CHL+15]. We confirm our results with practical experiments in Section 5.7.

5.3 Preliminary Remarks about Problems $\mathbb{A}, \mathbb{B}, \mathbb{C}, \mathbb{D}$

We shall first make some important considerations about Definition 5.1.1.

(i) Let $\{W_a\}_a$ be as in Definition 5.1.1, $\pi \in \mathfrak{S}_n$ be a permutation with associated matrix $A_\pi \in \{0,1\}^{n\times n}$ and D any invertible diagonal $n \times n$ matrix. Then $P' = PDA_\pi$ and $Q' = A_\pi^{-1}D^{-1}Q$ satisfy $W_0 = P'Q'$ and $W_a = P'U'_aQ'$ for all $1 \le a \le t$, where $U'_a = A_\pi^{-1}U_aA_\pi$ is obtained from U_a by permuting its diagonal entries via π. Thus, $P', \{U'_a\}_a$ and Q' satisfy the same problem. For this reason, we only ask to recover the set $\{(u_{1,i}, \ldots, u_{t,i}) : 1 \le i \le n\}$ in Definition 5.1.1.

(ii) If $t = 1$ in Problem $\mathbb{C}$, then the problem is not solvable because its solution is not unique. Namely, we write $W_1 = W_0Z_1$, where $Z_1 = Q^{-1}U_1Q$ is diagonalizable with eigenvalues the diagonal entries of U_1. But also, for every $v \in \ker(W_0)$ one has

$$W_1 = W_0(Z_1 + vw_1^T)$$

for some $w_1 \in \mathbb{Q}^n$. Now, Z_1 and $Z_1 + vw_1^T$ likely have different eigenvalues which means that the solution is not unique.

(iii) There are cases when the problem is clearly not solvable for $p < n$. For example, if $P = [1_p|0_{p\times(n-p)}]$ then for all a the matrix PU_a only involves the first p diagonal entries of U_a and the information on the remaining $n - p$ is lost. These cases will not occur for "generic" or "random" instances of the problem.

(iv) We call a matrix $W_0 = PQ$ as in Definition 5.1.1 a "special input". If such a matrix is not available as input, but instead, only matrices $\{W_a\}_a$ of the form PU_aQ are available, then one can recover ratios of diagonal entries of $\{U_a\}_a$, if $t \ge 3$. Namely, defining $P' = PU_1$ and assuming that U_1 is invertible, one obtains that $W'_0 := P'Q = W_1$ and for $2 \le a \le t$, $W'_a := P'(U_aU_1^{-1})Q = W_a$. Running the algorithm on input $\{W'_a : 0 \le a \le t-1\}$ therefore reveals the tuples of diagonal entries of the matrices $U_aU_1^{-1}$ for $1 \le a \le t-1$. We will use this approach in Section 5.6.2.1 to improve the cryptanalysis of the CLT13 multilinear map.

(v) For simplicity, we have stated Definition 5.1.1 over $\mathbb{Q}$. More generally, we can consider matrices over any field K with exact linear algebra (e.g. solving linear systems, diagonalizing matrices, etc.). Our algorithms in the forthcoming sections of this chapter still apply to that case.

5.4 An Algorithm for Problem $\mathbb{C}$

We describe an algorithm to solve Problem $\mathbb{C}$ of Definition 5.1.1. The main idea is to reduce the problem to a certain auxiliary problem for commuting matrices. Finally, based on our algorithm, we propose minimal parameters for the rank p of P, and the number t of input matrices. These parameters will turn useful for important improvements for our applications Section 5.6.

5.4.1 Description of our algorithm

Consider integers $n, t \geq 2$ and $2 \leq p < n$ and an instance of Problem $\mathbb{C}$. We aim at computing the set $\{(u_{1,i}, \ldots, u_{t,i}) : 1 \leq i \leq n\}$ corresponding to diagonal matrices $\{U_a\}_a$. We remark that it is enough to solve the following auxiliary problem, that we shall refer to Problem $\mathbb{C}'$. Proposition 5.4.2 below shows that it is sufficient to solve Problem $\mathbb{C}'$ in order to solve Problem $\mathbb{C}$.

Definition 5.4.1 (Problem $\mathbb{C}'$)**.** *Let integers $n, t \geq 2$ and $2 \leq p < n$. Given*

- *a matrix $V \in \mathbb{Q}^{p \times n}$ of rank p and a basis matrix $E \in \mathbb{Q}^{n \times (n-p)}$ of $\ker(V)$,*
- *a set of matrices $\{Y_a : 1 \leq a \leq t\} \subseteq \mathbb{Q}^{n \times n}$*

compute matrices $\{X_a : 1 \leq a \leq t\} \subseteq \mathbb{Q}^{(n-p) \times n}$, such that the matrices $Y_a + EX_a$ for $1 \leq a \leq t$ commute with each other.

Proposition 5.4.2. *Let $\{W_a : 0 \leq a \leq t\}$ as in Problem $\mathbb{C}$. Let $E \in \mathbb{Q}^{n \times (n-p)}$ be a basis matrix of the kernel of W_0. Let W_0^+ be a right-inverse[1] of W_0. Define $V = W_0$ and $Y_a = W_0^+ W_a$ for $1 \leq a \leq t$. Assume that Problem $\mathbb{C}'$ is uniquely solvable for the input matrices V, E and $\{Y_a : 1 \leq a \leq t\}$.*

Then Problem $\mathbb{C}$ is uniquely solvable for the input matrices $\{W_a : 0 \leq a \leq t\}$. Moreover, the matrix Q in the assumption of Problem $\mathbb{C}$ is unique up to multiplication by a permutation matrix and an invertible diagonal matrix if at least one of the matrices $\{U_a\}_a$ has pairwise distinct diagonal entries.

Proof. Write $W_0 = PQ$ and $W_a = PU_aQ$ as in Problem $\mathbb{C}$. For every $1 \leq a \leq t$, we write $W_a = (PQ)(Q^{-1}U_aQ) = W_0Z_a$, where $Z_a := Q^{-1}U_aQ$. The matrices $\{Z_a : 1 \leq a \leq t\}$ commute and are simultaneously diagonalizable. For every $1 \leq a \leq t$, the matrix Z_a can be written as $Z_a = Y_a + EX_a$ for some matrix $X_a \in \mathbb{Q}^{(n-p) \times n}$ since $W_0Y_a = W_a$. Since the matrices $\{Z_a\}_a$ commute, V, E and $\{Y_a\}_a$ define a valid input for Problem $\mathbb{C}'$. By assumption, we can compute the matrices $\{X_a\}_a$ by solving Problem $\mathbb{C}'$ and these are unique. From the knowledge of the matrices $\{X_a\}_a$, we compute $Z_a = Y_a + EX_a$ for $1 \leq a \leq t$. Then these

[1] If W_0 (of full rank p) is defined over the complex numbers, one can take $W_0^+ = W_0^*(W_0W_0^*)^{-1}$ where W_0^* is the conjugate transpose of W_0, and $W_0^* = W_0^T$ over the real numbers.

matrices are also unique. Thus the set of tuples of eigenvalues $\{(u_{1,i}, \ldots, u_{t,i}) : 1 \leq i \leq n\}$ is unique and can be computed by simultaneous diagonalization.

For the last part of the statement, assume that we have matrices P', Q', diagonal matrices $\{U'_a\}_a$, which are necessarily of the form $U'_a = A^{-1}U_aA$ for a permutation matrix A, such that $W_0 = P'Q'$ and $W'_a = P'U'_aQ'$ for every a. By uniqueness of the matrices $\{Z_a\}_a$, we must have

$$Z_a = Q^{-1}U_aQ = Q'^{-1}U'_aQ' = Q'^{-1}A^{-1}U_aAQ' \quad , \quad 1 \leq a \leq t$$

or, equivalently $U_a(QQ'^{-1}A^{-1}) = (QQ'^{-1}A^{-1})U_a$ for $1 \leq a \leq t$. Consequently, the matrix $D := QQ'^{-1}A^{-1}$ commutes with the matrices $\{U_a\}_a$ and so is diagonal itself, as one of $\{U_a\}_a$ has pairwise distinct entries. This gives $Q = DAQ'$ and terminates the proof. □

Solving Problem $\mathbb{C}'$. Let us now explain how we compute a solution to Problem $\mathbb{C}'$ of Definition 5.4.1. Consider matrices $V, E, \{Y_a\}_a$ as in Problem $\mathbb{C}'$. We want to compute matrices $\{X_a\}_a$ such that the matrices $Z_a = Y_a + EX_a$ commute for all $1 \leq a \leq t$, that is, the Jacobi bracket $[Z_a, Z_b] = Z_aZ_b - Z_bZ_a$ is the zero matrix for all $a < b$. Using $Z_a = Y_a + EX_a$, this is equivalent to

$$0 = Y_aY_b - Y_bY_a + E \cdot S_{ab} + Y_aEX_b - Y_bEX_a \,, \tag{5.1}$$

where $S_{ab} := X_aY_b + X_aEX_b - X_bY_a - X_bEX_a$. Left multiplication by V and the identity $VE = 0$ then imply $VY_aY_b - VY_bY_a + VY_aEX_b - VY_bEX_a = 0$, which is equivalent to

$$\Delta_{ab} = VY_bEX_a - VY_aEX_b \quad , \quad 1 \leq a < b \leq t \,, \tag{5.2}$$

where $\Delta_{ab} := VY_aY_b - VY_bY_a$ is completely explicit in terms of the input matrices. Equation (5.2) describes a system of linear equations over $\mathbb{Q}$ in the variables given by the entries of the matrices X_a and X_b. Since Δ_{ab} has size $p \times n$, this leads to a system of np linear equations in the $2(n-p)n$ variables given by the entries of X_a and X_b. Writing Equation (5.2) for every pair (a, b) with $1 \leq a < b \leq t$, that is, writing out every commutativity relation between pairs of matrices Z_a and Z_b), we obtain a system of $t(t-1)/2np$ linear equations and $t(n-p)n$ variables given by the entries of the matrices $\{X_a : 1 \leq a \leq t\}$. From what precedes and by Proposition 5.4.2, we hence deduce the following result.

Proposition 5.4.3. *A unique solution to Problem $\mathbb{C}$ is implied by the existence of a unique solution to the explicit system of linear equations given in Equation* (5.2)*, which is a system of $\frac{1}{2}t(t-1)np$ linear equations in $t(n-p)n$ variables. There are at least as many equations as variables as soon as*

$$\frac{p}{n} \geq \frac{2}{t+1} \,. \tag{5.3}$$

Since there is no obvious linear dependence in the equations of the system, we heuristically expect, in the generic case, to find a unique solution $\{X_a : 1 \leq a \leq t\}$ under Condition (5.3). This solves Problem $\mathbb{C}'$, and therefore Problem $\mathbb{C}$.

Algorithm for Problem $\mathbb{C}$. We summarize our algorithm for Problem $\mathbb{C}$ as follows and will refer to it as Algorithm $\mathcal{A}_{\mathbb{C}}$ for the rest of this chapter.

Algorithm 10 Algorithm for Problem $\mathbb{C}$ (Algorithm $\mathcal{A}_{\mathbb{C}}$)

Parameters: Integers $2 \leq p \leq n$ and $t \geq 2$
Input: A valid input for Problem $\mathbb{C}$ as in Definition 5.1.1
Output: "Success" or "Fail"; in case of "Success", also output a solution. "Success" means uniqueness of the solution; "Fail" means that no solution was found.

Compute a basis matrix E of $\ker(W_0)$
Define $W_0^+ = W_0^T(W_0W_0^T)^{-1}$
for (a, b) with $1 \leq a < b \leq t$ **do**
    Compute the matrix $\Delta_{ab} := W_aW_0^+W_b - W_bW_0^+W_a$
end for
Solve the system of linear equations described in Equation (5.2)
if the solution to the system of equations is not unique **then**
    Return "Fail" and break
else
    Denote by $\{X_a : 1 \leq a \leq t\}$ the unique solution
end if
for $1 \leq a \leq t$ **do**
    Compute $Z_a := W_0^+W_a + EX_a$
end for
Perform simultaneous diagonalization of $\{Z_a\}_a$
Return "Success" with the tuples of eigenvalues of the matrices $\{Z_a\}_a$

5.4.2 Optimization of the parameters

We shall now find minimal possible (with respect to n) values for the parameters t and p. In our applications which we describe in Section 5.6 we are led to minimizing $p+t$ as a function of n. Following Proposition 5.4.3, we set

$$F_n(t) := p_n(t) + t := \frac{2n}{t+1} + t\ , \quad t \in \mathbb{R}_{>0}\ ,\ n \in \mathbb{Z}_{\geq 2}\ .$$

By elementary calculations, we see that F_n admits a minimum at $t_0 = \sqrt{2n} - 1$, which gives

$$p = p_n(t_0) = \sqrt{2n}\ .$$

This shows that minimal values for p and t are $O(\sqrt{n})$. This claim is in line with our practical observations in Section 5.7.

5.5 An Algorithm for Problem $\mathbb{D}$

We now describe an algorithm to solve Problem $\mathbb{D}$ of Definition 5.1.1. Here, the main idea is to reduce the problem to an instance of Problem $\mathbb{C}$ and solve the latter by Algorithm $\mathcal{A}_{\mathbb{C}}$. We again conclude this section by proposing minimal parameters for the rank p of P and Q, and the number t of input matrices. These parameters will turn useful for important improvements for our applications Section 5.6.

5.5.1 Description of our algorithm

Consider integers $n, t \geq 2$ and $2 \leq p < n$ and an instance of Problem $\mathbb{D}$. The main idea of our algorithm is a reduction to Problem $\mathbb{C}$ which can be solved using Algorithm $\mathcal{A}_{\mathbb{C}}$ (Algorithm 10). To do so, we exhibit matrices which are certain augmentations of the input matrices $\{W_a\}_a$ and admit the same solution as the input matrices. More precisely, following the notation of Definition 5.1.1, we will construct public matrices $W_0' = PQ'$ and $W_a' = PU_aQ'$ for $1 \leq a \leq t$, for the same diagonal matrices $\{U_a\}_a$ and for some $n \times n$ invertible matrix Q'.

Reducing Problem $\mathbb{D}$ to Problem $\mathbb{C}$. For every $1 \leq a, b \leq t$, we define the matrices

$$\Delta_{ab} = W_a W_0^{-1} W_b - W_b W_0^{-1} W_a \,. \tag{5.4}$$

Note that $\Delta_{ab} = -\Delta_{ba}$. We have the following lemma.

Lemma 5.5.1. *Let $W_0 = PQ$ and $W_a = PU_aQ$ for $1 \leq a \leq t$ as in Problem $\mathbb{D}$. Define the matrix $B = QW_0^{-1}P - 1_n \in \mathbb{Q}^{n\times n}$ and let r denote its rank. Then the following hold:*

(i) the matrix B has rank $r = n - p$

(ii) there exist matrices $V_a \in \mathbb{Q}^{p\times r}$ and $G_a \in \mathbb{Q}^{r\times p}$ for $1 \leq a \leq t$ such that for all $1 \leq a < b \leq t$, one has $\Delta_{ab} = V_aG_b - V_bG_a$.

Proof. (i) Let $C = QW_0^{-1}P$. Then $CQ = Q$ and the column-image of Q is contained in the eigenspace, say $\mathcal{E}$, of C for the eigenvalue 1. So, $\mathcal{E}$ has dimension at least p. However, the rank of C is bounded above by the rank of Q, i.e. by p. Finally, $\mathcal{E}$ has dimension exactly p and the rank r of $B = C - 1_n$ equals $n - p$.

(ii) For every $1 \leq a, b \leq t$, we can write

$$\begin{aligned} \Delta_{ab} &= PU_a(QW_0^{-1}P - 1_n)U_bQ - PU_b(QW_0^{-1}P - 1_n)U_aQ \\ &= PU_aBU_bQ - PU_bBU_aQ \end{aligned} \tag{5.5}$$

since U_a and U_b commute. Since B has rank r, there exist matrices $B_1 \in \mathbb{Q}^{n\times r}, B_2 \in \mathbb{Q}^{r\times n}$ with $B = B_1B_2$. Setting $V_a = PU_aB_1$ and $G_a = B_2U_aQ$ gives the claim. □

The following properties of the matrix B defined in Lemma 5.5.1 are key points for our reduction of Problem $\mathbb{D}$ to Problem $\mathbb{C}$.

Lemma 5.5.2. *Let $W_0 = PQ$ and $W_a = PU_aQ$ for $1 \leq a \leq t$ as in Problem $\mathbb{D}$. Let $B \in \mathbb{Q}^{n\times n}$ be the matrix of Lemma 5.5.1 with respect to P and Q and let $r = n - p$. Let $B_1 \in \mathbb{Q}^{n\times r}$ and $B_2 \in \mathbb{Q}^{r\times n}$ be such that $B = B_1B_2$. Then the following hold:*

(i) $PB_1 = 0_{p\times r}$

(ii) The matrix $Q' := [Q|B_1]$ is an $n \times n$ invertible matrix.

Proof. (i) The matrix B_2 defines a surjection $B_2 : \mathbb{Q}^n \to \mathbb{Q}^r$. Therefore, for every $x \in \mathbb{Q}^r$, we can write $x = B_2y$ for some $y \in \mathbb{Q}^n$ and obtain $PB_1x = PB_1(B_2y) = (PB)y = 0$.

(ii) Since $r = n - p$, the matrix Q' clearly has size $n \times n$. To show its invertibility, we establish that $\mathrm{im}(Q) \cap \mathrm{im}(B_1) = \{0\}$. Since B_2 is surjective, the images of B_1 and $B_1B_2 = B$ coincide. Let $Qx = By \in \mathrm{im}(Q) \cap \mathrm{im}(B_1)$, with $x \in \mathbb{Q}^p$ and $y \in \mathbb{Q}^n$. This gives $Qx = (QW_0^{-1}P - 1_n)y = QW_0^{-1}Py - y$. Thus $y = QW_0^{-1}Py - Qx = Qz$ with $z = W_0^{-1}Py - x$. Therefore, $Qx = By = B(Qz) = 0$ because $BQ = 0$. □

Remark 5.5.3. The matrix Q' depends on the matrix B once we have fixed a rank factorization $B = B_1B_2$ for B. Such matrices B_1 and B_2 are unique up to a matrix in $R \in \mathrm{GL}(r, \mathbb{Q})$: setting $B_1' = B_1R$ and $B_2' = R^{-1}B_2$ gives another factorization.

We now show that finding matrices $\{V_a\}_a$ such that there exist $\{G_a\}_a$ satisfying

$$\Delta_{ab} = V_aG_b - V_bG_a$$

for every a, b is sufficient to solve Problem $\mathbb{D}$. We view these matrices as being complementary to $\{W_a\}_a$ because they define themselves an instance of Problem $\mathbb{D}$ with the same solution as $\{W_a\}_a$, as seen in the proof of Lemma 5.5.1. This allows us to increase the rank of the matrix Q. We thus now formulate Problem $\mathbb{D}'$. Proposition 5.5.5 below bridges Problem $\mathbb{D}$ and Problem $\mathbb{C}$.

Definition 5.5.4 (Problem $\mathbb{D}'$)**.** *Let $n, t \geq 2$ and $2 \leq p < n$ be integers. For every $1 \leq a, b \leq t$, let $\Delta_{ab} \in \mathbb{Q}^{p\times p}$ be such that $\Delta_{ab} = V_aG_b - V_bG_a$ for $V_a \in \mathbb{Q}^{p\times(n-p)}$ of rank $n-p$ and $G_a \in \mathbb{Q}^{(n-p)\times p}$.*

The problem states as follows: Given the matrices Δ_{ab} for all $1 \leq a, b \leq t$, compute such matrices $\{V_a : 1 \leq a \leq t\}$.

Proposition 5.5.5. *Let $W_0 = PQ$ and $W_a = PU_aQ$ for $1 \leq a \leq t$ be as in Problem $\mathbb{D}$. For $1 \leq a, b \leq t$, let Δ_{ab} be the matrix defined in Equation* (5.4)*. Moreover, assume that*

(i) Problem $\mathbb{D}'$ is uniquely solvable for the input matrices $\{\Delta_{ab} : 1 \leq a < b \leq t\}$ and denote by $\{V_a : 1 \leq a \leq t\}$ the unique solution.

(ii) Problem $\mathbb{C}$ is uniquely solvable for the input matrices $W_0' = [W_0|0_{p\times(n-p)}] \in \mathbb{Q}^{p\times n}$ and $W_a' = [W_a|V_a] \in \mathbb{Q}^{p\times n}$ for $1 \leq a \leq t$.

Then Problem $\mathbb{D}$ is uniquely solvable on input $\{W_a : 0 \leq a \leq t\}$ and the unique solution is given by the unique solution to Problem $\mathbb{C}$ on input $\{W_a' : 0 \leq a \leq t\}$.

Proof. By Lemma 5.5.1 there exist matrices $V_a \in \mathbb{Q}^{p\times r}$ and $G_a \in \mathbb{Q}^{r\times p}$ for $1 \leq a \leq t$ such that $\Delta_{ab} = V_aG_b - V_bG_a$ for all $1 \leq a < b \leq t$. Therefore the matrices $\{\Delta_{ab}\}_{a,b}$ define an instance of Problem $\mathbb{D}'$. By assumption *(i)*, we compute the unique solution $\{V_a\}_a$ for this problem.

Now, let $W_0' = [W_0|0_{p\times(n-p)}] \in \mathbb{Q}^{p\times n}$ and $W_a' = [W_a|V_a] \in \mathbb{Q}^{p\times n}$ for $1 \leq a \leq t$. Let $B = QW_0^{-1}P - 1_n$ as in Lemma 5.5.1 of rank $r = n - p$. Let $B_1 \in \mathbb{Q}^{n\times r}$ and $B_2 \in \mathbb{Q}^{r\times n}$ be a rank factorization of B; i.e. $B = B_1B_2$. Letting $Q' := [Q|B_1] \in \mathbb{Q}^{n\times n}$, we have $PQ' = P[Q|B_1] = [W_0|0_{p\times r}] = W_0'$ and, by uniqueness of $\{V_a\}_a$ (see proof of Lemma 5.5.1), we have

$$PU_aQ' = PU_a[Q|B_1] = [W_a|V_a] = W_a'$$

for $1 \leq a \leq t$, as $PB_1 = 0_{n\times r}$ by Lemma 5.5.2 *(i)*. The matrix Q' is invertible by Lemma 5.5.2 *(ii)*. Therefore, W_0' and $\{W_a'\}_a$ define a valid input for Problem $\mathbb{C}$. By assumption *(ii)*, this problem is uniquely solvable and the solution must be the tuples of diagonal entries of the matrices $\{U_a\}_a$. This is also a solution to Problem $\mathbb{D}$ since the matrices $\{U_a\}_a$ are the same for the input matrices $\{W_a\}_a$ for Problem $\mathbb{D}$ and $\{W_a'\}_a$ for Problem $\mathbb{C}$. □

Solving Problem $\mathbb{D}'$. In view of Proposition 5.5.5, it remains to compute matrices $\{V_a\}_a$ from the matrices $\{\Delta_{ab}\}_{a,b}$. We achieve this by standard linear algebra, and combining with Algorithm $\mathcal{A}_{\mathbb{C}}$ (Algorithm 10)) describes a full algorithm for Problem $\mathbb{D}$.

From now on we assume $t \geq 3$. This is not a major restriction and we explain in Remark 5.5.6 *(ii)* why we exclude $t = 2$. Let $\Delta_{ab} = V_a G_b - V_b G_a$ for $1 \leq a, b \leq t$ as in Problem $\mathbb{D}'$. Let $r = n - p$ and r_{ab} be the rank of Δ_{ab}; clearly, $r_{ab} \leq \min(2r, p)$. We further assume $p > 2n/3$, or equivalently $2r < p$, which is a necessary condition as otherwise the matrices Δ_{ab} likely have full rank p and thus cannot reveal any information. We define $\mathcal{K}_{ab} := \mathrm{im}(\Delta_{ab}) = \mathcal{K}_{ba} \subseteq \mathbb{Q}^p$ and

$$\mathcal{K}_a = \bigcap_{1 \leq b \leq t, b \neq a} \mathcal{K}_{ab} \ , \ 1 \leq a \leq t \ .$$

For $1 \leq a \leq t$, denote by $\mathcal{V}_a$ the image of the matrix V_a. We first argue that, heuristically, $\mathcal{V}_a \subseteq \mathcal{K}_{ab}$ for every $b \neq a$. Let $v \in \mathcal{V}_a$. If there exists $x \in \mathbb{Q}^p$ such that $v = V_a G_b x$ and $V_b G_a x = 0$ then $v = \Delta_{ab} x$, i.e. $v \in \mathcal{K}_{ab}$. Such an element x must therefore lie in $(x_0 + \ker(V_a G_b)) \cap \ker(V_b G_a)$, where $x_0 \in \mathbb{Q}^p$ is any vector such that $v = V_a G_b x_0$. It is easy to see that this intersection is non-empty if $\ker(V_a G_b) + \ker(V_b G_a) = \mathbb{Q}^p$. Heuristically, as the matrices $\{V_a\}_a$ have rank r, $\ker(V_a G_b) + \ker(V_b G_a)$ has dimension at least $2(p-r)$; accordingly we can heuristically expect that $\ker(V_a G_b) + \ker(V_b G_a) = \mathbb{Q}^p$ as soon as $2(p - r) > p$, i.e.

$$\frac{p}{n} > \frac{2}{3} \ .$$

We now justify that, heuristically under a suitable parameter selection, $\mathcal{K}_a = \mathcal{V}_a$ for every $1 \leq a \leq t$. For a fixed a, we compute $\mathcal{K}_a$ modulo $\mathcal{V}_a$ and consider $\overline{\mathcal{K}_{ab}} := \mathcal{K}_{ab}/\mathcal{V}_a \subseteq \mathbb{Q}^{p-r}$ for $b \neq a$. Then $\mathcal{K}_a = \mathcal{V}_a$ if and only if $\overline{\mathcal{K}_a} := \bigcap_{b \neq a} \overline{\mathcal{K}_{ab}} = \{0\}$. Since $\mathcal{V}_a$ has dimension r, $\overline{\mathcal{K}_{ab}}$ has dimension $r_{ab} - r$. For every $b \neq a$, we view $\overline{\mathcal{K}_{ab}}$ as the kernel of $\mathbb{Q}^{p-r} \to \mathbb{Q}^{p-r}/\overline{\mathcal{K}_{ab}}$, represented by a matrix $A_{ab} \in \mathbb{Q}^{(p-r_{ab}) \times (p-r)}$. Therefore $\overline{\mathcal{K}_a}$ is represented by an augmented matrix $A_a = [A_{a1}| \ldots |A_{a,a-1}|A_{a,a+1}| \ldots |A_{at}]$ describing the kernel of $\mathbb{Q}^{p-r} \to \bigoplus_{b \neq a} \mathbb{Q}^{p-r}/\overline{\mathcal{K}_{ab}}$. The matrix A_a has $\sum_{1 \leq b \leq t, b \neq a}(p - r_{ab})$ rows and $p - r$ columns. Now, $\mathcal{K}_a = \mathcal{V}_a$ if and only if A_a has full rank. Heuristically, we expect this to be the case as soon as:

$$\sum_{1 \leq b \leq t, b \neq a} (p - r_{ab}) \geq p - r \ . \tag{5.6}$$

Remark 5.5.6. *(i)* In fact, we expect that $r_{ab} = 2r$ for every a, b. Then, from Equation (5.6), we heuristically expect that $\mathcal{K}_a = \mathcal{V}_a$ for every a, if $(t-1)(p-2r) \geq p - r$, i.e.

$$\frac{p}{n} \geq \frac{2t-3}{3t-5} \text{ or, equivalently, } t \geq \frac{2p-n}{3p-2n} + 1 \ . \tag{5.7}$$

(ii) We assumed $t \geq 3$ so that the intersections $\{\mathcal{K}_a\}_a$ are well-defined. If $t = 2$, $\mathcal{K}_1$ coincides with the image of Δ_{12}, which will not reveal V_1 and V_2.

We compute bases of $\{\mathcal{K}_a\}_a$ by standard linear algebra. For the rest of this section, we will make the assumption that $\mathcal{K}_a = \mathcal{V}_a$ for every a, and let C_a be a basis matrix for $\mathcal{K}_a$. Thus, there exists $M_a \in \mathrm{GL}(r, \mathbb{Q})$ such that $V_a = C_a M_a$. This gives for $a < b$:

$$\Delta_{ab} = V_a G_b - V_b G_a = C_a(M_a G_b) - C_b(M_b G_a) = C_a N_{ab} - C_b N_{ba} \tag{5.8}$$

with $N_{ab} = M_a G_b$. In Equation (5.8), Δ_{ab} and C_a, C_b are known, which allows to compute the unknown matrix $N^{(ab)} = [N_{ab}|N_{ba}]^T$ as a solution to $\Delta_{ab} = [C_a| - C_b] \cdot N^{(ab)}$. Once the matrices $\{N_{ab}\}_{a,b}$ are computed, we obtain a system of linear equations over $\mathbb{Q}$, given by

$$M_a^{-1} \cdot N_{ab} = G_b \quad , \quad 1 \leq a < b \leq t \; . \tag{5.9}$$

Since there are $\frac{1}{2}t(t-1)$ choices for pairs (a, b) and for each pair the matrix equality in Equation (5.9) gives rp equations, the system of equations has $\frac{1}{2}t(t-1)rp$ equations. Moreover, it has $tr^2 + trp = trn$ variables, given by the tr^2 entries of the matrices $\{M_a^{-1} : 1 \leq a \leq t\}$ and the trp entries of the matrices $\{G_b : 1 \leq b \leq t\}$. Heuristically, if the system has at least as many equations as variables, that is, $trn \leq \frac{1}{2}t(t-1)rp$, or equivalently, $2n \leq (t-1)p$, then it is expected to have a unique solution. This bound is automatically satisfied if Equation (5.7) holds. This reveals the matrices $\{M_a : 1 \leq a \leq t\}$ and thus the matrices $\{V_a : 1 \leq a \leq t\}$ by computing $V_a = C_a M_a$.

Proposition 5.5.7. *Assume that $\mathcal{K}_a = \mathcal{V}_a$ for every $1 \leq a \leq t$ (see Remark 5.5.6 (i)). Then, a unique solution to Problem $\mathbb{D}'$ is implied by the existence of a unique solution to the explicit system of linear equations given in Equation* (5.9)*, which is a system of $\frac{1}{2}t(t-1)(n-p)p$ linear equations in $t(n-p)n$ variables. There are at least as many equations as variables as soon as $p(t-1) \geq 2n$.*

Algorithm for Problem $\mathbb{D}$. We summarize our algorithm for Problem $\mathbb{D}$ as follows and will refer to it as Algorithm $\mathcal{A}_{\mathbb{D}}$ for the rest of this chapter.

Remark 5.5.8. Problem $\mathbb{D}$ of Definition 5.1.1 is symmetric in the sense that the matrices P and Q have the same rank. An asymmetric variant consists in having P and Q of ranks $p \neq q$. Note that our algorithm $\mathcal{A}_{\mathbb{D}}$ adapts to that case: if $p < q$, then "cutting" the last $q - p$ columns of $\{W_a\}_a$ means "cutting" the last $q-p$ columns of Q, which reduces to the symmetric case. This approach is however not very genuine, as it "cuts" information instead of possibly exploiting it. We leave it open to design a better algorithm for the asymmetric case.

5.5.2 Optimization of the parameters

We find minimal possible values for t and p with respect to a given n. In Section 5.6.2.1 we will see that it is of interest to minimize $2p+t$ in order to minimize the number of public encodings in [CLT13]. According to Equation (5.7), the main (heuristic) condition to be ensured is $p \geq \frac{2t-3}{3t-5}n$. We therefore set

$$F_n(t) := 2p_n(t) + t := \frac{2t-3}{3t-5}n + t \; , \quad t \in \mathbb{R}_{>0} \setminus \{5/3\} \; , \; n \geq 2 \; .$$

Direct calculations show that F_n has a minimum at $t_0 = \frac{1}{3}(\sqrt{2n} + 5)$, and we have

$$p_n(t_0) = \frac{2}{3}n + \frac{1}{3\sqrt{2}}\sqrt{n} \; , \; F_n(t_0) = \frac{4}{3}n + \frac{2}{3}\sqrt{2n} + \frac{5}{3} \; .$$

In conclusion, we heuristically expect Algorithm $\mathcal{A}_{\mathbb{D}}$ to succeed for the choice $p = \lceil p_n(t_0) \rceil$ and $t = \lceil t_0 \rceil$.

Algorithm 11 Algorithm for Problem $\mathbb{D}$ (Algorithm $\mathcal{A}_{\mathbb{D}}$)

Parameters: Integers $2 \leq p \leq n$ and $t \geq 3$

Input: A valid input for Problem $\mathbb{D}$ as in Definition 5.1.1

Output: "Success" or "Fail", and in case of "Success", additionally output a solution. "Success" means that the computed solution is unique; "Fail" means that a solution was not found.

for (a, b) with $1 \leq a \neq b \leq t$ **do**
  Compute the matrix $\Delta_{ab} = W_a W_0^{-1} W_b - W_b W_0^{-1} W_a$
end for
for $1 \leq a \leq t$ **do**
  Compute a basis matrix C_a of $\mathcal{K}_a := \bigcap_{1 \leq b \leq t, b \neq a} \mathrm{im}(\Delta_{ab})$
end for
if $\dim(\mathcal{K}_a) \neq n - p$ for all $1 \leq a \leq t$ **then**
  Output "Fail" and break
else
  for $1 \leq a < b \leq t$ **do**
    Compute N_{ab} as solutions to $\Delta_{ab} = [C_a| - C_b] \cdot [N_{ab}|N_{ba}]^T$
  end for
end if
Solve for $\{M_a\}_a$ the system of equations $M_a^{-1} N_{ab} = G_b$ for $(a, b) \in \{1, \ldots, t\}^2, a < b$.
if a unique solution is not found **then**
  output "Fail" and break
end if
for $1 \leq a \leq t$ **do**
  Compute $V_a = C_a \cdot M_a$
  Compute $W'_a = [W_a|V_a]$
end for
Run Algorithm $\mathcal{A}_{\mathbb{C}}$ on the input matrices $W'_0 = [W_0|0]$ and $\{W'_a\}_a$ and return its output

5.6 Applications of our algorithms

In this section, we describe two applications for our algorithms and obtain significant improvements on previous works. In the first place, we consider the CRT-ACD Problem stated in Definition 2.2.7, and the algorithm designed by Coron and Pereira, which we recalled in Section 2.2.3.2. In second place, we consider the cryptanalysis of the CLT13 multilinear map by Cheon et al. which we recalled in Section 3.2.2.2. For both problems, we describe improved algorithms requiring a smaller input size.

5.6.1 Improved algorithm for the CRT-ACD Problem

We again consider the Approximate Common Divisor Problem from Definition 2.2.7 based on the Chinese Remainder Theorem, already considered in Section 4.8.1 in the context of the hidden lattice problem. This time, our improvement is of a different nature. Following the same notation, let $N = \prod_{i=1}^{n} p_i$ be a composite squarefree public modulus, and let $\mathcal{S}$ be a set of integers which have somewhat small residues in $\{0, \ldots, p_i - 1\}$, modulo every prime p_i; formally their size is controlled by the parameter ρ, as in Definition 2.2.7. Clearly, the larger the set $\mathcal{S}$ is, the more information one can exploit to factor N. Our interest is therefore to *minimize* the cardinality of the set $\mathcal{S}$ with respect to n. As seen in the algorithm by Coron and Pereira (see Algorithm 1 in Section 2.2.3.2), there is a practical algorithm to factor N using $\#\mathcal{S} = n + 1$, where n is the number of prime factors of N.

5.6.1.1 A naive improvement

There is a naive generalization of the algorithm in [CP19] using only $O(\sqrt{n})$ public instances in $\mathcal{S}$, while, however, requiring a range of parameters which is worse than in [CP19].

For integers $p \geq 2$ and $t \geq 1$ of size $O(\sqrt{n})$, let x be the vector $(y_1 z, \ldots, y_t z) \in \mathbb{Z}^{tp}$ (note that the dimension tp is $O(n)$) for elements $y_1, \ldots, y_t \in \mathcal{S}$ and $z \in \mathcal{S}^p$. This variant clearly reduces $\#\mathcal{S}$ considerably, as $\#\mathcal{S} = p + t = O(\sqrt{n})$. However, the algorithm from [CP19] (see Section 2.2.3.2) requires to construct the vector $b = (x, y \cdot x)$ for $y \in \mathcal{S}$. This gives rise to mod-p_i-residue vectors $\{b^{(i)}\}_i$ of Euclidean norm approximately $2^{3\rho}$ instead of $2^{2\rho}$ as in [CP19]. This impacts the first step of the algorithm to perform worse than originally described. Namely, in rough terms, the stronger condition $3\rho < \eta$ will be required for the orthogonal lattice attack to succeed. If this condition is not fulfilled, the lattice reduction step in the orthogonal lattice algorithm fails to compute a basis of the completed lattice $\overline{\mathcal{L}}$ (where $\mathcal{L}$ again denotes the lattice generated by the vectors $\{b_i\}_i$).

In our improvement, we would like to avoid this situation, that is, we will lower $\#\mathcal{S}$, while continuing to use the condition $2\rho < \eta$, exactly as in [CP19].

5.6.1.2 Our improved algorithm

Our main point in improving the algorithm from [CP19] lies in recognizing that Equation (2.6) defines an instance of Problem $\mathbb{A}$ of Definition 5.1.1 with $t = 1$ because the matrices P and Q have rank n. Our improvement lies in generalizing the vector b as to obtain an instance of Problem $\mathbb{C}$, instead.

We consider $\#\mathcal{S} < n+1$. For convenience, we write $\mathcal{S} = \{x_1, \ldots, x_p, y_1, \ldots, y_t\}$ with integers $2 \leq p < n$ and $2 \leq t < n$ satisfying $2n \leq (t+1)p$. We let $x = (x_1, \ldots, x_p) \in \mathcal{S}^p$ and

$$b = (x, y_1 \cdot x, \ldots, y_t \cdot x) \in \mathbb{Z}^{(t+1)p} \,. \tag{5.10}$$

As previously, let $\{b^{(i)}\}_i \subseteq \mathbb{Z}^{(t+1)p}$ denote the short residue vectors modulo the primes $\{p_i\}_i$ and $x \equiv x^{(i)} \pmod{p_i}$, $y_a \equiv y_a^{(i)} \pmod{p_i}$ for $1 \leq a \leq t$ and $1 \leq i \leq n$. By the Chinese Remainder Theorem, we observe that b lies in the lattice $\mathcal{L} = \bigoplus_{i=1}^n \mathbb{Z} b^{(i)}$ modulo N. Namely, there are integers $c_1, \ldots, c_n \in \mathbb{Z}$ such that

$$b \equiv \sum_{i=1}^n c_i \begin{bmatrix} x^{(i)} \\ y_1^{(i)} \cdot x^{(i)} \\ \vdots \\ y_t^{(i)} \cdot x^{(i)} \end{bmatrix} =: \sum_{i=1}^n c_i b^{(i)} \,. \pmod{N}$$

As in [CP19], the orthogonal lattice algorithm reveals a basis $\{b'^{(i)}\}_i$ of $\overline{\mathcal{L}}$ and the Euclidean norm of $\{b^{(i)}\}_i$ is still approximately $2^{2\rho}$. As explained in Chapter 4 (in particular, Section 4.8.1), this equation defines a hidden lattice problem, and a basis of $\overline{\mathcal{L}}$ may therefore be computed by using Algorithm II (see Algorithm 8).

Contrary to Equation (2.6), the new choice of our vector b in Equation (5.10) now gives rise to matrix equations

$$W_0 = P \cdot Q \,,\; W_a = P \cdot U_a \cdot Q \,,\; 1 \leq a \leq t \,, \tag{5.11}$$

where $P \in \mathbb{Z}^{p \times n}$ has columns $\{x^{(i)}\}_i$ and $\{U_a\}_a \subseteq \mathbb{Z}^{n \times n}$ are diagonal matrices with entries $\{y_a^{(i)}\}_{a,i}$. The matrix Q is a base change matrix from $\{b'^{(i)}\}_i$ to $\{b^{(i)}\}_i$, again, due to the base change from $\overline{\mathcal{L}}$ to $\mathcal{L}$, as explained in Section 2.2.3.2. If W_0 has rank p, then Equation (5.11) defines a valid input for Problem $\mathbb{C}$ of Definition 5.1.1. Then Algorithm $\mathcal{A}_{\mathbb{C}}$ (Algorithm 10) from Section 5.4 reveals the diagonal entries $\{y_a^{(i)}\}_{a,i}$ of the matrices $\{U_a\}_a$. One finally factors N by computing $\gcd(y_a - y_a^{(i)}, N)$.

Parameters. From Section 5.4.2 we see that $\#\mathcal{S} = p + t$ is minimized when $p = \lceil \sqrt{2n} \rceil$ and $t + 1 = \lceil \sqrt{2n} \rceil$. Consequently,

$$\#\mathcal{S} = 2\lceil \sqrt{2n} \rceil = O(\sqrt{n}) \,.$$

In summary, letting n be the number of secret prime factors in the public modulus N, we can factor N given only $O(\sqrt{n})$ input samples, whereas the algorithm [CP19] uses $O(n)$ input samples. This gives a quadratic improvement on the input length.

Remark 5.6.1. More generally, the work [CP19] proposes a multiparty key-exchange protocol built from the CLT13 multilinear map, by relying on the hardness of the CRT-ACD Problem. We remark however that our algorithms do not impact the security of the key-exchange from [CP19]. Namely, it uses certain encodings of matrices, which would us require to work with sets of matrices instead. However, the product of matrices does not commute, so our techniques do not apply to that case.

Algorithm 12 Algorithm for the CRT-ACD Problem with $\#\mathcal{S} = O(\sqrt{n})$

Parameters: The CRT-ACD parameters (Definition 2.2.7)
Input: An integer $N = \prod_{i=1}^{n} p_i$ and a set $\mathcal{S}$ as in Definition 2.2.7 with $\#\mathcal{S} = 2\lceil\sqrt{2n}\rceil$
Output: The prime factors $\{p_i : 1 \leq i \leq n\}$ of N

1: Let $p = \lceil\sqrt{2n}\rceil$ and $t = p = \lceil\sqrt{2n}\rceil - 1$. Write the elements in $\mathcal{S}$ as $x_1, \ldots, x_p, y_1, \ldots, y_t$ and let $x = (x_1, \ldots, x_p)$. Construct the vector $b = (x, y_1 \cdot x, \ldots, y_t \cdot x) \in \mathbb{Z}^{(t+1)p}$ from $\mathcal{S}$.
2: Given b, compute a basis of $\overline{\mathcal{L}}$, where $\mathcal{L}$ is the lattice generated by the vectors $b^{(i)}$ for $1 \leq i \leq n$; denote by $[W_0|W_1|\ldots|W_t] \in \mathbb{Z}^{n \times p(t+1)}$ the computed basis of $\overline{\mathcal{L}}$. This can be carried out by Algorithm I 7 or Algorithm 8 from Chapter 4.
3: Run Algorithm $\mathcal{A}_{\mathbb{C}}$ (Algorithm 10) on the input matrices W_0 and $\{W_a : 1 \leq a \leq t\}$. Denote the computed tuples of diagonal entries by $\{y_a^{(i)}\}_{a,i}$.
4: **for** $1 \leq i \leq n$ **do**
5: Compute and return $\gcd(y - y_1^{(i)}, N)$
6: **end for**

Algorithm. Our new algorithm improving Algorithm 1 is summarized in Algorithm 12.

5.6.2 Improved Cryptanalysis of CLT13 Multilinear Maps

We consider the CLT13 multilinear map scheme by Coron et al. from [CLT13], with which already Chapter 3 deals more extensively. Our second application concerns the algorithm by Cheon et al. from Section 3.2.2.2, a polynomial-time attack against the Diffie-Hellman key exchange protocol based on CLT13 when enough encodings of zero are available to the attacker.

It is of interest to investigate this line of cryptanalysis on CLT13 when only a limited number of such encodings is available. Namely, not every CLT13-based construction necessarily reveals enough such encodings, and the attack of Cheon et al. is in such a case prevented.

5.6.2.1 Attacking CLT13 with fewer encodings

We formally introduce the following CLT13-based problem, of interest in the framework of the Cheon et al. attack.

Definition 5.6.2 (CLT13 Problem)**.** *Let $n \geq 2$ be the dimension of* CLT13 *and $x_0 = \prod_{i=1}^{n} p_i$ following the classical CLT13-notation. Let $\mathcal{E}$ be a finite non-empty set of encodings at level 1 and $\mathcal{E}_0 \subseteq \mathcal{E}$ a non-empty subset such that every element of $\mathcal{E}_0$ is an encoding of zero.*
The CLT13 *Problem is as follows: Given the sets $\mathcal{E}$ and $\mathcal{E}_0$, factor x_0.*

We refer to the sets $\mathcal{E}$ and $\mathcal{E}_0$ as the sets of "available encodings" and "available encodings of zero", respectively. It is not a loss of generality to consider level-one encodings in these sets, as encodings can always be multiplied by other encodings to increase their level. As in Section 3.2.2.2, we write $\mathcal{E} = \mathcal{A} \cup \mathcal{B} \cup \mathcal{C}$ with $\mathcal{A} \subseteq \mathcal{E}_0$, for disjoint sets $\mathcal{A}, \mathcal{B}$ and $\mathcal{C}$. As mentioned in Section 3.2.2.2, the algorithm by Cheon et al. requires $\#\mathcal{E}_0 \geq n$ to factor x_0, and a total number of public encodings $\#\mathcal{E} = 2n + 2$, where n is the number of prime factors dividing x_0.

For an improved algorithm, we aim at reducing the number of encodings needed for the factorization of x_0, and therefore raise the following problems related to Definition 5.6.2, and which we shall treat independently below:

($\star$) Factor x_0 with fewer available encodings of zero, i.e. $\#\mathcal{E}_0 < n$

($\star\star$) Factor x_0 with fewer available encodings, i.e. $\#\mathcal{E} < 2n + 2$

A naive improvement. Somewhat similarly to the CRT-ACD Problem treated in Section 5.6.1, there is a naive improvement of the Cheon et al. attack, using fewer encodings (of zero), but with the cost of assuming $\kappa = 4$. Namely, when $\kappa = 4$, one can form product encodings of the form $\alpha_j\beta_a\gamma_k\delta_\ell$ at level 4, where every encoding is at level 1. Such encodings can then be partitioned into sets $\mathcal{A}, \mathcal{B}$ and $\mathcal{C}$ such that $\mathcal{A}$ contains only encodings of zero and $\#\mathcal{A} = O(\sqrt{n})$. However, this approach has the inconvenience of using $\kappa = 4$ and our improved attack aims at lowering the number of public encodings while keeping $\kappa = 3$.

5.6.2.2 Our improved algorithm

We treat Questions ($\star$) and ($\star\star$) independently. In particular, we explain how to positively answer Question ($\star$) by making use of Algorithm $\mathcal{A}_{\mathbb{C}}$, and Question ($\star\star$) by making use of Algorithm $\mathcal{A}_{\mathbb{D}}$.

Minimizing the number of encodings of zero

Let us first treat Question ($\star$) and explain how to use our Algorithm $\mathcal{A}_{\mathbb{C}}$ to factor x_0 using only $\#\mathcal{E}_0 = O(\sqrt{n})$ level-one encodings of zero.

We fix integers $2 \leq p < n$ and $3 \leq t < n$ and assume again $\kappa = 3$. As in Section 3.2.2.2, we write $\mathcal{E} = \mathcal{A} \cup \mathcal{B} \cup \mathcal{C}$ with $\mathcal{A} \subseteq \mathcal{E}_0$. Further, we let $\#\mathcal{A} = p$, $\#\mathcal{B} = t$ and $\#\mathcal{C} = n$; we claim that we can factor x_0 with $p = O(\sqrt{n})$.

Every product encoding $c = \alpha_j\beta_a\gamma_k$ with $(\alpha_j, \beta_a, \gamma_k) \in \mathcal{A} \times \mathcal{B} \times \mathcal{C}$ is an encoding of zero and by correctly zero-testing c using the zero-test parameter, we obtain integer matrix relations

$$W_a = P \cdot U_a \cdot Q \quad , \quad 1 \leq a \leq t \tag{5.12}$$

where $P \in \mathbb{Z}^{p\times n}$ and $Q \in \mathbb{Z}^{n\times n}$ are matrices corresponding to the encodings in $\mathcal{A}$ and $\mathcal{C}$, respectively, and $\{U_a\}_a \subseteq \mathbb{Z}^{n\times n}$ are diagonal matrices corresponding to encodings in $\mathcal{B}$. Exactly as in Section 3.2.2.2, the matrices $\{U_a\}_a$ contain diagonal entries β_{ai} such that $\beta_a \equiv \beta_{ai} \pmod{p_i}$ for $1 \leq i \leq n$, where $\{p_i\}_i$ are the prime factors of x_0. With high probability the ranks of P and Q are p and n, respectively. Defining $W'_0 = W_1$ and $W'_a = W_{a-1}$ for $2 \leq a \leq t$ we obtain an instance similar to Problem $\mathbb{C}$ of Definition 5.1.1, but without a "special input matrix" PQ (see Section 5.3). Using Algorithm $\mathcal{A}_{\mathbb{C}}$, we reveal the eigenvalues (the diagonal entries) of the matrix products $\{U_aU_1^{-1}\}_a$ as it is likely that U_1 be invertible. We finally deduce the prime factorization of x_0 by computing greatest common divisors, as in Section 3.2.2.2.

By the optimization of the parameters in Section 5.4.2, we can choose $\#\mathcal{B} = t = \lceil\sqrt{2n}\rceil$ and $\#\mathcal{A} = p = \lceil\sqrt{2n}\rceil$. It is worth noticing that while the set $\mathcal{B}$ now has $O(\sqrt{n})$ elements compared

to $\#\mathcal{B} = 2$ for the original attack by Cheon et al. attack, the total number of encodings in $\mathcal{E}$ remains $p + t + n = 2\lceil\sqrt{2n}\rceil + n = O(n)$, as for the Cheon et al. attack. Therefore lowering the number of encodings in $\mathcal{A}$ has no impact on the total number of encodings.

Minimizing the total number of encodings

We now consider Question $(\star\star)$. We explain how to use Algorithm $\mathcal{A}_{\mathbb{D}}$ to factor x_0 using $\#\mathcal{E} = \frac{4}{3}n + \mathcal{O}(\sqrt{n})$ instead of $\#\mathcal{E} = 2n + 2$ as in the original Cheon et al. attack. Namely, the original attack, as described in Section 3.2.2.2 uses sets $\mathcal{A}, \mathcal{B}, \mathcal{C}$ with in total $2n + 2$ encodings.

This time, we consider $\mathcal{A}$ and $\mathcal{B}$ as before, but consider a set $\mathcal{C}$ with $\#\mathcal{C} = p$ instead of $\#\mathcal{C} = n$ as before; therefore we have a total number of encodings $\#\mathcal{E} = 2p + t$. It is now easy to see that upon correct zero-testing we derive matrix equations as in Equation (5.12) but with $Q \in \mathbb{Z}^{n\times p}$ this time. Hence, if both matrices P and Q have rank p, we obtain an instance of Problem $\mathbb{D}$ of Definition 5.1.1 without "special input matrix" W_0. Then Algorithm $\mathcal{A}_{\mathbb{D}}$ reveals ratios of diagonal entries of $\{U_a U_1^{-1}\}$ (again assuming the likely invertibility of U_1) and we consequently factor x_0.

Following Section 5.5.2, we are led to minimize the number of encodings, i.e. the quantity $\#\mathcal{E}(n) = 2p + t$, viewed as a function of n. We can let $p = \lceil \frac{2}{3}n + \frac{1}{3\sqrt{2}}\sqrt{n}\rceil$ and $t = \lceil \frac{1}{3}\sqrt{2n} + \frac{5}{3}\rceil$ and therefore obtain

$$\#\mathcal{E}(n) = 2\left\lceil \frac{2}{3}n + \frac{1}{3\sqrt{2}}\sqrt{n}\right\rceil + \left\lceil \frac{1}{3}\sqrt{2n} + \frac{5}{3}\right\rceil = \frac{4}{3}n + O(\sqrt{n})\ .$$

Algorithms. We summarize our improvements of the Cheon et al. attack (compare with Algorithm 2) below.

Algorithm 13 Improved Cheon et al. attack against CLT13 with fewer encodings

Parameters: The CLT13 parameters with $\kappa = 3$

Input: Disjoint sets $\mathcal{A}, \mathcal{B}, \mathcal{C}$ of encodings at level one, where the encodings in $\mathcal{A}$ encode zero, and $\#\mathcal{A} = p, \#\mathcal{B} = t$ and $\#\mathcal{C} = n$ in the framework of $(\star)$, and $\#\mathcal{A} = \#\mathcal{C} = p, \#\mathcal{B} = t$ in the framework of $(\star\star)$, for integers $2 \le p < n$ and $3 \le t < n$; and $x_0 = \prod_{i=1}^{n} p_i$

Output: The prime factors $\{p_i : 1 \le i \le n\}$ of x_0

1: **for** $1 \le a \le t$ **do**
2: Construct the matrix $W_a \in \mathbb{Z}^{p\times n}$ as in Equation (5.12) in the framework of $(\star)$. These matrices are in $\mathbb{Z}^{p\times p}$ in the framework of $(\star\star)$.
3: **end for**
4: Run Algorithm $\mathcal{A}_{\mathbb{C}}$ (Algorithm 10) on the input matrices $W_0' = W_1$ and $\{W_a'\}_{2\le a\le t}$ with $W_a' = W_{a-1}$ in the case of framework $(\star)$; in the case of framework $(\star\star)$, run Algorithm $\mathcal{A}_{\mathbb{D}}$ (Algorithm 11). Denote by $\{\beta_{ai}/\beta_{1i}\}_{a,i}$ the computed tuples of diagonal entries
5: **for** $1 \le i \le n$ **do**
6: Compute coprime integers x_i, y_i such that $\beta_{2i}/\beta_{1i} = x_i/y_i$
7: Compute and return $\gcd(x_i\beta_2 - y_i\beta_1, x_0)$, where $\mathcal{B} = \{\beta_i\}_i$
8: **end for**

Cryptanalysis with independent slots. As explained in Chapter 3 (see Section 3.3.9.1), one can adapt the Cheon et. al attack in order to cryptanalyze the CLT13 Scheme when no encodings of zero are available beforehand, but instead only "partial-zero" encodings. We can improve this cryptanalysis (i.e. Algorithm 6) following the same techniques as above. Following the notation from Section 3.3.9.1, let ℓ the number of partial-zero encodings. More precisely, in order to factor x_0, we can now replace the Cheon et al. attack used as a sub-algorithm (Step 4) in Algorithm 6 by Algorithm $\mathcal{A}_{\mathbb{C}}$, once ℓ encodings of zero are created. This means that we can set $\ell = O(\sqrt{n})$, which brings a twofold improvement: first, lattice reduction (in the orthogonal lattice attack from Section 3.3.6) is only run on a lattice of dimension $O(\sqrt{n})$, and second, the number of partial-zero encodings is reduced to $O(\sqrt{n})$.

5.7 Computational Aspects and Practical Results

Let us now describe practical parameters for our algorithms $\mathcal{A}_{\mathbb{C}}$ and $\mathcal{A}_{\mathbb{D}}$ designed in Sections 5.4 and 5.5 of this chapter. We have implemented our algorithms in SageMath [S$^+$17]; our source code is provided in [CNW20b]. Our experiments are done on a standard $3,3$ GHz Intel Core i7 processor.

5.7.1 Instance Generation of Problems $\mathbb{C}$ and $\mathbb{D}$

As is the case for many applications in cryptanalysis, we consider matrices with integer entries in Definition 5.1.1. To generate instances of Problems $\mathbb{C}$ and $\mathbb{D}$, given fixed integers n, t, p, we uniformly at random generate matrices P, Q and $\{U_a\}_a$ as in Definition 5.1.1 with entries in $[-k, k] \cap \mathbb{Z}$ for some $k \in \mathbb{Z}_{\geq 1}$. Setting $W_0 = PQ$ and $W_a = PU_aQ$ for $1 \leq a \leq t$ gives rise to instances of Problems $\mathbb{C}$ or $\mathbb{D}$.

In order to speed up our algorithms, note that we can perform the linear algebra over $\mathbb{Z}/\ell\mathbb{Z}$ for a sufficiently large prime ℓ, instead of over $\mathbb{Q}$. It suffices to choose ℓ slightly larger than the size of the output, i.e. the diagonal entries of $\{U_a\}_a$. Thus, we may choose for example $\ell = O(\max_{a,i} |u_{ai}|)$, where u_{ai} for $1 \leq i \leq n, 1 \leq a \leq t$ denote the diagonal entries of U_a. The running time of our algorithms clearly depends on the entry size of the generated matrices. The overall computational cost of our algorithms $\mathcal{A}_{\mathbb{C}}$ and $\mathcal{A}_{\mathbb{D}}$ is dominated by the cost of solving systems of linear equations and performing simultaneous diagonalization, which can be done by standard algorithms for non-sparse linear algebra.

5.7.2 Practical Experiments

We here gather practical parameters for Problems $\mathbb{C}$ and $\mathbb{D}$, as well as our applications of Section 5.6.

Problems $\mathbb{C}$ and $\mathbb{D}$. We create instances of Problems $\mathbb{C}$ and $\mathbb{D}$ and run Algorithm $\mathcal{A}_{\mathbb{C}}$ and $\mathcal{A}_{\mathbb{D}}$ with increasing values for p and t. We then select the minimal practical value of p and t such that our algorithms reveal the diagonal entries of $\{U_a\}_a$ successfully. We then compare these practical values p, t with the minimal theoretical values $p_0(n), t_0(n)$ derived in Sections 5.4, 5.5. Table 5.1 and Table 5.2 gather practical data for Problem $\mathbb{C}$ and Problem $\mathbb{D}$, respectively. Therefore, for the data in Table 5.1, we have selected $p_0(n) = \lceil \sqrt{2n} \rceil$ and $t_0(n) = \lceil \sqrt{2n} - 1 \rceil$ (as predicted by Section 5.4.2), while for the data in Table 5.2, we have selected $p_0(n) = \lceil \frac{2}{3}n + \frac{\sqrt{n}}{3\sqrt{2}} \rceil$

and $t_0(n) = \lceil \frac{1}{3}(\sqrt{2n} + 5) \rceil$ (as predicted by Section 5.5.2). In both tables, "Entry size" is an approximation of the bit-size of the max-norm of each input matrix.

We have chosen an increasing size n for the diagonal matrices (up to $n = 500$), together with a decreasing entry size as to keep a somewhat balanced running time. We observe that the practical values for p and t almost perfectly align with the theoretical values $p_0(n)$ and $t_0(n)$ in both cases. This also allows to point out the practical optimality of our algorithms: it seems from these practical outputs that our theoretical bounds are not only optimal from a theoretical perspective, but also from a practical one.

Applications. Table 5.3 gathers data for our two applications described in Section 5.6. More precisely, Series 1 of Table 5.3 compares our new algorithm (Algorithm 12) for the CRT-ACD Problem with [CP19] (see Algorithm 1 in Section 2.2.3.2). The parameters provided in Series 1 allow a complete factorization of the public modulus $N = \prod_{i=1}^{n} p_i$ with prime numbers of bit size η, and noise bit size ρ in the elements of $\mathcal{S}$. In particular, p and t are chosen to be minimal with respect to revealing the prime factors of N successfully. The column "Num. of samples" compares the cardinality of the public set $\mathcal{S}$ for both algorithms; the column "this work" is thus set as $\#\mathcal{S} = p + t$, while for [CP19] one has $\#\mathcal{S} = n + 1$. Our algorithm is practical and significantly improves on [CP19]; for instance, for $n = 50$, our algorithm factors N using only 19 public samples in $\mathcal{S}$, whereas [CP19] requires 51 samples to factor N for the same instance.

Similarly, Series 2 and 3 of Table 5.3 compare our new cryptanalysis of CLT13 (Algorithm 13) with the previous work of [CHL$^+$15] (see Algorithm 2 in Section 3.2.2.2). Here, Series 2 deals with the case of minimizing the number of total encodings, which bases on Algorithm $\mathcal{A}_{\mathbb{D}}$, and Series 3 deals with the case of minimizing the number of encodings of zero available to the attacker, which bases on Algorithm $\mathcal{A}_{\mathbb{C}}$. The values of p and t in Series 2 and 3 are again selected minimal as to allow a complete factorization of $x_0 = \prod_{i=1}^{n} p_i$, as for the Cheon et al. attack. The columns "Num. of encodings" resp. "Num. of encodings of zero" compare the number of encodings in the set $\mathcal{E}$ resp. $\mathcal{E}_0$ in both algorithms. For Series 2, "this work" is therefore set as $\#\mathcal{E} = 2p + t$ and Series 3 uses $\#\mathcal{E}_0 = p$.

In particular, our new algorithm breaks CLT13 with only 10 public encodings of zero, while the original algorithm [CHL$^+$15] uses 50 encodings of zero. While not being as remarkable, also the number of total encodings (Series 2) is improved, as predicted by the bound $O(4/3n)$.

		Practice		Theory		
n	Entry size	p	t	$p_0(n)$	$t_0(n)$	Running time
15	1000	6	4	6	5	4 min 4 s
25	750	8	7	8	7	3 min 45 s
50	600	10	9	10	9	4 min 34 s
100	200	15	14	15	14	1 h 17 min
150	100	18	16	18	17	6 h 29 min
500	20	32	31	32	31	29 min 3 s

Table 5.1: Experimental data for Algorithm $\mathcal{A}_{\mathbb{C}}$

		Practice		Theory		
n	Entry size	p	t	$p_0(n)$	$t_0(n)$	Running time
15	1000	11	4	11	4	4 min 2 s
25	750	18	4	18	5	1 min 54 s
50	600	35	5	35	5	1 min 39 s
100	200	70	7	70	7	5 min 14 s
150	100	103	8	103	8	23 min 14 s
500	20	339	13	339	13	6 min 57 s

Table 5.2: Experimental data for Algorithm $\mathcal{A}_{\mathbb{D}}$

Series 1					Num. of samples	
n	η	ρ	p	t	this work	[CP19]
20	1000	200	7	6	13	21
30	1000	100	8	7	15	31
50	800	100	10	9	19	51
Series 2					Num. of encodings	
n	η	ρ	p	t	this work	[CHL$^+$15]
20	1000	200	15	4	34	42
30	1000	100	22	5	49	62
50	800	100	35	5	75	102
Series 3					Num. of encodings of zero	
n	η	ρ	p	t	this work	[CHL$^+$15]
20	1000	200	7	6	7	20
30	1000	100	8	7	8	30
50	800	100	10	9	10	50

Table 5.3: Experimental data for the CRT-ACD Problem and the CLT13 Problem

CHAPTER 6

Questions related to Edwards Curves

Elliptic curves are important objects in modern number theory and lie at the heart of Elliptic Curve Cryptography (ECC). In particular, mathematicians and cryptographers put a lot of effort in finding suitable models for elliptic curves to guarantee security and efficiency. In 2007, Harold Edwards presented a *normal form* for elliptic curves, building on historical results of Euler and Gauss. Only a few months later this new form for elliptic curves has been proposed for ECC by Bernstein-Lange (see [Edw07, BL07]). Since then, these curves carry Edwards' name and take a central place in modern algorithms for ECC.

In this chapter, we study some questions of geometric, arithmetic and computational nature related to Edwards curves. We provide the reader with several examples illustrating our results.

The content of the present chapter is based on a collaboration with Samuele Anni[1] and an article will be submitted in the near future.

6.1 Introduction

Twisted Edwards curves over a (non-binary) field K are defined by the equation $ax^2 + y^2 = 1 + dx^2y^2$ for distinct elements $a \in K, d \in K \setminus \{0, 1\}$. When $a = 1$, the curve is simply called *Edwards curve* and denoted by $\mathcal{E}_d$. What makes these curves, originally introduced by the work of Harold Edwards [Edw07], interesting for cryptographic purposes is that they carry an efficient group law, and are birationally equivalent to elliptic curves, i.e. there is a birational map $\phi : \mathcal{E}_d \to E$ from $\mathcal{E}_d$ to an elliptic curve E. With the work of Bernstein-Lange and Bernstein et al. [BL07, BBJ$^+$08], Edwards curves entered ECC and compete today with many other forms of elliptic curves, such as Jacobi intersections, Montgomery curves, Hessian curves, or Huff curves. The work [BL07] shows that the group law on Edwards curves gives very efficient arithmetic on elliptic curves and has considerable speed-up factors against other forms of elliptic curves. This also gives a simplified protection against side-channel attacks. Nowadays, there is a vast literature on Edwards curves. The Diffie-Hellman protocol $X25519$ [Ber06] and the signature scheme EdDSA (Edwards-curve Digital Signature Algorithm) [BDL$^+$11] are based on (twists) of the Edwards curve $\mathcal{E}_d$, with $d = 121665/121666$ in the prime field $\mathbb{F}_p$ where $p = 2^{255} - 19$, and offer high security.

[1] Maître de conférences at Aix-Maseille Université

The arithmetic of (twisted) Edwards curves is also of interest from a purely mathematical viewpoint, as for example the study of morphisms between these curves, such as isomorphisms or isogenies of Edwards curves (see e.g. [AG12]). In this chapter, we shall study Edwards curves from an arithmetic-geometric standpoint, without necessarily targetting cryptographic applications.

6.2 Our Contributions

Abstract Construction of the Edwards Model and its Twists. Abstractly, without specification to any affine or projective model, an elliptic curve over a field K is a pair $(E, \mathcal{O})$ where E is a non-singular projective curve over K of genus 1 and $\mathcal{O}$ is a K-rational point. The theorem of Riemann-Roch can be used to find a Weierstrass model for E (see e.g. [Sil09, Chapter III, §3, Proposition 3.1]). Our first contribution is to rely on the theorem of Riemann-Roch to derive the Edwards model for elliptic curves. This way of establishing the Edwards model is new and differs conceptually from the construction of Bernstein-Lange (see [BL07, BBJ$^+$08]). We believe that such more abstract constructions potentially help to construct new elliptic curve models from the understanding of the underlying geometry. Also twists of Edwards curves can be derived more abstractly. It is well-known that the curve $\mathcal{E}_d^a : ax^2 + y^2 = 1 + dx^2y^2$ can be interpreted as a quadratic twist of $\mathcal{E}_{d/a}$ ([BBJ$^+$08]). In the language of group cohomology, there is a natural way of defining *twists* of curves, and we describe a construction for the twist $\mathcal{E}_d^a$ of $\mathcal{E}_{d/a}$ using this formalism.

Explicit construction of Edwards curves from 2-torsion points. An elliptic curve E defined over a field K with a K-rational point of order 4 is birationally equivalent over K to an Edwards curve $\mathcal{E}_d$ defined over K. Following [BL07, Theorem 2.1] and [BBJ$^+$08, Theorem 3.3], we can explicitly construct a birational map from an elliptic curve E over K in Weierstrass form with a K-rational point P of order 4, to an Edwards curve $\mathcal{E}_d$ over K. The map and $d \in K \setminus \{0, 1\}$ are explicit in terms of E, that is, the coefficients of the Weierstrass model of E and the coordinates of P. A natural question is to extend this construction when E does not have a point of order 4 over its base field K. To tackle this question, we are constrained to work over larger fields, over which $E(K)$ has 4-torsion. Our contribution in this direction is an explicit description of the smallest field extension of the base field K over which E admits an Edwards model. One may expect this field to be the extension of K where E has full 4-torsion (the 4-division field of E), but we prove that in some cases it actually is a strict subfield thereof. Our construction only requires 2-torsion points of E, which gives some important advantages compared to [BL07, BBJ$^+$08], which work with 4-torsion. Intuitively, we associate to every point $(e_i, 0)$ ($i = 1, 2, 3$) of order 2 on $E(\overline{K})$ (in characteristic away from $2, 3$ we can always write the 2-torsion in this way), a pair of Edwards curves $\mathcal{E}_{d_i}$, $\mathcal{E}_{1/d_i}$ defined over an explicit field K_i/K, with an explicit formula for d_i depending on e_1, e_2, e_3. We also make precise how large the field K_i is by studying the 2-division field of E. We remark that our construction is compatible also from a purely algebraic (vs. geometric) study of the *j-invariant of E*. Namely, the expression for the j-invariant of an elliptic curve E birationally equivalent to an Edwards curve $\mathcal{E}_d$ is a rational function in d, and can be used to derive all $\overline{K}$-isomorphism classes of E (and pf $\mathcal{E}_d$).

On a different level, this result allows us to recover a known result by Ahmahi and Granger

on $\overline{K}$-isomorphism classes of Edwards curves [AG12]. We also describe a natural Galois action on the Edwards curves, induced from the Galois action on the points on E. Together with this observation, we obtain a stronger version than the result in [AG12].

Rank statistics and invariants for the Edwards family. The twisted Edwards curve $\mathcal{E}_d^a$ is birationally equivalent to the elliptic curve $E_{a,d} : y^2 = x^3 + 2(a+d)x^2 + (a-d)^2x$. We refer to the family $\{E_{a,d}\}_{a,d}$ as the (elliptic) *Edwards family*. For our last contribution in this chapter, we consider this family over $\mathbb{Q}$, with $d \in \mathbb{Z}$ and, mostly, $a = 1$. By Mordell's Theorem, the group $E_{1,d}(\mathbb{Q})$ is finitely generated, and it is a natural question to study its torsion part and its free part. We provide a precise description for the torsion subgroup of $E_{1,d}(\mathbb{Q})$, by means of explicit formulae for torsion points. For the free part, we analyse the rank of these curves computationally. Previous extensive computer calculations were for example carried out by Zagier and Kramarz for the ranks in the family of rational curves $x^3 + y^3 = m$ for cubefree integers m, [ZK87]. Before studying the rank of $E_{1,d}$ we establish a formula for the *conductor* of $E_{1,d}$, an important invariant of E for the arithmetic of its L-function, for example. Together with this remark we see that the conductor of $E_{1,d}$ grows quickly with d, and quadratically for a well-described subfamily of curves. We provide statistical distributions for the ranks in the family $\{E_d\}_d$, which appear to point out to a rather rapid accumulation of curves of rank at least 2. Along with our statistical distribution for the rank, we study the *analytic* order of the Shafarevich-Tate group, which can be computed under the Birch and Swinnerton-Dyer Conjecture. Despite of the large conductors, the L-functions associated to $E_{1,d}$ have a rather slow convergence (in the sense that many terms need be precomputed). To overcome this obstacle, we design a general algorithm (i.e. for general elliptic curves) to compute the order of the Shafarevich-Tate group by using a truncation of the L-series. This algorithm is of independent interest and extends previous work from [HY15].

6.3 Background and notation on Edwards Curves

Throughout the whole chapter, K is a field of characteristic not 2. We denote by $\overline{K}$ an algebraic closure of K. We will sometimes work with $K = \mathbb{Q}$ and will mention it explicitly then. We present some background material on Edwards curves.

6.3.1 Edwards Curves and Twisted Edwards Curves

We consider the definition of twisted Edwards curves from [BL07, BBJ+08].

Definition 6.3.1. *For distinct $a, d \in K^\times$ and $c \in K^\times$ with $dc^4 \neq 1$, the* twisted Edwards curve *over K with coefficients c, d and twisted by a, is the curve given by the equation*

$$\mathcal{E}_{c,d}^a : ax^2 + y^2 = c^2(1 + dx^2y^2) .$$

An Edwards curve *is a twisted Edwards curve with $a = 1$.*

The excluded cases for a, c, d define singular curves. The equation in Definition 6.3.1 uses affine coordinates (x, y) and we will always work with the affine model. The projectivized model is $(aX^2 + Y^2)Z^2 = Z^4 + dX^2Y^2$ in projective coordinates $[X : Y : Z]$. In Lemma 6.3.2 we show certain isomorphisms between the different families of Edwards curves. These results are known, but we gather them for completeness here.

Lemma 6.3.2. *Let $a, c, d \in K^\times$ with $a \neq d$ and $dc^4 \neq 1$.*

(i) There is a K-isomorphism $\Phi_c : \mathcal{E}^a_{c,d} \to \mathcal{E}^a_{1,dc^4}$, $(x, y) \mapsto (x/c, y/c)$, with inverse $\Phi_c^{-1} : \mathcal{E}^a_{1,dc^4} \to \mathcal{E}^a_{c,d}$ sending (X, Y) to (cX, cY).

(ii) There is an isomorphism $\Psi_a : \mathcal{E}^a_{c,d} \to \mathcal{E}^1_{c,d/a}$, $(x, y) \mapsto (\sqrt{a}x, y)$ defined over $K(\sqrt{a})$.

(iii) Assume that $a \neq dc^4$. There is an isomorphism $\mathcal{E}^a_{c,d} \to \mathcal{E}^1_{1,dc^4/a}$, $(x, y) \mapsto (\sqrt{a}x/c, y/c)$, defined over $K(\sqrt{a})$.

Proof. (i) As $c \neq 0$, $ax^2 + y^2 = c^2(1 + dx^2y^2)$ can be rewritten as $a\left(\frac{x}{c}\right)^2 + \left(\frac{y}{c}\right)^2 = 1 + dc^4\left(\frac{x}{c}\right)^2\left(\frac{y}{c}\right)^2$, which after the change of variables $(X, Y) = (x/c, y/c)$ gives $aX^2 + Y^2 = 1 + dc^4X^2Y^2$, which is the curve $\mathcal{E}^a_{1,dc^4}$, as $dc^4 \neq 1$ by assumption.

(ii) We write the equation $ax^2+y^2 = c^2(1+dx^2y^2)$ as $(\sqrt{a}x)^2+y^2 = c^2(1+(d/a)(\sqrt{a}x)^2y^2)$ and $(X, Y) = (\sqrt{a}x, y)$ gives $X^2 + Y^2 = c^2(1 + (d/a)X^2Y^2)$, the curve $\mathcal{E}^1_{c,d/a}$ (as $d/a \neq 1$).

(iii) is obtained by composing Ψ_c and Φ_a. Note that $dc^4/a \neq 1$. □

These isomorphisms explain why the parameter c is often omitted in the literature. It is sufficient to replace d by dc^4 to obtain general results including the parameter c. In view of Lemma 6.3.2 we shall mainly restrict our study to twisted Edwards curves with $c = 1$ and we will write $\mathcal{E}_d$ for the curve $\mathcal{E}^1_{1,d}$ and $\mathcal{E}^a_d$ for the curve $\mathcal{E}^a_{1,d}$.

Let now $c = 1$. Let $\mathcal{E}^a_d(K)$ be the group of K-points on $\mathcal{E}^a_d$. The sum of $(x_1, y_1), (x_2, y_2) \in \mathcal{E}^a_d(K)$ is defined by (see [BBJ+08, Section 6])

$$(x_1, y_1) + (x_2, y_2) = \left(\frac{x_1y_2 + y_1x_2}{1 + dx_1x_2y_1y_2}, \frac{y_1y_2 - x_1x_2 - ax_1x_2}{1 - dx_1x_2y_1y_2}\right) .$$

The neutral element for the addition is $(0, 1)$ and the inverse of (x, y) is $(-x, y)$. In particular, $(0, -1)$ has order 2. If $a = 1$ then $(\pm 1, 0)$ have order 4. The addition law is *complete* if a is a square and d a non-square in K, meaning that the denominators in this formula are not zero, and hence the formula is defined for all pairs of points on $\mathcal{E}^a_d$ (for example, duplication or inversion). In [ALNR11, Section4] the authors give a geometric description of this group law.

6.3.2 Models for elliptic curves

An elliptic curve over K is a non-singular projective curve of genus 1 over K together with a K-rational point $\mathcal{O} \in E(K)$. An elliptic curve in the short Weierstrass model is given by an equation $y^2 = x^3+\alpha x+\beta$ where $\alpha, \beta \in K$ satisfy $4\alpha^3+27\beta^2 \neq 0$. Another well-known model for elliptic curves is the Montgomery model, introduced in [Mon87]. Its arithmetic has been studied independently in several works (see for instance [CS18] for a survey).

Definition 6.3.3. *For $A \in K \setminus \{\pm 2\}$, $B \in K^\times$, the* Montgomery curve *$\mathcal{M}_{A,B}$ defined over K with parameters A, B is given by the equation*

$$\mathcal{M}_{A,B} : Bv^2 = u^3 + Au^2 + u .$$

The cases $A^2 = 4, B = 0$ are excluded because they define singular curves.

We say that two algebraic varieties X, Y over K are *birationally equivalent over* K if there exist rational maps $f : X \to Y$ and $g : Y \to X$ defined over K such that $(f \circ g)(P) = P$ for all $P \in Y(\overline{K})$ such that $(f \circ g)(P)$ is defined, and $(g \circ f)(Q) = Q$ for all $Q \in X(\overline{K})$ such that $(g \circ f)(Q)$ is defined. The map f is called a *birational equivalence* from X to Y. It induces an isomorphism from a non-empty subset $U \subseteq X$ to a non-empty subset $V \subseteq Y$. The points where f is not defined are *exceptional points*.

Every twisted Edwards curve $\mathcal{E}_d^a$ over K is birationally equivalent over K to a Montgomery curve by [BBJ+08, Theorem 3.2]. Explicitly, the map

$$\begin{aligned} \mathcal{E}_d^a : ax^2 + y^2 = 1 + dx^2y^2 &\xrightarrow{\psi_{a,d}} \mathcal{M}_{A,B} : Bv^2 = u^3 + Au^2 + u\,, \\ (x, y) &\longmapsto (u, v) = \left(\frac{1+y}{1-y}, \frac{1+y}{(1-y)x}\right) \end{aligned} \tag{6.1}$$

with $(A, B) := (A(a, d), B(a, d)) = \left(2\frac{a+d}{a-d}, \frac{4}{a-d}\right)$ is a birational equivalence. Moreover, there is a birational equivalence between $\mathcal{E}_d^a$ and the curve in medium Weierstrass form

$$\begin{aligned} \mathcal{E}_d^a : ax^2 + y^2 = 1 + dx^2y^2 &\longrightarrow E_{a,d} : y^2 = x^3 + 2(a+d)x^2 + (a-d)^2x\,, \\ (x, y) &\longmapsto \left((a-d)\frac{1+y}{1-y}, (a-d)\frac{2(1+y)}{(1-y)x}\right)\,. \end{aligned} \tag{6.2}$$

Note that the map is undefined at the points $(x, y) \in \mathcal{E}_d^a(K)$ with $x = 0$ or $y = 1$; defining the *exceptional points* for the birational equivalence.

One can interpret elliptic curves with given torsion data from the viewpoint of moduli spaces. Let $K \subseteq \mathbb{C}$ be a subfield. The congruence subgroup of $\mathrm{SL}_2(\mathbb{Z})$ of level 4 is defined by

$$\Gamma_1(4) = \gamma = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z}) : \gamma \equiv \begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix} \pmod 4 \right\}\,.$$

It acts on the complex upper half-plane $\mathfrak{h} = \{\tau \in \mathbb{C} : \Im(\tau) > 0\}$ via the usual fractional linear transformations. The *modular curve* $Y_1(4)$ is defined as $\Gamma_1(4) \setminus \mathfrak{h} = \{\Gamma_1(4)\tau : \tau \in \mathfrak{h}\}$, which is non-compact. It can be compactified by considering the action on the extended upper half-plane $\mathfrak{h}^* = \mathfrak{h} \cup \mathbb{P}^1(\mathbb{Q})$: there exists a smooth projective curve $X_1(4)/\mathbb{Q}$ (the *(compactified) modular curve of level* 4) over $\mathbb{Q}$ and an isomorphism $\Gamma_1(4) \setminus \mathfrak{h}^* \simeq X_1(4)(\mathbb{C})$. There is an interpretation as moduli spaces for elliptic curves with points of order 4: there is a bijection

$$Y_1(4)(K) \leftrightarrow \{(E, P) : E/K \text{ elliptic curve}, P \in E(K) \text{ point of order } 4\}/\simeq\,,$$

that is, an equivalence class $\Gamma_1(4)\tau$ in $Y_1(4)$ corresponds to a pair (E, P), where E is an elliptic curve over K defined up to isomorphism and P is a point of exact order 4 in $E(K)$ preserved by the isomorphism. We define the projective d-line over K by $\mathbb{P}^1(K) \setminus \{\infty, 0, 1\}$, and over every point on the d-line we consider the elliptic curve $E_d := E_{1,d}$ defined in (6.2). Let $P_4 \in E_d(K)$ be one of the points of order 4 in $E_d(K)$. The Weierstrass equation $y^2 = x^3 + 2(1 + d)x^2 + (1 - d)^2x$ of E_d can be seen as a surface over the d-line with fibers given by the curves $\{E_d\}_d$. The singular fibers are E_0 and E_1 lying over $d = 0$ and $d = 1$. The points on $Y_1(4)(K)$ are pairs (E_d, P_4) up to isomorphism.

6.3.3 Isomorphism classes of Edwards curves

Let E be an elliptic curve defined over K. We denote by $\mathrm{Edw}_{\overline{K}}(E)$ the set of Edwards curves $\mathcal{E}_d$ defined over $\overline{K}$ (i.e. with $d \in \overline{K}$) such that there exists a birational equivalence over $\overline{K}$ between $\mathcal{E}_d$ and E; note that any birational equivalence over $\overline{K}$ becomes an isomorphism over $\overline{K}$. For a subfield $L \subseteq \overline{K}$ we also denote by $\mathrm{Edw}_L(E)$ the subset of $\mathrm{Edw}_{\overline{K}}(E)$ of Edwards curves defined over L and birationally equivalent to E over a finite extension L'/L. We treat two curves $\mathcal{E}_d$ and $\mathcal{E}_{d'}$ as distinct whenever $d \neq d'$.

The set $\mathrm{Edw}_{\overline{K}}(E)$ defines E only up to isomorphism: if $f : E \to E'$ is an isomorphism, then any birational equivalence $\mathcal{E}_d \to E$ induces a birational equivalence $\mathcal{E}_d \to E'$ after composition with f. Therefore the following lemma is clear.

Lemma 6.3.4. *Let E,E' be elliptic curves defined over K.*

(i) If E and E' are $\overline{K}$-isomorphic, then $\mathrm{Edw}_{\overline{K}}(E) = \mathrm{Edw}_{\overline{K}}(E')$.

(ii) If E and E' are not $\overline{K}$-isomorphic, then $\mathrm{Edw}_{\overline{K}}(E) \cap \mathrm{Edw}_{\overline{K}}(E') = \emptyset$.

The j-invariant defines a $\overline{K}$-isomorphism class of elliptic curves. For $j \in K$, define:

$$\mathrm{Edw}_{\overline{K}}(j) = \{\mathcal{E}_d \mid d \in \overline{K},\ \exists\, E/K \text{ with } j(E) = j \text{ and } \exists\, \overline{K}\text{-birational equivalence } \mathcal{E}_d \to E\}\,, \tag{6.3}$$

and, similarly, for the set $\mathrm{Edw}_L(j)$ for a subfield $L \subseteq \overline{K}$. We shall in the sequel also use the notation $\mathfrak{Ell}(j)$ for the set of $\overline{K}$-isomorphism classes of elliptic curves over K with j-invariant j. We denote by $[E] \in \mathfrak{Ell}(j)$ the equivalence class of E.

The curves in $\mathrm{Edw}_{\overline{K}}(j)$ are all birationally equivalent over $\overline{K}$, and thus isomorphic. For distinct curves $\mathcal{E}_d$ and $\mathcal{E}_{d'}$ in $\mathrm{Edw}_{\overline{K}}(j)$, letting ϕ_d denote a birational equivalence $\mathcal{E}_d \to E$ (with $[E] \in \mathfrak{Ell}(j)$) and $\phi_{d'}$ a birational equivalence $\mathcal{E}_{d'} \to E$, then $\phi_{d'}^{-1}\phi_d$ is a birational equivalence $\mathcal{E}_d \to \mathcal{E}_{d'}$.

The following result of Ahmadi and Granger [AG12] gives explicit relations between isomorphic Edwards curves.

Proposition 6.3.5 ([AG12], Proposition 4.2)**.** *Let $\mathcal{E}_d$ and $\mathcal{E}_{d'}$ be two Edwards curves defined over K. Then $\mathcal{E}_d$ is isomorphic to $\mathcal{E}_{d'}$ over $\overline{K}$ if and only if*

$$d' \in \Sigma(d) := \left\{ d, \frac{1}{d}, \left(\frac{1 \pm d^{1/4}}{1 \mp d^{1/4}}\right)^4, \left(\frac{1 \pm \sqrt{-1}d^{1/4}}{1 \mp \sqrt{-1}d^{1/4}}\right)^4 \right\}. \tag{6.4}$$

As remarked in [AG12], this result also follows from Edwards' original paper [Edw07, Proposition 6.1]. Proposition 6.3.5 implies that $\mathrm{Edw}_{\overline{K}}(j)$ is of cardinality 6, and of smaller cardinality if and only if some elements in $\Sigma(d)$ coincide. The proof technique of Proposition 6.3.5 in [AG12] is to establish isomorphisms between Edwards curves and Legendre curves. In what follows we shall establish a stronger version of this result.

Remark 6.3.6. If $\mathrm{Edw}_{\overline{K}}(E)$ contains $\mathcal{E}_d$ defined over K with $d = \ell^4$ a fourth power in K, then $\mathrm{Edw}_{\overline{K}}(E)$ contains four curves defined over K, defined by the fourth powers $1/d = (1/\ell)^4$ and $\left(\frac{1\pm\ell}{1\mp\ell}\right)^4$. Remark also that by Proposition 6.3.5, the curves in $\mathrm{Edw}_{\overline{K}}(E)$ cannot all be defined over the base field K unless -1 is a square in K.

6.4 Abstract constructions of Edwards and twisted Edwards curves

We prove an abstract construction of the Edwards model from the Riemann-Roch Theorem. Then, we use a classical cohomological approach to obtain the twisted Edwards model. Let K be a perfect field of characteristic different from 2.

6.4.1 Universal construction of the Edwards model

We prove a universal construction of the Edwards model for elliptic curves with a rational point of order 4, from the theorem of Riemann-Roch (see [Sil09, Chapter II, §5, Theorem 5.4]) only, without relying on any model (such as the Weierstrass model) for elliptic curves. Along the same concept, the Weierstrass model (see [Sil09, Chapter III, §3, Proposition 3.1]), and also Huff's model (see [JTV10, Theorem 2]) can be derived from the Riemann-Roch Theorem only. An elliptic curve is a pair $(E, \mathcal{O})$ where E is a non-singular projective curve of genus 1 over K, and $\mathcal{O} \in E(K)$ is a rational point. Let $\mathrm{Div}_K(E)$ be the group of K-rational divisors on E. For $D \in \mathrm{Div}_K(E)$, let

$$\mathcal{L}(D) = \{f \in K(E)^\times : \mathrm{div}(f) + D \geq 0\} \cup \{0\}$$

be the Riemann-Roch space of D, a finite-dimensional K-vector space of dimension $\ell(D)$. If $D = \sum_{P \in E(K)} a_P(P)$ with $a_P \in \mathbb{Z}$, we write $\deg(D) = \sum_P a_P \in \mathbb{Z}$ for its degree and $\sum(D) = \sum_P a_P P \in E(K)$ for its sum.

Theorem 6.4.1. *Let $(E, \mathcal{O})$ be an elliptic curve defined over K. Assume that there exists $P \in E(K)$ of order 4. Define the divisors*

$$D_1 := 2(\mathcal{O}) - 2(2P)\ ,\ D_2 := 2(\mathcal{O}) - (P) - (3P)\ \in \mathrm{Div}_K(E)$$

(i) The Riemann-Roch spaces $\mathcal{L}(D_1)$ and $\mathcal{L}(D_2)$ are one-dimensional over K.
(ii) Let $x, y \in K(E)^\times$ be generators of $\mathcal{L}(D_1)$ and $\mathcal{L}(D_2)$, respectively. The functions x^2y^2 and $x^2 + y^2 - 1$ in $K(E)^\times$ are proper multiples of each other, i.e. there exists $d \in K \setminus \{0, 1\}$ such that $x^2y^2 = 1 + dx^2y^2$.
(iii) The map

$$\phi : E(K) \to \mathbb{P}^2(K), Q \mapsto [x(Q) : y(Q) : 1]$$

gives a birational equivalence from E to the Edwards curve $\mathcal{E}_d$ defined over K. In particular, x, y from part (ii) are coordinate functions for the Edwards model of $(E, \mathcal{O})$.

Proof. (i) The divisors D_1 and D_2 satisfy $\deg(D_1) = \deg(D_2) = 0$ and $\sum(D_1) = \sum(D_2) = \mathcal{O}$ since P has order 4. We use the following consequence of the Riemann-Roch Theorem: if D is a divisor of degree 0 then $\ell(D) \in \{0, 1\}$, and more precisely, we have for $D = \sum_Q a_Q(Q)$:

$$\ell(D) = \begin{cases} 1 & \text{if } \sum_Q a_Q Q = \mathcal{O} \\ 0 & \text{if } \sum_Q a_Q Q \neq \mathcal{O}\ , \end{cases} \tag{6.5}$$

where $\mathcal{O}$ here is the rational neutral element of E ([Was03, Proposition 11.14]). It follows that the Riemann-Roch spaces $\mathcal{L}(D_1)$ and $\mathcal{L}(D_2)$ are 1-dimensional over K.

(ii) Define $D := 2D_1 + 2D_2$. By *(i)*, there exist non-zero rational functions x, y on E such that $\mathcal{L}(D_1) = \langle x \rangle$ and $\mathcal{L}(D_2) = \langle y \rangle$ and we may assume the divisor of x is exactly $\mathrm{div}(x) = -D_1 = 2(2P) - 2(\mathcal{O})$ and the divisor of y is exactly $\mathrm{div}(y) = -D_2 = (P) + (3P) - 2(\mathcal{O})$. Therefore, x is non-zero and well-defined at P and y is non-zero and well-defined at $2P$. Therefore, we may consider x and y to be normalized such that $x(P) = y(2P) = 1$.

Let f, g be the rational functions $f = x^2y^2$ and $g = x^2 + y^2 - 1$. We show that these functions are K-multiples of each other. Clearly, we have

$$\mathrm{div}(f) = 2\mathrm{div}(x) + 2\mathrm{div}(y) = -2D_1 - 2D_2 = -D$$

so that $f \in \mathcal{L}(D)$. Since $D = 8(\mathcal{O}) - 2(P) - 4(2P) - 2(3P)$ has degree zero and sums to $\mathcal{O}$, we have $\ell(D) = 1$, by Equation (6.5).

Let us now consider the function $g = x^2+y^2-1$. We establish below that $\mathrm{div}(g) = -D$. The poles of g are those of x and y counted with double multiplicity, that is, $\mathcal{O}$ with multiplicity 8. Therefore there exists a divisor Z_g on E such that $\mathrm{div}(g) = Z_g - 8(\mathcal{O})$ where $Z_g = \sum_Q a_Q(Q)$ and the sum runs over the zeros of g; it remains to show that $Z_g = 2(P) + 4(2P) + 2(3P)$. First, we see that $P, 2P$ and $3P$ are indeed zeroes of g. Namely, by our choices of x and y, we have that $x(P) = 1$ and $y(2P) = 1$, implying that $(x^2 - 1)(P) = 0$ and $(y^2 - 1)(2P) = 0$. Then, it follows from the divisors of x and y that $g(P) = (x^2-1)(P)+y^2(P) = 0$ and $g(2P) = x^2(2P)+(y^2-1)(2P) = 0$. To see that g is zero at $3P$, we note already that $y^2(3P) = 0$, thus it is sufficient to show that x^2-1 is zero at $3P$. Consider the endomorphism $\rho : K(E)^\times \to K(E)^\times$ sending a function h to the function $\rho(h)$ defined by $\rho(h)(Q) = h(-Q)$ for every $Q \in E(K)$. Since $\mathcal{O} = -\mathcal{O}$ and $2P = -2P$ (P has order 4), ρ is an endomorphism of the vector space $\mathcal{L}(D_1)$, containing x. Since $\mathcal{L}(D_1)$ is of dimension 1 and ρ^2 is the identity on $\mathcal{L}(D_1)$, ρ is either the identity or minus the identity on $\mathcal{L}(D_1)$. We distinguish both cases and compute $x(3P)$ accordingly. If ρ is the identity then $\rho(x) = x$ and in particular, $x(3P) = \rho(x)(P) = x(P) = 1$. If ρ is minus the identity then $x(3P) = \rho(x)(P) = -x(P) = -1$. In both cases, we have that $(x^2 - 1)(3P)$ equals 0, whence $g(3P) = 0$.

Therefore there exist integers $a_1, a_2, a_3 \geq 1$ such that $Z_g = a_1(P) + a_2(2P) + a_3(3P)$. We now prove that $(a_1, a_2, a_3) = (2, 4, 2)$.

It is enough to prove that $a_1, a_3 \geq 2$ and $a_3 \geq 4$. In that case we have that $Z_g \geq 2(P) + 4(2P) + 2(3P)$ and $\mathrm{div}(g) = Z_g - 8(\mathcal{O}) \geq 2(P) + 4(2P) + 2(3P) - 8(\mathcal{O})$. By principality of $\mathrm{div}(g)$ we then obtain the desired equality $\mathrm{div}(g) = 2(P) + 4(2P) + 2(3P) - 8(\mathcal{O}) = -D$.

Let $b_1 := \mathrm{ord}_P(x^2 - 1)$, $b_3 := \mathrm{ord}_{3P}(x^2 - 1)$ and $b_2 := \mathrm{ord}_{2P}(y^2 - 1)$. By the choice of x and y we have that $b_i \geq 1$ for all $i \in \{1, 2, 3\}$. We shall prove that $b_2 = 4$ and $b_1 = b_3 = 2$.

To see that $b_2 = 4$ we write the divisor of $y^2 - 1$ as $Z_{y^2-1} - 4(\mathcal{O})$ for some divisor $Z_{y^2-1} \geq 0$. ($y^2 - 1$ has the same poles as y with double multiplicity). We clearly see from the divisor of y, that $y^2 - 1$ is non-zero at the points $P, 3P$ and $\mathcal{O}$; thus, $Z_{y^2-1} = b_2(2P)$ and $\mathrm{div}(y^2 - 1) = b_2(2P) - 4(\mathcal{O})$. By principality, we deduce $b_2 = 4$.

To see that $b_1 = b_3 = 2$ we consider the divisor of $x^2 - 1$. From the divisor of x, we can write $\mathrm{div}(x^2 - 1) = Z_{x^2-1} - 4(\mathcal{O})$ for some divisor $Z_{x^2-1} \geq 0$. Since $x(2P) = 0$, $x^2 - 1$ does not vanish at $2P$ and we have $Z_{x^2-1} = b_1(P)+b_3(3P)$ and thus $\mathrm{div}(x^2-1) = b_1(P)+b_3(3P)-4(\mathcal{O})$. By principality, the integers $b_1, b_3 \geq 1$ must satisfy $b_1 + b_3 = 4$ (so that the divisor is of degree 0) and $b_1 + 3b_3 \equiv 0 \pmod 4$ (so that the divisor sums to $\mathcal{O}$). By the first condition the only possible solutions are $(b_1, b_3) \in \{(1, 3), (3, 1), (2, 2)\}$, but the first two solutions are excluded by the second condition. Consequently, $b_1 = b_3 = 2$. Putting all together, we obtain, as

desired:

$$\begin{aligned} a_1 = \mathrm{ord}_P(g) &\geq \min\{2\mathrm{ord}_P(y), \mathrm{ord}_P(x^2-1)\} = \min\{2, b_1\} = 2\,, \\ a_2 = \mathrm{ord}_{2P}(g) &\geq \min\{2\mathrm{ord}_{2P}(x), \mathrm{ord}_{2P}(y^2-1)\} = \min\{4, b_2\} = 4\,, \\ a_3 = \mathrm{ord}_{3P}(g) &\geq \min\{2\mathrm{ord}_{3P}(y), \mathrm{ord}_{3P}(x^2-1)\} = \min\{2, b_3\} = 2\,. \end{aligned}$$

This establishes that $\mathrm{div}(f) = \mathrm{div}(g) = -D$. Therefore there exists a constant $d \in K^\times$ such that $f = dg$. The case $d = 1$ amounts to the relation $(x^2-1)(1-y^2) = 0$ in the function field of E, thus for every Q, Q is either a zero of x^2-1 or a zero of $1-y^2$. Taking $Q = \mathcal{O}$ gives a contradiction; thus $d \neq 1$.

(iii) We show that $\phi : E(K) \to \mathbb{P}^2(K)$ sending a point Q in $E(K)$ to $[x(Q) : y(Q) : 1]$ extends to a morphism $E(K) \to \mathcal{E}_d(K) \subseteq \mathbb{P}^2(K)$, by showing that ϕ is a morphism of degree 1 between smooth curves ([Sil09, Chapter II, §2, Corollary 2.4.1]). Since E is smooth and the image of ϕ is contained in $\mathcal{E}_d(K)$ by (ii), ϕ extends to a morphism $E \to \mathcal{E}_d$ ([Sil09, Chapter II, §2, Proposition 2.1]). Since $d \neq 0, 1$, $\mathcal{E}_d$ is smooth. The map ϕ is non-constant, thus surjective, and induces an injection $\phi^* : K(\mathcal{E}_d) \to K(E), h \mapsto h \circ \phi$ on the function fields. The degree $[K(E) : \phi^*K(\mathcal{E}_d)]$ of ϕ is 1. Otherwise two distinct points $Q \neq Q' \in E(K)$ would be mapped to the same point $[x(Q') : y(Q') : 1] = [x(Q) : y(Q) : 1]$ in $\mathcal{E}_d(K)$ under ϕ; thus $x - x(Q)$ has zeroes at Q and Q' and a double pole at $\mathcal{O}$ (as $\mathrm{div}(x) = 2(2P) - 2(\mathcal{O})$), so that $\mathrm{div}(x - x(Q)) = (Q) + (Q') - 2(\mathcal{O})$, and similarly, one establishes $\mathrm{div}(y - y(Q)) = (Q) + (Q') - 2(\mathcal{O})$. Since these divisors coincide, there exists $\nu \in K^\times$ such that $x - x(Q) = \nu(y - y(Q))$, or equivalently, $x - \nu y = x(Q) - \nu y(Q)$, showing that $x - \nu y$ is constant, which gives a contradiction with the divisors of x and y. □

6.4.2 Twisted Edwards curves from Galois cohomology

Let $d \in K \setminus \{0, 1\}$ and $\mathcal{E}_d$ be an Edwards curve defined over K. The parameter a in Definition 6.3.1 is referred to as the *twist* parameter for the curve. By Lemma 6.3.2 (ii) and as noted in [BBJ+08, Section 2], $\mathcal{E}_d^a$ is a quadratic twist of the curve $\mathcal{E}^1_{d/a} : x^2 + y^2 = 1 + (d/a)x^2y^2$; an isomorphism between these curves is defined over the quadratic extension $K(\sqrt{a})$. If a is a square in K then the curves are isomorphic over K, therefore we assume that a is not a square. Twisted Edwards curves were also studied in [HWCD08].

We here derive the quadratic twist $\mathcal{E}_d^a$ of $\mathcal{E}_{d/a}$ more abstractly from Galois cohomology. In particular, this theory justifies why the quadratic twist equation for $\mathcal{E}_{d/a}$ is given by $\mathcal{E}_d^a$ (up to isomorphism). Let us first recall the connection between twists of curves and Galois cohomology, mainly following [Sil09, Chapter X].

Let $\mathcal{C}/K$ be a smooth projective curve defined over K. We say that a smooth projective curve $\mathcal{C}'/K$ is a *twist* of $\mathcal{C}/K$ if there is a $\overline{K}$-isomorphism $\psi : \mathcal{C}' \to \mathcal{C}$. Two twists $\mathcal{C}'$, $\mathcal{C}''$ are equivalent (we write $\sim$ for the equivalence) if there is a K-isomorphism $\mathcal{C}' \to \mathcal{C}''$. We denote the set of twists of $\mathcal{C}$ modulo K-isomorphism by

$$\mathrm{Tw}_K(\mathcal{C}) = \{\psi : \mathcal{C}' \to \mathcal{C}\ \overline{K}\text{-isomorphism}\}/\sim .$$

We further say that $\mathcal{C}'/K$ together with a $\overline{K}$-isomorphism $\psi : \mathcal{C}' \to \mathcal{C}$ is a twist of degree $\delta \in \mathbb{N}_{\geq 1}$ of $\mathcal{C}/K$ if δ is the degree of the smallest extension field of K over which ψ is defined.

The well-known relation between the set $\mathrm{Tw}(\mathcal{C}/K)$ of twists of $\mathcal{C}/K$ and a certain cohomology set is recalled in Theorem 6.4.2. We denote by $\mathrm{Aut}(\mathcal{C})$ the automorphism group of $\mathcal{C}$, that is $\overline{K}$-isomorphisms $\mathcal{C}(K) \to \mathcal{C}(K)$, under composition of maps. The absolute Galois group $G_K := \mathrm{Gal}(\overline{K}/K)$ of K acts on $\mathrm{Aut}(\mathcal{C})$ via conjugation, that is, for $\sigma \in G_K$ and $\psi \in \mathrm{Aut}(\mathcal{C})$, $\psi^\sigma \psi^{-1} \in \mathrm{Aut}(\mathcal{C})$. Here, for a $\overline{K}$-isomorphism $\psi : \mathcal{C}' \to \mathcal{C}$, ψ^σ is the automorphism obtained by acting with $\sigma \in G_K$ as $\psi^\sigma : \mathcal{C}' \to \mathcal{C}$ (via $\psi^\sigma(x,y) = \psi(\sigma(x), \sigma(y))$). One shows that the map $\sigma \mapsto \psi^\sigma \psi^{-1}$ is a 1-cocycle with values in $\mathrm{Aut}(\mathcal{C})$, i.e. an element of $H^1(G_K, \mathrm{Aut}(\mathcal{C}))$.

Theorem 6.4.2 (Theorem 2.2, Chapter X, §2 [Sil09])**.** *There is a bijection*

$$\mathrm{Tw}_K(\mathcal{C}) \to H^1(G_K, \mathrm{Aut}(\mathcal{C})) \, .$$

The correspondence is as follows: for $\mathcal{C}'/K \in \mathrm{Tw}_K(\mathcal{C})$ together with a $\overline{K}$-isomorphism $\psi : \mathcal{C}' \to \mathcal{C}$, the corresponding element in $H^1(G_K, \mathrm{Aut}(\mathcal{C}))$ is the class $[\xi]$ of the 1-cocycle $\xi^{(\psi)} : G_K \to \mathrm{Aut}(\mathcal{C})$, $\sigma \mapsto \xi_\sigma^{(\psi)} = \psi^\sigma \psi^{-1}$, determined by the K-isomorphism class of $\mathcal{C}$ and independent of the chosen isomorphism ψ.

Conversely, for $\xi \in H^1(G_K, \mathrm{Aut}(\mathcal{C}))$, a curve $\mathcal{C}'/K$ together with a $\overline{K}$-isomorphism $\psi : \mathcal{C}' \to \mathcal{C}$ such that $\xi_\sigma = \psi^\sigma \psi^{-1}$ for every $\sigma \in G_K$, is constructed as follows. Denote by $\overline{K}(\mathcal{C})_\xi$ a field which is isomorphic to $\overline{K}(\mathcal{C})$, via an abstract isomorphism which we denote by $\eta : \overline{K}(C) \to \overline{K}(\mathcal{C})_\xi$. Now we extend the Galois action on $\overline{K}(\mathcal{C})$ to one on $\overline{K}(\mathcal{C})_\xi$, by setting $\eta(f)^\sigma = \eta(f^\sigma \xi_\sigma)$ for every $f \in \overline{K}(\mathcal{C})$ viewed as a map $\mathcal{C} \to \mathbb{P}^1$, and every $\sigma \in G_K$. This Galois action on $\overline{K}(\mathcal{C})_\xi$ is said "twisted" by ξ. The subfield $\overline{K}(\mathcal{C})_\xi^{G_K} \subset \overline{K}(\mathcal{C})_\xi$ of elements fixed under this action of G_K is then the function field of $\mathcal{C}'$. Finally, an isomorphism $\mathcal{C}' \to \mathcal{C}$ comes from the isomorphism $\overline{K}(\mathcal{C}') \simeq \overline{K}(\mathcal{C})$.

Specializing to Edwards curves, we derive the quadratic twist model for $\mathcal{E}_{d/a}$.

Theorem 6.4.3. *Let $d \in K \setminus \{0,1\}$ and $a \in K$ not a square. Let $\chi_a : G_K \to \{\pm 1\}$, mapping $\sigma \mapsto \sigma(\sqrt{a})/\sqrt{a}$ be the quadratic character associated to $K(\sqrt{a})$.*
Define $\xi : G_K \to \mathrm{Aut}(\mathcal{E}_{d/a})$, sending $\sigma \mapsto \chi_a(\sigma)$, where $1 \in \mathrm{Aut}(\mathcal{E}_{d/a})$ is the identity and $-1 \in \mathrm{Aut}(\mathcal{E}_{d/a})$ is the negation. Then, a representative for the equivalence class of twists in $\mathrm{Tw}_K(\mathcal{E}_{d/a})$ associated to $[\xi] \in H^1(G_K, \mathrm{Aut}(\mathcal{E}_{d/a}))$, is given by the curve

$$\mathcal{E}_d^a : aX^2 + Y^2 = 1 + dX^2Y^2 \, .$$

Proof. Let $a \in K$ and $\xi : \sigma \mapsto \chi_a(\sigma)$ be as in the statement. Let $\mathcal{C}/K$ be the twist of $\mathcal{E}_{d/a}$, defined according to the bijection of Theorem 6.4.2, realizing the 1-cocycle ξ, that is, if $\psi : \mathcal{C} \to \mathcal{E}_{d/a}$ is an isomorphism defined over $K(\sqrt{a})$, then $\psi^\sigma \psi^{-1} = \xi_\sigma$ for every $\sigma \in G_K$.

We now show that an equation for $\mathcal{C}$ is given by $aX^2 + Y^2 = 1 + dX^2Y^2$. For $(x,y) \in \mathcal{E}_{d/a}(K)$, we have $x^2 + y^2 = 1 + (d/a)x^2y^2$. Using that $-(x,y) = (-x,y)$ in $\mathcal{E}_{d/a}(K)$, the action of G_K on $P = (x,y)$ is $P^\sigma = (\chi_a(\sigma)x, y)$. It is easy to see that the functions $X = x/\sqrt{a}$ and $Y = y$ are fixed by G_K in $\overline{K}(\mathcal{C})_\xi$, i.e. $\sigma(X) = X, \sigma(Y) = Y$ for every $\sigma \in G_K$, and satisfy the relation

$$aX^2 + Y^2 = x^2 + y^2 = 1 + (d/a)x^2y^2 = 1 + dX^2Y^2$$

in the function field of $\mathcal{C}$, which gives for $\mathcal{C}$ the equation $aX^2 + Y^2 = 1 + dX^2Y^2$. Therefore, $\mathcal{E}_d^a$ is a quadratic twist of $\mathcal{E}_{d/a}$ by the quadratic character χ_a associated to $K(\sqrt{a})$. □

6.5 Algebraic construction of Edwards curves

Let $j \in K$ be a j-invariant defining a $\overline{K}$-isomorphism class of elliptic curves. In this section, we give an algebraic description of $\mathrm{Edw}_{\overline{K}}(j)$.

6.5.1 Edwards covering of the j-line

As noticed in [BBJ+08], the elliptic curve $E_{a,d}$ from Equation (6.2), birationally equivalent to $\mathcal{E}_d^a$, has j-invariant

$$j(E_{a,d}) = \frac{16(a^2 + 14ad + d^2)^3}{ad(a-d)^4} \,. \tag{6.6}$$

We consider $a = 1$ and write j_d for $j(E_{1,d})$. It follows that $\mathcal{E}_d \in \mathrm{Edw}_{\overline{K}}(j)$ if and only if d satisfies (6.6), or equivalently, d is a root of the polynomial

$$\gamma_j(T) := 16(1 + 14T + T^2)^3 - jT(1-T)^4 \in K[T] \,.$$

In particular, every d with $\mathcal{E}_d \in \mathrm{Edw}_{\overline{K}}(j)$ is algebraic over K. Moreover, there are 6 roots of γ_j over $\overline{K}$ counted with multiplicity and thus $\mathrm{Edw}_{\overline{K}}(j)$ contains at most 6 curves (which is, of course, necessary by Proposition 6.3.5). Our algebraic realization of $\mathrm{Edw}_{\overline{K}}(j)$ is therefore $\mathrm{Edw}_{\overline{K}}(j) = \{\mathcal{E}_d/\overline{K} \mid \gamma_j(d) = 0\}$.

Let us now study properties of γ_j. By expanding, we write $\gamma_j(T)$ as $a_0T^6 + a_1(j)T^5 + a_2(j)T^4 + a_3(j)T^3 + a_2(j)T^2 + a_1(j)T + a_0$ with coefficients in $K[j]$ given by

$$a_0 = 16 \,,\; a_1(j) = 672 - j \,,\; a_2(j) = 9456 + 4j \,,\; a_3(j) = 45248 - 6j \tag{6.7}$$

The polynomial γ_j is self-reciprocal, i.e.:

$$\gamma_j(T) = T^6\gamma_j\left(\frac{1}{T}\right) \,,\; T \neq 0 \tag{6.8}$$

Consequently, if $\gamma_j(d) = 0$ (then $d \neq 0$), then $\gamma_j(1/d) = 0$ and both are zeroes with the same multiplicity. Thus, if $\mathrm{Edw}_{\overline{K}}(j)$ contains $\mathcal{E}_d$ then it also contains $\mathcal{E}_{1/d}$ and both curves are defined over the same field.

Example 6.5.1. Consider the special j-invariants $j \in \{0, 1728\}$. Edwards curves associated to these j-invariants are studied for example in [LWZ13]. We have $\gamma_0(T) = 16(1+14T+T^2)^3$, of which the unique root is $T = -7 \pm 4\sqrt{3}$ with multiplicity 3. Therefore $\mathrm{Edw}_{\overline{K}}(0) = \{\mathcal{E}_{-7\pm4\sqrt{3}}\}$ and both curves are defined over $K(\sqrt{3})$. For $j = 1728$, solutions to $\gamma_{1728}(T) = 0$ are $\{-1, 17 \pm 12\sqrt{2}\}$ defined over K and $K(\sqrt{2})$, respectively. Thus, $\mathrm{Edw}_{\overline{K}}(1728) = \{\mathcal{E}_{-1}, \mathcal{E}_{17\pm12\sqrt{2}}\}$.

In view of Example 6.5.1, the map

$$\overline{K} \setminus \{0,1\} \to \overline{K} \,,\; d \mapsto j(E_d) = \frac{16(1 + 14d + d^2)^3}{d(1-d)^4} \tag{6.9}$$

is surjective and exactly three-to-one above $j(E_d) = 1728$ and two-to-one above $j(E_d) = 0$. If $j(E_d) \neq 0, 1728$, the map is six-to-one. If $\mathrm{char}(K) = 3$, then $1728 = 0$ and the map is one-to-one. In other terms the map (6.9) describes a six-fold cover of the (projective) j-line

$\mathbb{P}^1(K) \setminus \{\infty\}$ by the d-line, with ramifications of index 2 at $j = 0$ and of index 3 at $j = 1728$. We refer to this as the Edwards covering of the j-line.

The following proposition now recovers Proposition 6.3.5 by [AG12] from a purely algebraic approach.

Proposition 6.5.2. *Let $j \in K$ be a j-invariant, and $d \in \overline{K}$ such that $\mathcal{E}_d \in \mathrm{Edw}_{\overline{K}}(j)$. Then*

$$\mathrm{Edw}_{\overline{K}}(j) = \{\mathcal{E}_{d'} \mid d' \in \Sigma(d)\} ,$$

where $\Sigma(d)$ is the set defined in Equation (6.4)*.*

Proof. Let $j \in K$ and $d \in \overline{K}$; since $\mathcal{E}_d \in \mathrm{Edw}_{\overline{K}}(j)$, we must have $\gamma_j(d) = 0$. Let

$$\varphi_\epsilon(d) = \left(\frac{1 + \epsilon d^{1/4}}{1 - \epsilon d^{1/4}}\right)^4 , \ \epsilon \in \{1, \sqrt{-1}\}$$

so that $\Sigma(d) = \{d, 1/d, \varphi_1(d), \varphi_1(d)^{-1}, \varphi_{\sqrt{-1}}(d), \varphi_{\sqrt{-1}}(d)^{-1}\}$. It remains to show that $\varphi_1(d)$ and $\varphi_{\sqrt{-1}}(d)$ are zeroes of γ_j. These are direct computations. By (6.8), also $d^{-1}, \varphi_1(d)^{-1}$ and $\varphi_{\sqrt{-1}}(d)^{-1}$ are zeroes of γ_j. Since γ_j has degree 6, these are all its zeroes. □

6.5.2 Computing the roots of γ_j

We now address the problem of computing the roots of γ_j explicitly. By the self-reciprocity of γ_j, its roots come in inverse pairs $\{d_i, 1/d_i\}$ for $i \in \{1, 2, 3\}$. For $i \in \{1, 2, 3\}$, define

$$\alpha_i := d_i + 1/d_i \in \overline{K} .$$

Lemma 6.5.3. *Let $j \in K$ and write $\mathrm{Edw}_{\overline{K}}(j) = \bigcup_{i\in\{1,2,3\}}\{\mathcal{E}_{d_i}, \mathcal{E}_{1/d_i}\}$. Then $\alpha_i := d_i + 1/d_i$ for $i \in \{1, 2, 3\}$ are the three zeroes over $\overline{K}$ of the polynomial*

$$f_j(T) := T^3 - (j/16 - 42)T^2 + (j/4 + 588)T - j/4 + 2744 . \tag{6.10}$$

Moreover, f_j is separable if and only if $j \notin \{0, 1728\}$. If in addition f_j is irreducible over $K[T]$, then a splitting field of f_j is of degree 3 over K if $j - 1728$ is a square in K and of degree 6 over K if $j - 1728$ is a non-square in K.

Proof. We write $(1/16)\gamma_j(T) = \prod_{i=1}^3 (T - d_i)(T - 1/d_i) = \prod_{i=1}^3 (T^2 - \alpha_i T + 1)$. This gives

$$\frac{1}{16}\gamma_j(T) = T^6 - \sigma_1 T^5 + (\sigma_2 + 3)T^4 - (2\sigma_1 + \sigma_3)T^3 + (\sigma_2 + 3)T^2 - \sigma_1 T + 1 ,$$

where $\sigma_1 = \alpha_1 + \alpha_2 + \alpha_3$, $\sigma_2 = \alpha_1\alpha_2 + \alpha_1\alpha_3 + \alpha_2\alpha_3$, and $\sigma_3 = \alpha_1\alpha_2\alpha_3$ are the elementary symmetric polynomials in the variables $\alpha_1, \alpha_2, \alpha_3$. Comparing coefficients with (6.7) implies

$$\sigma_1 = \frac{-a_1(j)}{16} = \frac{j}{16} - 42 , \ \sigma_2 = \frac{a_2(j)}{16} - 3 = \frac{j}{4} + 588 ,$$
$$\sigma_3 = \frac{2a_1(j) - a_3(j)}{16} = \frac{j}{4} - 2744 .$$

The polynomial f_j in the statement is precisely $T^3 - \sigma_1 T^2 + \sigma_2 T - \sigma_3$ with $\sigma_1, \sigma_2, \sigma_3$ as above, and its zeroes are $\alpha_1, \alpha_2, \alpha_3$ by construction. The discriminant of f_j is $4j^2(j - 1728)$, from

which it follows that f_j has repeated zeroes in $\overline{K}$ if and only if $j \notin \{0, 1728\}$. Assume now that f_j is irreducible over $K[T]$ and separable. Then its Galois group over K is isomorphic to $\mathfrak{A}_3$ if $j - 1728$ is a square in K and to $\mathfrak{S}_3$ if $j - 1728$ is a non-square in K. The degree of the splitting field of f_j over K follows from this fact. □

The fact that f_j is separable if and only if $j \notin \{0, 1728\}$ is in accordance with Example 6.5.1. The zeroes of f_j can easily be computed using for instance Cardano's formulas (in characteristic not 3). Let $i \in \{1, 2, 3\}$ and let $L_i = K(\alpha_i)$. From α_i, d_i and $1/d_i$ are obtained as the zeroes of $q_i(T) = T^2 - \alpha_i T + 1$. If $(\alpha_i - 2)(\alpha_i + 2)$ is a square in L_i, then $d_i, 1/d_i$ are defined over L_i, otherwise over the quadratic extension of L_i obtained by adjoining $\sqrt{(\alpha_i - 2)(\alpha_i + 2)}$.

In Section 6.7 we will show that the splitting type of $\gamma_j(T)$ over K is entirely characterized by the 2-torsion behaviour of the underlying isomorphism class of elliptic curves.

6.6 Geometric construction of Edwards curves

Given an elliptic curve E/K in Weierstrass form together with a point P of order 4, [BL07, Theorem 2.1] shows how to construct an Edwards curve birationally equivalent to (a quadratic twist of) E. In this section, we generalize this construction over extensions of the base field. Indeed not every elliptic curve has a point of order 4 over its base field. We consider all the 12 points of order 4 on E defined over $\overline{K}$, and to each of these points, we associate the Edwards curves birationally equivalent to E, obtained from the construction of [BL07]. We shall then describe explicit relations between these curves and thereby obtain a description of the set $\mathrm{Edw}_{\overline{K}}(E)$.

6.6.1 The construction of Bernstein-Lange

We start by reviewing the construction of [BL07]. For full details, we refer to the proof of [BL07, Theorem 2.1].

Assume that E is given in long Weierstrass model $y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$ with $a_i \in K$. Let $P = (x_P, y_P) \in E(K)$ be a point of order 4. By a suitable translation of E, one may assume that $2P = (0, 0)$ and thus E is given by a medium Weierstrass model. Denote the new coordinates on E by $\bar{x}, \bar{y}$. Setting $d := 1 - 4\bar{x}_P^3/\bar{y}_P^2$, the Edwards curve $\mathcal{E}_d$ is birationally equivalent to a quadratic twist of E. In [BBJ+08, Theorem 3.2] the twist was shown to be unnecessary, so that $\mathcal{E}_d$ is birationally equivalent to E over K. Because of the clear dependence on P, we shall rather write d_P instead of d. Applying the necessary changes of variables, one easily shows the more general result:

Proposition 6.6.1. *Let E be an elliptic curve over K defined by $y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$, with $a_1, \ldots, a_6 \in K$. Assume that $E(K)$ has a point $P = (x_P, y_P)$ of exact order 4. Set*

$$d_P = 1 - \frac{4(x_P - x_{2P})^3}{[y_P + \frac{1}{2}(a_1x_P + a_3)]^2}\,. \tag{6.11}$$

(i) Let $\overline{x} = x - x_{2P}$ and $\overline{y} = y + \frac{1}{2}(a_1x_P + a_3)$. The Edwards curve $\mathcal{E}_{d_P}$ is birationally equivalent over K to the quadratic twist E^{t_P} of E given by $t_P \cdot \overline{y}^2 = \overline{x}^3 + a_2\overline{x}^2 + a_4\overline{x}$, with $t_P = x_P/(1 - d_P)$.
(ii) The Edwards curve $\mathcal{E}_{d_P}$ is birationally equivalent over K to E.

Proof. Using the change of coordinates $(\overline{x}, \overline{y}) \mapsto (x - x_{2P}, y + \frac{1}{2}(a_1 x_P + a_3))$, the formula for d_P easily follows from the proof of Theorem 2.1 in [BL07]. Point (i) is Theorem 2.1 of [BL07]. Point (ii) is Theorem 3.3 of [BBJ$^+$08]. □

Remark 6.6.2. The birational equivalences between $\mathcal{E}_{d_P}$ and (the quadratic twist of) E are made explicit in the proofs of [BL07, Theorem 1.2] and [BBJ$^+$08, Theorem 3.3]. Note that the birational equivalence in (ii) depends on the chosen point P.

An alternative way to construct Edwards curves from elliptic curves having a point of order 4 over its base field is to use the Tate normal form (see [Hus04, §4] and [Kna92, §V.5]), parametrizing elliptic curves with a torsion point of prescribed order. We briefly elaborate this here. For $b, c \in K$, define the two-parameter family of elliptic curves in Tate normal form

$$E(b,c) : y^2 + (1-c)xy - by = x^3 - bx^2$$

with discriminant $\Delta(b,c) = (1-c)^4b^3 - (1-c)^3b^3 - 8(1-c)^2b^4 + 36(1-c)b^4 - 27b^4 + 16b^5$; thus if $b \neq 0$ then $E(b,c)$ is elliptic. Let $Q = (0,0)$ on $E(b,c)$; one checks that $2Q$ and $3Q$ are not the point at infinity. Let now E be an elliptic curve over K with a point P of order $N \geq 4$. Then there is a unique pair $(E(b,c), Q)$ with $Q = (0,0)$ of order N and such that E is isomorphic to $E(b,c)$ and the isomorphism sends P to Q (see e.g. [Str15, Lemma 2.1]). Requiring that $E(b,c)$ has a point of order $N \geq 4$ gives rise to algebraic equations that b and c must satisfy. For our case, $Q = (0,0)$ has order 4 on $E(b,c)$ if and only if $c = 0$, giving E the Tate normal form

$$E(b,0) : y^2 + xy - by = x^3 - bx^2$$

with discriminant $\Delta(b,0) = b^4(1+16b)$ and j-invariant $j(b,0) = (16b^2+16b+1)^3/((16b+1)b^4)$. One checks that $2Q = (b, bc)$ on $E(b,c)$; i.e. $2Q = (b,0)$ on $E(b,0)$. To obtain a birational Edwards model, we apply Proposition 6.6.1 to $E(b,0)$ and $Q = (0,0)$ and obtain

$$d_Q = 1 - \frac{4(-b)^3}{1/4b^2} = 1 + 16b\,,$$

showing that $E(b,0)$ is K-birationally equivalent to the Edwards curve $\mathcal{E}_{1+16b}$.

6.6.2 Edwards curves from the two-torsion group

Let E be an elliptic curve over K be given in Weierstrass form as in Proposition 6.6.1. Denote by $E_4(\overline{K}) := E(\overline{K})[4] \setminus E(\overline{K})[2]$ the set of twelve points of order exactly 4 in $E(\overline{K})$. We define

$$d : E_4(\overline{K}) \to \overline{K}\ ,\ P \mapsto d_P \tag{6.12}$$

with d_P as in Equation (6.11). We obtain the following geometric description

$$\mathrm{Edw}_{\overline{K}}(E) = \{\mathcal{E}_{d_P} \mid P \in E_4(\overline{K})\}\ .$$

Remark 6.6.3. *(i)* Clearly, the map d in Equation (6.12) depends on E. In fact, it only depends on the $\overline{K}$-isomorphism class of E. More precisely, let E/K, E'/K be elliptic curves and $\phi : E \to E'$ a $\overline{K}$-isomorphism. Assume that E is given by the long Weierstrass equation $y^2 +$

$a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$. Denote the map in (6.12) by $d^{(E)}$ for E and by $d^{(E')}$ for E'. Then:

$$d_P^{(E)} = d_{\phi(P)}^{(E')} \quad , \quad \forall\, P \in E_4(\overline{K})\,.$$

This is be established by explicit calculations on the Weierstrass model, which we omit here. Since ϕ is an isomorphism, it is given by a change of variables $(x, y) \mapsto (u^2x + r, u^3y + su^2x + t)$, for some $u, r, s, t \in \overline{K}$ and $u \neq 0$ ([Sil09, Chapter III, §3, Proposition 3.1]). The substitution in the Weierstrass model for E gives for E' the Weierstrass model with coefficients $a_1' = (a_1 + 2s)/u$, $a_2' = (a_2 - sa_1 + 3r - s^2)/u^2$, $a_3' = (a_3 + ra_1 + 2t)/u^3$, $a_4' = (a_4 - sa_3 + 2ra_2 - (t+rs)a_1 + 3r^2 - 2st)/u^4$, $a_6' = (a_6 + ra_4 + r^2a_2 + r^3 - ta_3 - t^2 - rta_1)/u^6$ (see [Sil09, Chapter III, §1, Table 3.1]). Let $Q = \phi(P) = (x_Q, y_Q)$. The x-coordinate x_{2Q} of $2Q$ can be computed by the duplication formula on E', and is explicit in $a_1, \ldots, a_6$ via the above relations. A substitution in the formula for $d_Q^{(E')}$ gives $d_Q^{(E')} = 1 - 4(x_{2Q} - x_Q)^3/[y_Q + 1/2(a_1'x_Q + a_3')]^2 = d_P^{(E)}$.

(ii) We provide an example in SageMath [S$^+$20] for the computation of the map d, and an illustration of point *(i)* of the current remark.

```
sage: def compute_d(E,P):
....: #Input: E elliptic curve with 4-torsion, P point of order 4 on E
....: #Output: parameter d for birational Edwards model
....:     (xP,yP,x2P)=E(P).xy()[0],E(P).xy()[1],E(2*P).xy()[0]
....:     return 1-4*(xP-x2P)^3/(yP+1/2*(E.a1()*xP+E.a3()))^2

sage: E=EllipticCurve([0,0,0,-11,14]);E
Elliptic Curve defined by y^2 = x^3 - 11*x + 14 over Rational Field
sage: F=E.change_weierstrass_model([2,1,5,-1/3]);F
Elliptic Curve defined by y^2 + 5*x*y - 1/12*y = x^3 - 11/2*x^2 - 7/24*x
+ 35/576 over Rational Field
sage: E.j_invariant(),F.j_invariant()
(287496, 287496)
sage: phi=E.isomorphism_to(F)
sage: P=E(1,2);Q=phi(P); P,Q,P.order(),Q.order()
((1 : 2 : 1), (0 : 7/24 : 1), 4, 4)
sage: dP=compute_d(E,P); dQ=compute_d(F,Q); dP,dQ
(2, 2)
```

In light of Remark 6.6.3, one can work (as in Section 6.4) with an isomorphism class of elliptic curves $[E] \in \mathfrak{Ell}(j)$ for a fixed j-invariant in $j \in K$, and construct the set $\mathrm{Edw}_{\overline{K}}(j)$.

The points P of order 4 (hence the corresponding curves $\mathcal{E}_{d_P}$) are defined over $K(E[4])$, the 4-division field of E. The extension $K(E[4])/K$ is Galois and

$$\mathrm{Gal}(K(E[4])/K) \to \mathrm{Aut}(E(\overline{K})[4])\ ,\ \sigma \mapsto ((x,y) \mapsto (\sigma(x), \sigma(y)))$$

is an injection. The choice of a basis of $E(\overline{K})[4]$ induces an isomorphism $\mathrm{Aut}(E(\overline{K})[4]) \to \mathrm{GL}_2(\mathbb{Z}/4\mathbb{Z})$, and thus, $[K(E[4]) : K] \leq \#\mathrm{GL}_2(\mathbb{Z}/4\mathbb{Z}) = 96$. In the sequel, we shall see that the Edwards curves can be defined over much smaller extensions of K. Namely, since the denominator of d_P is a square it can be expressed in terms of x_P, which implies $d_P \in K(x_P)$, the extension of K obtained by adjoining the x-coordinate of P.

Remark 6.6.4. Let P, Q be generators for $E(\overline{K})[4] \simeq \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}$. Then $E_4(\overline{K})$ is the set $E_4(\overline{K}) = \{\pm P, \pm Q, \pm(P+Q), \pm(P-Q), 2P \pm Q, \pm P + 2Q\}$, where $\pm P, 2P \pm Q$ double to $2P$, $\pm Q, 2Q \pm P$ double to $2Q$ and $\pm(P \pm Q)$ double to $2(P+Q)$. Writing E again in usual long Weierstrass equation, we have for every $R \in E_4(\overline{K})$ that $-R = (x_R, -a_1 x_R - a_3 - y_R)$, and obtain $d_{-R} = 1 - 4(x_R - x_{2R})^3/(-y_R - 1/2(a_1 x_R + a_3))^2 = d_R$. It follows again that $\#\mathrm{Edw}_{\overline{K}}(j) \leq 6$, as already observed in Section 6.5.

In the sequel we write

$$E(\overline{K})[2] = \{\mathcal{O}, (e_1, 0), (e_2, 0), (e_3, 0)\} \simeq \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}\ ,$$

with $e_1, e_2, e_3 \in \overline{K}$ two by two distinct. Namely, since $\mathrm{char}(K) \neq 2$, we can find a rational transformation from a long Weierstrass form for E to a medium Weierstrass form $y^2 = f(x)$, and the 2-torsion subgroup has the above form, with e_1, e_2, e_3 being the roots of $f(x)$. For $i \neq j$ in $\{1,2,3\}$, set $e_{ij} := e_i - e_j \in \overline{K}^\times$. The following is the main theorem of this section.

Theorem 6.6.5. *Let E/K be an elliptic curve and let $E(\overline{K})[2] = \{\mathcal{O}, (e_1, 0), (e_2, 0), (e_3, 0)\}$ be its 2-torsion subgroup. If $\mathcal{E}_d$ is an Edwards curve birationally equivalent to E, then*

$$d \in \left\{ \left(\frac{\sqrt{e_{12}} \pm \sqrt{e_{13}}}{\sqrt{e_{32}}} \right)^4, \left(\frac{\sqrt{e_{21}} \pm \sqrt{e_{23}}}{\sqrt{e_{31}}} \right)^4, \left(\frac{\sqrt{e_{31}} \pm \sqrt{e_{32}}}{\sqrt{e_{21}}} \right)^4 \right\}\ .$$

Define $z_1 = e_{12}e_{13}, z_2 = e_{21}e_{23}$ and $z_3 = e_{31}e_{32}$. The Edwards curves are defined over

$$K_1 := K(\sqrt{z_1})\ ,\ K_2 := K(\sqrt{z_2})\ ,\ K_3 := K(\sqrt{z_3})\ ,$$

respectively, and birationally equivalent to E over $K(\sqrt{e_{12}}, \sqrt{e_{13}}), K(\sqrt{e_{21}}, \sqrt{e_{23}}), K(\sqrt{e_{31}}, \sqrt{e_{32}})$, respectively. In particular:

(i) if $E(K)[2] \simeq \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$, then for all $i \in \{1,2,3\}$:

$$[K_i : K] = \begin{cases} 1 & \text{if } z_i \in (K^\times)^2 \\ 2 & \text{if } z_i \notin (K^\times)^2 \end{cases}$$

(ii) if $E(K)[2] \simeq \mathbb{Z}/2\mathbb{Z}$ then there exists $i \in \{1,2,3\}$ such that $[K_i : K] \in \{1,2\}$ (as in (i)) and for $j \neq j' \in \{1,2,3\} \setminus \{i\}$, $K_j \simeq K_{j'} \simeq K(\sqrt{\Delta(E)})$ and

$$[K_j : K] = \begin{cases} 2 & \text{if } z_j \in (K(\sqrt{\Delta(E)})^\times)^2 \\ 4 & \text{if } z_j \notin (K(\sqrt{\Delta(E)})^\times)^2 \end{cases}$$

(iii) if $E(K)[2] \simeq \{\mathcal{O}\}$ then $K_1 \simeq K_2 \simeq K_3$ and for all $i \in \{1,2,3\}$:

$$[K_i : K] = \begin{cases} 3 & \text{if } z_i \in (K(E[2]))^\times)^2 \\ 6 & \text{if } z_i \notin (K(E[2]))^\times)^2 \end{cases} \quad \textit{and} \quad \Delta(E) \in (K^\times)^2$$

and

$$[K_i : K] = \begin{cases} 6 & \text{if } z_i \in (K(E[2]))^\times)^2 \\ 12 & \text{if } z_i \notin (K(E[2]))^\times)^2 \end{cases} \quad \textit{and} \quad \Delta(E) \notin (K^\times)^2$$

Remark 6.6.6. Let us make some remarks concerning how Theorem 6.6.5 compares to the literature.

(i) One important gain over [BL07, Theorem 2.1] is that our formulae only amount to the computation of the 2-torsion of E and do not require the knowledge of its 4-torsion. Further, we see by the cases (i) and (ii) that the "converse" of the Bernstein-Lange construction is not true: d_P can be defined over K without P lying in $E(K)$. Moreover, our formulae can be applied to every elliptic curve, and describe explicitly the smallest extensions of their base field, over which they admit an Edwards model.

(ii) In [BBJ$^+$08, Theorem 5.1], the authors showed that an elliptic curve E/K having no point of order 4 over K, but all its 2-torsion points defined over K (i.e. $E(K)[2] \simeq \mathbb{Z}/2\mathbb{Z}\times\mathbb{Z}/2\mathbb{Z}$) is 2-isogenous over K to an elliptic curve birationally equivalent to a twisted Edwards curve. The main interest of this result is to find Edwards models for E when E does not have points of order 4 over its base field. We will extend this result over larger fields in Section 6.8.1. Point (ii) of Theorem 6.6.5 contributes in this direction and strengthens [BBJ$^+$08, Theorem 5.1]: when $E(K)$ has a unique point of order 2, and say $K_i = K$ (i.e. $[K_i : K] = 1$ in (ii)) for some $i \in \{1, 2, 3\}$, then E admits an Edwards model defined over K. Nevertheless, note that having a unique point of order 2 is not enough to define an Edwards model over K applying [BBJ$^+$08, Theorem 5.1].

Proof of Theorem 6.6.5

The proof of Theorem 6.6.5 is essentially the content of the Lemmas 6.6.7 and 6.6.9 below. The following result describes the 4-torsion of E in terms of the 2-torsion.

Lemma 6.6.7. *Let E/K be an elliptic curve and let $E[2](\overline{K}) = \{\mathcal{O}, (e_1, 0), (e_2, 0), (e_3, 0)\}$. For every $i \in \{1, 2, 3\}$ the four points*

$$\pm P_i^{\epsilon} := \left(e_i + \epsilon\sqrt{e_{ij}}\sqrt{e_{ik}}\ ,\ \pm\sqrt{e_{ij}}\sqrt{e_{ik}}(\sqrt{e_{ij}} + \epsilon\sqrt{e_{ik}})\right)\ ,\ j \neq k \in \{1, 2, 3\} \setminus \{i\}$$

with $\epsilon \in \{\pm 1\}$, have order exactly 4 in $E(\overline{K})$ and double to $(e_i, 0)$. We let $\pm P_i^+ := \pm P_i^1$ and $\pm P_i^- := \pm P_i^{-1}$.

Proof. This follows from [Sch14, page 112]. □

Remark 6.6.8. *(i)* As a consequence of Lemma 6.6.7, the 4-division field of E is

$$K(E[4]) = K(\sqrt{-1}, \sqrt{e_{12}}, \sqrt{e_{13}}, \sqrt{e_{23}})\ .$$

If e_{12}, e_{13} and e_{23} are simultaneously squares in K up to sign, then $K(E[4]) = K(\sqrt{-1}) = K(\zeta_4)$. In [GJLR16] the authors classify, up to isomorphism, elliptic curves over $K = \mathbb{Q}$ with this property. Assume now that E has full K-rational 2-torsion. Then by Lemma 6.6.7, E has a point of order 4 over K if either e_{12}, e_{13}, or e_{21}, e_{23}, or e_{31}, e_{32} are both squares in K. In Section 6.10.1 we provide examples of such families of elliptic curves.

(ii) This remark concerns Proposition 6.6.1 (i.e. [BL07, Theorem 2.1] and [BBJ$^+$08, Theorem 3.3]). If E has a K-rational point of order 4, then it has at least four such points, appearing in pairs of opposite points. Hence, to one point P as in Proposition 6.6.1 one associates two Edwards curves via the map d. We establish an explicit relation between these curves in what follows.

(iii) The 6 different x-coordinates of the points of Lemma 6.6.7 are the roots over $\overline{K}$ of the 4-division polynomial $\Phi_4 \in K[x]$ of E, of degree 6.

Lemma 6.6.9. *Let E/K be an elliptic curve and let $E[2](\overline{K}) = \{\mathcal{O}, (e_1, 0), (e_2, 0), (e_3, 0)\}$. For $i \in \{1, 2, 3\}$, let $\{\pm P_i^\epsilon \mid \epsilon \in \{\pm 1\}\}$ be the points of order 4 on E over $\overline{K}$ such that $2P_i = (e_i, 0)$. Let $d : E_4(\overline{K}) \to \overline{K}$ be the map from* (6.12)*. Then for every $i \in \{1, 2, 3\}$ and $\epsilon \in \{\pm 1\}$, the image of $\pm P_i^\epsilon$ under d is*

$$d_{\pm P_i^\epsilon} = \left(\frac{\sqrt{e_{ij}} - \epsilon\sqrt{e_{ik}}}{\sqrt{e_{kj}}} \right)^4 \in K_i = K(\sqrt{e_{ij}e_{ik}}) \ , \ j < k \in \{1, 2, 3\} \setminus \{i\} \ . \tag{6.13}$$

Proof. The fact that $d_{\pm P_i^\pm}$ lies in $K(x_{\pm P_i^\pm})$ is clear by the definition of d, and we have that that $K(x_{\pm P_i^\pm}) = K(\sqrt{e_{ij}e_{ik}})$ by Lemma 6.6.7.

We show (6.13); for easier notation we prove the statement for $i = 1$. By Remark 6.6.4, it suffices to establish the formula for $d_{P_1^\pm}$. For $\epsilon \in \{\pm 1\}$, we have by Lemma 6.6.7, $P_1^\epsilon = (e_1 + \epsilon\sqrt{e_{12}}\sqrt{e_{13}}, \sqrt{e_{12}}\sqrt{e_{13}}(\sqrt{e_{12}} + \epsilon\sqrt{e_{13}})) = (x_{P_1}, y_{P_1})$. It follows that

$$d_{P_1^\epsilon} = 1 - 4\frac{(x_{P_1^\epsilon} - x_{2P_1^\epsilon})^3}{y_{P_1^\epsilon}^2} = 1 - 4\epsilon\frac{\sqrt{e_{12}}\sqrt{e_{13}}}{(\sqrt{e_{12}} + \epsilon\sqrt{e_{13}})^2} = \left(\frac{\sqrt{e_{12}} - \epsilon\sqrt{e_{13}}}{\sqrt{e_{12}} + \epsilon\sqrt{e_{13}}} \right)^2 \in K(\sqrt{e_{12}e_{13}}) = K_1 \ .$$

Further, since $e_{12} - e_{13} = e_{32}$,

$$\left(\frac{\sqrt{e_{12}} - \epsilon\sqrt{e_{13}}}{\sqrt{e_{12}} + \epsilon\sqrt{e_{13}}} \right)^2 = \left(\frac{(\sqrt{e_{12}} - \epsilon\sqrt{e_{13}})^2}{e_{12} - e_{13}} \right)^2 = \left(\frac{\sqrt{e_{12}} - \epsilon\sqrt{e_{13}}}{\sqrt{e_{32}}} \right)^4$$

□

This concludes the first part of the proof of Theorem 6.6.5. In order to show the applicability of these formulas, we treat in detail the special j-invariants 0 and 1728.

Example 6.6.10. We detail the geometric construction of the sets $\mathrm{Edw}_{\overline{K}}(j)$ for $j \in \{0, 1728\}$. Since the involved computations are a little lengthy, we defer the details to Section 6.10.2. Compare with Example 6.5.1 for the algebraic construction of these sets.

To show the second part of Theorem 6.6.5, notice that the extensions $\{K_i\}_i$ are at most quadratic over $K(E[2]) = K(e_1, e_2, e_3)$ and the claims follow from the classification of $K(E[2])$. Thus, for every $i \in \{1, 2, 3\}$ we have $[K_i : K] = [K_i : K(E[2])] \cdot [K(E[2]) : K]$ where the first factor is 1 or 2.

If E has full K-rational 2-torsion, then $K(E[2]) = K$ and $[K_1 : K] = 1$ if $z_1 \in (K^\times)^2$ and $[K_1 : K] = 2$ otherwise, and similarly for K_2 and K_3.

If E has a unique K-rational point of order 2, say $(e_1, 0) \in E(K)$ without loss of generality, then e_2 and e_3 are conjugates in the quadratic extension $K(\sqrt{\delta})$, where $\delta = (e_2 - e_3)^2 \in K^\times \setminus (K^\times)^2$ is the discriminant of $(x - e_2)(x - e_3)$. Since $\Delta(E) = 16\delta(e_1 - e_2)^2(e_1 - e_3)^2$, it is easy to see that $K(\sqrt{\delta}) = K(\sqrt{\Delta(E)})$, i.e. $K(E[2]) = K(\sqrt{\Delta(E)})$. We deduce that $[K_1 : K]$ equals 1 or 2 depending on whether z_1 is or not a square in K. Write $z_2 := e_{21}e_{31}$, $z_3 := e_{31}e_{32}$ and $K_2 = K(\sqrt{z_2}), K_3 = K(\sqrt{z_3})$. Denoting by σ the conjugation automorphism of $K(\sqrt{\Delta(E)})$, and writing $e_3 = \sigma(e_2)$, we get $\sigma(z_2) = \sigma((e_2 - e_1)(e_2 - \sigma(e_2)) = (\sigma(e_2) - e_1)(\sigma(e_2) - e_2) = z_3$. Therefore, z_2 and z_3 are simultaneously squares or non-squares in $K(\sqrt{\Delta(E)})$, and moreover, if $z_2 = u^2$ for $u \in K(\sqrt{\Delta(E)})$, then $z_3 = \sigma(u)^2$, so there is an isomorphism $K_2 \simeq K_3$. The degree over K is either 2 or 4, according to the indicated cases.

If E has trivial K-rational 2-torsion then none among e_1, e_2, e_3 lie in K, and we distinguish the cases $\Delta(E) \in (K^\times)^2$ and $\Delta(E) \notin (K^\times)^2$. In the first case, the Galois group of $K(E[2])/K$ is isomorphic to $\mathfrak{A}_3$, so $[K(E[2]) : K] = 3$. If $\mathfrak{A}_3 = \langle\sigma\rangle$, then without loss of generality, $e_2 = \sigma(e_1)$ and $e_3 = \sigma^2(e_1)$. This implies, as before, that $\sigma(z_1) = z_2$ and $\sigma^2(z_1) = z_3$. Therefore K_1, K_2 and K_3 are all isomorphic to $K(E[2])(\sqrt{z_1})$. In the second case, the Galois group of $K(E[2])/K$ is isomorphic to $\mathfrak{S}_3$, i.e. $[K(E[2]) : K] = 6$. By a similar reasoning, $K_1 \simeq K_2 \simeq K_3 \simeq K(E[2])(\sqrt{z_1})$, which is of degree 6 over K if $z_1 \in (K(E[2])^\times)^2$ and 12 otherwise. This completes the proof of Theorem 6.6.5.

6.6.3 Link with the algebraic construction

Our Theorem 6.6.5 gives an explicit factorisation of the polynomial f_j from Equation (6.10) when j denotes the j-invariant of E.

Corollary 6.6.11. *Let E be as in Theorem 6.6.5 and let j be the j-invariant of E. The roots $\alpha_1, \alpha_2, \alpha_3$ of the polynomial f_j are defined over $K(E[2])$, and given by the formulae*

$$\alpha_1 = \frac{2(e_{12}^2 + 6z_1 + e_{13}^2)}{e_{32}^2}\ ,\ \alpha_2 = \frac{2(e_{21}^2 + 6z_2 + e_{23}^2)}{e_{31}^2}\ ,\ \alpha_3 = \frac{2(e_{31}^2 + 6z_3 + e_{32}^2)}{e_{21}^2}$$

In particular:

(i) if $E(K)[2] \simeq \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$, then f_j splits completely over $K[T]$,

(ii) if $E(K)[2] \simeq \mathbb{Z}/2\mathbb{Z}$ then f_j splits in $K[T]$ in one factor of degree 1 and one irreducible factor of degree 2,

(iii) if $E(K)[2] \simeq \{\mathcal{O}\}$ then f_j is irreducible in $K[T]$.

Proof. These are direct computations of $d_i + 1/d_i$ for $i \in \{1, 2, 3\}$ using the formulae of Theorem 6.6.5. The second part on the splitting of f_j also easily follows. □

Remark 6.6.12. Without the use of Theorem 6.6.5, Corollary 6.6.11 can be obtained as follows. For E with 2-torsion as in Theorem 6.6.5, one computes that $j(E) = 2^8(e_1^2 - e_1e_2 + e_2^2 - e_1e_3 - e_2e_3 + e_3^2)^3)/(e_{12}^2e_{13}^2e_{23}^2)$ and one computes a factorsiation of $f_{j(E)}(T)$. A direct computation shows that $j(E) - 1728$ is a square in K if and only if $\Delta(E)$ is, whence the splitting properties.

6.6.4 Modular d-function

In his original exposition [Edw07], Edwards considered an elliptic function which is associated to the Edwards model. In this section, we make a remark about this function in relation with our Theorem 6.6.5.

Let $\mathfrak{h} = \{\tau \in \mathbb{C} : \Im(\tau) > 0\}$ be the complex upper-half plane. Let $\Gamma = \mathrm{SL}_2(\mathbb{Z})$ and consider its congruence subgroup $\Gamma_1(4)$ of level 4 acting on $\mathfrak{h}$ via fractional linear transformations. Recall that a function $f : \mathfrak{h} \to \mathbb{C}$ is said a *modular function* or *modular form of weight* 0 for $\Gamma_1(4)$ if f is meromorphic on $\mathfrak{h}$, $f(\gamma(\tau)) = f(\tau)$ for every $\gamma \in \Gamma_1(4)$, and f has a Fourier expansion of the form $f(\tau) = \sum_{n=-m}^{\infty} a(n) \exp(2\pi i n\tau)$ for some $m \in \mathbb{N}$.

Denote by $X_1(4)$ the projective compactification of $Y_1(4) = \Gamma_1(4)\backslash\mathfrak{h}$ in $\mathbb{P}^2(\mathbb{C})$ by considering the action of $\Gamma_1(4)$ on the extended complex upper half-plane $\mathfrak{h}^* = \mathfrak{h} \cup \mathbb{Q} \cup \{\infty\}$. Points

on this curve parametrize isomorphism classes of pairs (E, P) where E is a complex elliptic curve and $P \in E(\mathbb{C})$ is a point of order 4. As projective curve, $X_1(4)$ is known to have genus 0, and thereby has a *rational* function field, i.e. the function field over $\mathbb{C}$ is generated by a unique function. This generator has been studied in [KK98] and denoted by $j_{1,4}$. We here give the link between this generator and the Edwards model of elliptic curves.

Define the classical *theta constants* obtained by specialization at 0 of the Jacobi theta functions:

$$\vartheta_2(\tau) = \sum_{n\in\mathbb{Z}} q^{(n+1/2)^2} \ , \ \vartheta_3(\tau) = \sum_{n\in\mathbb{Z}} q^{n^2} \ , \ \vartheta_4(\tau) = \sum_{n\in\mathbb{Z}} (-1)^n q^{n^2} \ .$$

Define the function $d : \mathfrak{h} \to \mathbb{C}$ by

$$d(\tau) := \frac{\vartheta_2(2\tau)^4}{\vartheta_3(2\tau)^4} \ . \tag{6.14}$$

In [KK98] this function is denoted by $j_{1,4}$. The functions $\vartheta_2(2\tau)$ and $\vartheta_3(2\tau)$ are modular forms of weight $1/2$ for $\Gamma_1(4)$, thus $d = j_{1,4}$ is a modular function for $\Gamma_1(4)$. Moreover, $j_{1,4}$ is a Hauptmodul for $X_1(4)$, i.e. a generator for the function field of $X_1(4)$ over $\mathbb{C}$, see [KK98, Theorem 3]. By using our Theorem 6.6.5, we can link d to the Weierstrass $\wp$-function of a complex torus.

Proposition 6.6.13. *For $\tau \in \mathfrak{h}$, let $\Lambda_\tau = \mathbb{Z} + \tau\mathbb{Z}$. The complex elliptic curve $\mathbb{C}/\Lambda_\tau$ is isomorphic to the complex Edwards curve $x^2 + y^2 = 1 + d(\tau)x^2y^2$.*

Proof. We write $\tau = \omega_2/\omega_1$ where $\omega_1, \omega_2 \in \mathbb{C}$ are a fundamental pair of periods (i.e. both non-zero and such that $\Im(\omega_2/\omega_1) > 0$). Then the period lattice $\Lambda = \omega_1\mathbb{Z} \oplus \omega_2\mathbb{Z}$ in $\mathbb{C}$ is homothetic to Λ_τ. Let $\wp(z) := \wp(z;\tau) = \wp_{\Lambda_\tau}(z)$ with complex variable z be the Weierstrass $\wp$-function associated to Λ_τ. It satisfies the differential equation $\wp'(z)^2 = 4\wp(z)^3 - g_2\wp(z) - g_3$ where the invariants $g_2 = 60\sum_{\omega\in\Lambda_\tau\setminus\{0\}} \omega^{-4}$ and $g_3 = 140\sum_{\omega\in\Lambda_\tau\setminus\{0\}} \omega^{-6}$ are the elliptic invariants, viewed as functions of τ. It is well-known (e.g. [Cha85, Chapter III, §3]) that the zeroes of the cubic $4\wp^3(z) - g_2\wp(z) - g_3$ characterizing the 2-torsion on $\mathbb{C}/\Lambda_\tau$, are given in terms of the half-periods

$$e_1 := e_1(\tau) = \wp\left(\frac{\omega_1}{2}\right) \ , \ e_2 := e_2(\tau) = \wp\left(\frac{\omega_2}{2}\right) \ , \ e_3 := e_3(\tau) = \wp\left(\frac{\omega_3}{2}\right) \tag{6.15}$$

where we have set $\omega_3 = \omega_1 + \omega_2$. By [Cha85, Chapter V, §6, Corollary 1], e_1, e_2, e_3 are related to the theta constants via the relations

$$e_{12}(\tau) = \left(\frac{\pi}{\omega_1}\right)^2 \vartheta_3(\tau)^4 \ , \ e_{13}(\tau) = \left(\frac{\pi}{\omega_1}\right)^2 \vartheta_4(\tau)^4 \ , \ e_{32}(\tau) = \left(\frac{\pi}{\omega_1}\right)^2 \vartheta_2(\tau)^4 \tag{6.16}$$

with the usual notation $e_{12} = e_1 - e_2, e_{13} = e_1 - e_3, e_{32} = e_3 - e_2$. Now observe that

$$d(\tau) = \frac{\vartheta_2(2\tau)^4}{\vartheta_3(2\tau)^4} = \left(\frac{\vartheta_3(\tau)^2 - \vartheta_4(\tau)^2}{\vartheta_3(\tau)^2 + \vartheta_4(\tau)^2}\right)^2 = \left(\frac{\vartheta_3(\tau)^2 - \vartheta_4(\tau)^2}{\vartheta_2(\tau)^2}\right)^4$$

where for the second equality we have used the duplication formulae $2\vartheta_2(2\tau)^2 = \vartheta_3(\tau)^2 - \vartheta_4(\tau)^2$ and $2\vartheta_3(2\tau)^2 = \vartheta_3(\tau)^2 + \vartheta_4(\tau)^2$ (see [Liu09, Equations (4-1) and (4-2)] with $x = 0$), and

for the last equality we have multiplied the numerator and denominator by $\vartheta_3(\tau)^2 - \vartheta_4(\tau)^2$ and used the identity $\vartheta_3(\tau)^4 - \vartheta_4(\tau)^4 = \vartheta_2(\tau)^4$. Using the relations (6.16) gives therefore

$$d(\tau) = \left(\frac{\epsilon_1 \sqrt{e_{12}(\tau)} - \epsilon_2 \sqrt{e_{13}(\tau)}}{\epsilon_3 \sqrt{e_{32}(\tau)}} \right)^4$$

for some $\epsilon_1, \epsilon_2, \epsilon_3 \in \{\pm 1\}$. Direct computations show that without loss of generality ϵ_1 and ϵ_2 can be taken 1, so that

$$d(\tau) = \left(\frac{\sqrt{e_{12}(\tau)} \pm \sqrt{e_{13}(\tau)}}{\sqrt{e_{32}(\tau)}} \right)^4 \tag{6.17}$$

By Theorem 6.6.5 we conclude that there is an isomorphism $\mathbb{C}/\Lambda_\tau \simeq \mathcal{E}_{d(\tau)}$. □

Note that our proof also gives a formula for $d(\tau)$ in terms of the half-periods of Weierstrass' elliptic functions; to our knowledge, this representation was not established before. The remaining values of d from Theorem 6.6.5 are obtained by fractional linear transformations. The group $\Gamma_1(4)$ has index 12 in $\mathrm{SL}_2(\mathbb{Z})$. Further, d is left invariant by the transformation $\tau \mapsto \tau + 1$, which corresponds to permuting e_2 and e_3 in (6.17) (this can also be seen directly on the theta constants). Choosing representatives for the cosets of $\mathrm{SL}_2(\mathbb{Z})/\Gamma_1(4)$ gives rise to fractional linear transformations deriving the remaining values of d, and in particular $\Sigma(d)$. This is similar to [Edw07, Theorem 18.1] (with the change that the chosen function actually corresponds to a fourth root of $d(\tau)$).

Relation with other modular functions

The following proposition is a consequence of Proposition 6.6.13 and relates d to the modular j-invariant and the modular discriminant function.

Proposition 6.6.14. *Let $j : \mathfrak{h} \to \mathbb{C}$ be the modular j-invariant and $\Delta : \mathfrak{h} \to \mathbb{C}$ the modular discriminant. Let $d(\tau)$ be defined as above. For every $\tau \in \mathfrak{h}$, we have the identities:*

$$j(\tau) = 16 \frac{(1 + 14d(\tau) + d(\tau)^2)^3}{d(\tau)(d(\tau) - 1)^4} \quad , \quad \Delta(\tau) = 2^8 d(\tau)(d(\tau) - 1)^4$$

This proposition follows easily by computing the j-invariant and discriminant of the curve E_d birationally equivalent to $\mathcal{E}_d$ (see Section 6.9.2). For the well-known modular-λ function defined by $\lambda(\tau) = \frac{\vartheta_2(\tau)^4}{\vartheta_3(\tau)^4}$ (see [Cha85, Chapter VII, §7]) we obtain the relation

$$\lambda(\tau) = d(\tau/2) \, . \tag{6.18}$$

The modular-λ function can be derived from isomorphism classes of Legendre curves $\mathcal{L}_\lambda : y^2 = x(x-1)(x-\lambda)$ ([Cha85, Chapter VII, §7]), similarly as $d(\tau)$ is deduced from isomorphism classes of Edwards curves ([Edw07]). Notice that (6.18) can directly be established by the use of Theorem 6.6.5 since the 2-torsion on $\mathcal{L}_\lambda$ is $\{\mathcal{O}, (0,0), (1,0), (\lambda, 0)\}$. The modular λ-function is a Hauptmodul for the modular curve $X(2)$ parametrizing isomorphism classes of elliptic curves with full rational 2-torsion. We believe that in terms of suitable terminology, the function $d = j_{1,4}$ defined above could in correspondence be referred to as modular d-function in view of its link to the Edwards model $\mathcal{E}_d$ for elliptic curves.

6.7 Galois Conjugacy on isomorphic Edwards curves

We describe a natural Galois action on the Edwards curves induced from the Galois action on the points of elliptic curves. We then make a case study for conjugate Edwards curves depending on the size of $E[2](K)$, following the notation of Theorem 6.6.5. Let $G_K := \mathrm{Gal}(\overline{K}/K)$ denote the absolute Galois group of K.

Definition 6.7.1. *For $d \in \overline{K} \setminus \{0,1\}$ and $\sigma \in G_K$, we call ${}^{\sigma}(\mathcal{E}_d) := \mathcal{E}_{\sigma(d)}$ the* σ-conjugate Edwards curve of $\mathcal{E}_d$. *We also say that $\mathcal{E}_d$ and ${}^{\sigma}(\mathcal{E}_d)$ are* G_K-conjugates.

Definition 6.7.2. *Let $k \in \mathbb{N}_{\geq 1}$ and $r_1, r_2, \ldots, r_k \in \mathbb{N}_{\geq 1}$. Let K be a field and $f(T) \in K[T]$. We say that $f(T)$ has* splitting type $(r_1, r_2, \ldots, r_k)$ *(up to permutation) over K if f factors over $K[T]$ into irreducible factors of degree $r_1, r_2, \ldots, r_k$, i.e. $f(T) = \prod_{\ell=1}^{k} g_\ell(T)$ where $g_\ell(T) \in K[T]$ is irreducible over K of degree r_ℓ.*

We are interested in the splitting type of the polynomial $\gamma_{j(E)}(T)$. We explain below, that it depends on the size of the group $E[2](K)$.

The case $E(K)[2] \simeq \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$

Let $i \in \{1,2,3\}$. If z_i, as in Theorem 6.6.5, is a square in K, i.e. $[K_i : K] = 1$, then d_i and $1/d_i$ (and the corresponding Edwards curves) defined over K_i are not Galois conjugates. Otherwise, i.e. $[K_i : K] = 2$, d_i and $1/d_i$ are $\mathrm{Gal}(K_i/K)$-conjugates with minimal polynomial $T^2 - \alpha_i T + 1$ over K, and the corresponding Edwards curves are $\mathrm{Gal}(K_i/K)$-conjugates.

The case $E(K)[2] \simeq \mathbb{Z}/2\mathbb{Z}$

Let $i \in \{1,2,3\}$ such that $(e_i, 0)$ generates $E[2](K)$. Following Theorem 6.6.5, write $K_i = K(\sqrt{z_i})$ and $K_l \simeq K_{l'} \simeq K(\sqrt{\Delta(E)})(\sqrt{z_l})$ for $l \neq l' \in \{1,2,3\} \setminus \{i\}$. We distinguish whether z_1 is a square in K_1 or not and z_l is a square in $K(\sqrt{\Delta(E)})$ or not.

If z_1 is a square in K_1 then $K_1 = K$ and $\mathcal{E}_{d_1}$ and $\mathcal{E}_{1/d_1}$ over K are not conjugated. Otherwise, d_1 and $1/d_1$ have minimal polynomial $T^2 - \alpha_1 T + 1$ over K and the corresponding Edwards curves $\mathcal{E}_{d_1}$ and $\mathcal{E}_{1/d_1}$ are $\mathrm{Gal}(K_1/K)$-conjugates.

If z_l is a square in $K(\sqrt{\Delta(E)})$ (and thus also $z_{l'}$) then, similarly, the elements $d_l, 1/d_l$ and $d_{l'}, 1/d_{l'}$ have minimal polynomials $T^2 - \alpha_l T + 1$ and $T^2 - \alpha_{l'} T + 1$ over K, and the curves $\mathcal{E}_{d_l}, \mathcal{E}_{d_{l'}}$ and $\mathcal{E}_{1/d_l}, \mathcal{E}_{1/d_{l'}}$ are $\mathrm{Gal}(K(\sqrt{\Delta(E)})/K)$-conjugates.

If z_l is not a square in $K(\sqrt{\Delta(E)})$ then $K_l = K(\sqrt{\Delta(E)}, \sqrt{z_l})$ and $d_l, d_{l'}, 1/d_l, 1/d_{l'} \in K_l$ have minimal polynomial $(T^2 - \alpha_l T + 1)(T^2 - \alpha_{l'} T + 1) = T^4 - (\alpha_l + \alpha_{l'})T^3 + (\alpha_l \alpha_{l'} + 2)T^2 - (\alpha_l + \alpha_{l'})T + 1$ over K, and the corresponding curves $\mathcal{E}_{d_l}, \mathcal{E}_{d_{l'}}$ and $\mathcal{E}_{1/d_l}, \mathcal{E}_{1/d_{l'}}$ are $\mathrm{Gal}(K_l/K)$-conjugates.

The case $E(K)[2] \simeq \{\mathcal{O}\}$

In this case, $K_1 \simeq K_2 \simeq K_3 \simeq K(E[2])$.

Assume first that $\Delta(E)$ is a square in K, so that $\mathrm{Gal}(K(E[2])/K) \simeq \mathfrak{A}_3$. Theorem 6.6.5 says that $[K_i : K] \in \{3, 6\}$ for $i \in \{1,2,3\}$. Let σ be a generator for $\mathfrak{A}_3$. If $[K_i : K] = 3$ then z_1 (and also its conjugates $z_2 = \sigma(z_1), z_3 = \sigma^2(z_1)$) is a square in $K(E[2])$. Consequently, $f_{j(E)}$ splits completely as $(T - \alpha_1)(T - \alpha_2)(T - \alpha_3)$ in $K[T]$ with conjugate roots $\alpha_1, \alpha_2 = \sigma(\alpha_1), \alpha_3 =$

$\sigma^2(\alpha_1)$. Further, the (conjugate) quadratics $T^2-\alpha_i T+1$ for $i\in\{1,2,3\}$ lead to conjugate pairs of roots $\{d_1,1/d_1\},\{\sigma(d_1),\sigma(1/d_1)\},\{\sigma^2(d_1),1/\sigma^2(d_1)\}$. The elements $d_1,d_2,d_3\in K_1$ have minimal polynomial $(T-d_1)(T-d_2)(T-d_3)$ over K, and $1/d_1,1/d_2,1/d_3\in K_1$ have minimal polynomial $(T-1/d_1)(T-1/d_2)(T-1/d_3)$ over K. In particular the curves $\mathcal{E}_{d_1},\mathcal{E}_{d_2},\mathcal{E}_{d_3}$ and $\mathcal{E}_{1/d_1},\mathcal{E}_{1/d_2},\mathcal{E}_{1/d_3}$, respectively, are $\mathrm{Gal}(K(E[2])/K)$-conjugates. If $[K_i:K]=6$, then z_1 is not a square in $K(E[2])$ and d_1,d_2,d_3 and their inverses have minimal polynomial $1/16\gamma_{j(E)}$ over K and lead to $\mathrm{Gal}(K(E[2])/K)$-conjugated curves defined over $K(E[2])$, where σ sends $d_1\mapsto d_2\mapsto d_3$ and $1/d_1\mapsto 1/d_2\mapsto 1/d_3$. Letting $d_1=d$, with the notation of the set $\Sigma(d)$, we obtain that

$$d\,,\left(\frac{1+d^{1/4}}{1-d^{1/4}}\right)^4\,,\left(\frac{1+\sqrt{-1}d^{1/4}}{1-\sqrt{-1}d^{1/4}}\right)^4$$

are Galois-conjugates, as well as their inverses

$$\frac{1}{d}\,,\left(\frac{1-d^{1/4}}{1+d^{1/4}}\right)^4\,,\left(\frac{1-\sqrt{-1}d^{1/4}}{1+\sqrt{-1}d^{1/4}}\right)^4\,.$$

Assume now that $\Delta(E)$ is not a square in K, i.e. $\mathrm{Gal}(K(E[2])/K)\simeq\mathfrak{S}_3$. By Theorem 6.6.5, $[K_i:K]\in\{6,12\}$ for $i\in\{1,2,3\}$. Proceeding as above, it is immediate to establish that all the Edwards curves defined over $K(E[2])$ are $\mathrm{Gal}(K(E[2])/K)$-conjugates.

Example 6.7.3. *(i)* Consider the infinite family of elliptic curves over $\mathbb{Q}$ defined by

$$E_t: y^2=x^3-a(t)x^2+b(t)^2x-a(t)b(t)^2\,,\ t\in\mathbb{Q}\,,$$

with $a(t)=6+3t$ and $b(t)=8+4t$. The cubic in x factors over $\mathbb{Q}(\sqrt{-1})$ as $(x-a(t))(x-b(t)\sqrt{-1})(x+b(t)\sqrt{-1})$, so that $E_t(\mathbb{Q})[2]\simeq\mathbb{Z}/2\mathbb{Z}$ for every t. With the choice $e_1=a(t),e_2=b(t)\sqrt{-1},e_3=-b(t)\sqrt{-1}$, one obtains $z_1=25(t+2)^2\in(\mathbb{Q}^\times)^2$ and $z_2=(t+2)^2(2-6\sqrt{-1})^2\in(\mathbb{Q}(\sqrt{-1})^\times)^2$. Therefore, $\gamma_{j(E_t)}$ has splitting type $(1,1,2,2)$ for every $t\in\mathbb{Q}$.

(ii) Consider the elliptic curve $E: y^3=x^3+4x^2+x+1$ over $K=\mathbb{F}_5$ with $j(E)=2$. The curve has trivial 2-torsion over $\mathbb{F}_5$ and full 2-torsion over $\mathbb{F}_5(E[2])\simeq\mathbb{F}_5[x]/(x^3+4x^2+x+1)=\mathbb{F}_{5^3}$. The discriminant of E is 1, thus the Galois group of x^3+4x^2+x+1 is $\mathfrak{A}_3$. Let $u\in\mathbb{F}_5(E[2])$ be a root of x^3+4x^2+x+1; thus $(u,0)$ has order 2 on E over $\mathbb{F}_5(E[2])$. The roots of x^3+4x^2+x+1 are $e_1=u,e_2=3u^2+4u=\sigma(e_1)=e_1^5$ and $e_3=2u^2+1=\sigma^2(e_1)=e_1^{5^2}$, where $\langle\sigma\rangle=\mathfrak{A}_3$ and $\sigma: t\mapsto t^5$ is the Frobenius on $\mathbb{F}_{5^3}$. Direct computations show $z_1=3u^2+3u+1=(3u^2+u+1)^2$, a square in $\mathbb{F}_{5^3}$, and so for $z_2=\sigma(z_1)=z_1^5$ and $z_3=\sigma^2(u)=z_1^{5^2}$. By Theorem 6.6.5, the Edwards curves birationally equivalent to E are defined over a cubic extension of $\mathbb{F}_5$: $K_1\simeq K_2\simeq K_3\simeq\mathbb{F}_5(E[2])$. Computations give $(d_1,1/d_1)=(3u^2{+}4,2u^2{+}2u{+}2),(d_2,1/d_2)=(4u,u^2{+}4u{+}1)=(d_1^5,(1/d_1)^5)$ and $(d_3,1/d_3)=(2u^2+u,2u^2+4u+3)=(d_1^{5^2},(1/d_1)^{5^2})$, so that d_1,d_2,d_3 are conjugates in $\mathbb{F}_{5^3}$ with minimal polynomial $m_1(T)=T^3+T^2+T+4\in\mathbb{F}_5[T]$ and $1/d_1,1/d_2,1/d_3$ are conjugates in $\mathbb{F}_5(E[2])$ with minimal polynomial $m_2(T)=T^3+4T^2+4T+4\in\mathbb{F}_5[T]$, and it is readily checked that $\gamma_{j(E)}(T)=16m_1(T)m_2(T)\in\mathbb{F}_5[T]$. The polynomial $f_{j(E)}(T)=T^3+T+1$ is irreducible over $\mathbb{F}_5[T]$ and splits over $\mathbb{F}_5(E[2])[T]$ as $(T-\alpha_1)(T-\alpha_2)(T-\alpha_3)$ with $\alpha_1=2u{+}1,\alpha_2=u^2{+}3u{+}1=\alpha_1^5$ and $\alpha_3=4u^2+3=\alpha_1^{5^2}$. These elements define the (conjugate) quadratic polynomials $T^2-\alpha_1T+1,T^2-\alpha_2T+1,T^2-\alpha_3T+1$ with roots $\{d_1,1/d_1\},\{d_2,1/d_2\},\{d_3,1/d_3\}$. Let E_d

Splitting type	$\Delta(E)$	$E[2](K)$	properties of z_1, z_2, z_3
$(1,1,1,1,1,1)$		$\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$	$z_1, z_2, z_3 \in (K^\times)^2$
$(1,1,1,1,2)$		$\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$	$z_1, z_2 \in (K^\times)^2, z_3 \notin (K^\times)^2$
$(1,1,2,2)$		$\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$	$z_1 \in (K^\times)^2, z_2, z_3 \notin (K^\times)^2$
		$\mathbb{Z}/2\mathbb{Z}$	$z_1 \in (K^\times)^2, z_2 \in (K^\times)^2$
$(1,1,4)$		$\mathbb{Z}/2\mathbb{Z}$	$z_1 \in (K^\times)^2, z_2 \notin (K^\times)^2$
$(2,2,2)$		$\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$	$z_1, z_2, z_3 \notin (K^\times)^2$
$(2,4)$		$\mathbb{Z}/2\mathbb{Z}$	$z_1 \notin (K^\times)^2, z_2 \notin (K^\times)^2$
$(3,3)$	$\in (K^\times)^2$	$\{\mathcal{O}\}$	$z_1 \in (K(E[2])^\times)^2$
(6)	$\in (K^\times)^2$	$\{\mathcal{O}\}$	$z_1 \notin (K(E[2])^\times)^2$
	$\notin (K^\times)^2$	$\{\mathcal{O}\}$	$z_1 \in (K(E[2])^\times)^2$
	$\notin (K^\times)^2$	$\{\mathcal{O}\}$	$z_1 \notin (K(E[2])^\times)^2$

Table 6.1: Splitting type of $\gamma_{j(E)}$

be the elliptic curve from (6.2) (with $a = 1$). A direct verification shows that the curves E_d with $d \in \{d_1, d_2, d_3, 1/d_1, 1/d_2, 1/d_3\}$ all have j-invariant 2.

Table 6.1 summarizes the splitting types of $\gamma_{j(E)}$ in the different cases. Note since d and $1/d$ appear in pairs, the polynomial cannot have an odd number of linear factors. The splitting type $(1,1,1,1,1,1)$ occurs for example for the curve with label 5525.5-b9 in [LMF20a] defined over $\mathbb{Q}(\sqrt{-1})$.

The Galois action on Edwards curves can be seen naturally as arising from that on the points of E. The group G_K acts on the (4-torsion) points of E as follows: for $\sigma \in G_K$ and $P \in E(\overline{K})$ let $P^\sigma := (\sigma(x), \sigma(y))$ denote the image of P under the action of σ. For a point of order 4, its image under $\sigma \in G_K$ is again a point of order 4. By acting on $d \in \overline{K}$, G_K acts on $\mathrm{Edw}_{\overline{K}}(j)$. The next lemma shows a compatibility between these Galois actions.

Lemma 6.7.4. *Let $P \in E_4(\overline{K})$ and $\mathcal{E}_{d_P}$ be its associated Edwards curve. For every $\sigma \in G_K$, we have ${}^\sigma(\mathcal{E}_{d_P}) = \mathcal{E}_{d_{P^\sigma}}$, i.e. the map $d : E_4(\overline{K}) \to \overline{K}, P \mapsto d_P$ is G_K-equivariant.*

Proof. This follows from $\sigma(d_P) = 1 - \frac{4(\sigma(x_P)-\sigma(x_{2P}))^3}{[\sigma(y_P)+\frac{1}{2}(a_1\sigma(x_P)+a_3)]^2} = d_{P^\sigma}$. □

Theorem 6.7.5. *Let E be an elliptic curve over K and let $S, T \in E(\overline{K})$ be points of order 4. Let $\sigma \in \mathrm{Gal}(\overline{K}/K)$. The following hold:*

(i) if $S = \pm T^\sigma$ then $d_S = \sigma(d_T)$

(ii) if $S \neq \pm T^\sigma$ and $2S = 2T^\sigma$ then $d_S = 1/\sigma(d_T)$

(iii) if $2S \neq 2T^\sigma$ then $d_S = \left(\frac{1-\epsilon\sigma(d_T)^{1/4}}{1+\epsilon\sigma(d_T)^{1/4}}\right)^4$ for some $\epsilon \in \overline{K}$ with $\epsilon^4 = 1$.

Moreover, if $j(E) \neq 0, 1728$, then the implications in (i), (ii), (iii) *are equivalences.*

Proof. Since the characteristic of K is not 2, E may be given by $y^2 = (x - e_1)(x - e_2)(x - e_3)$ with $e_1, e_2, e_3 \in \overline{K}$ two by two distinct.

(i) If $S = \pm T^\sigma$ then $d_S = d_{T^\sigma} = \sigma(d_T)$ by Remark 6.6.4 and Lemma 6.7.4.
(ii) Assume $S \neq \pm T^\sigma$ and $2S = 2T^\sigma = (e_i, 0)$ for $i \in \{1, 2, 3\}$. Thus S, T^σ are among the points $\pm P_i^\pm$ of Lemma 6.6.7 and are not opposite; thus, $S = \pm P_i^+$ and $T^\sigma = \pm P_i^-$ or $S = \pm P_i^-$ and $T^\sigma = \pm P_i^-$. By Equation (6.13) of Lemma 6.6.9 we obtain that d_S and $\sigma(d_T)$ are opposite in each case.
(iii) Assume $2S \neq 2T^\sigma$. Without loss of generality we assume $2S = (e_1, 0)$ and $2T^\sigma = (e_2, 0)$. By Lemma 6.6.7, S equals one of $\pm P_1^\pm$ and T^σ equals one of $\pm P_2^\pm$. We show that

$$d_S^{1/4}(1 - \epsilon\sigma(d_T)^{1/4}) = (1 + \epsilon\sigma(d_T)^{1/4}) \tag{6.19}$$

for some fourth root of unity ϵ. We first treat the case $S = \pm P_1^+$ and $T^\sigma = \pm P_2^-$. By Lemma 6.6.9,

$$d_S = \left(\frac{\sqrt{e_{12}} + \sqrt{e_{13}}}{\sqrt{e_{32}}}\right)^4 , \; d_{T^\sigma} = \sigma(d_T) = \left(\frac{\sqrt{e_{21}} - \sqrt{e_{23}}}{\sqrt{e_{31}}}\right)^4 .$$

The fourth roots of unity of d_S and $\sigma(d_T)$ are $\zeta\frac{\sqrt{e_{12}}-\sqrt{e_{13}}}{\sqrt{e_{32}}}$ and $d_{T^\sigma} = \zeta'\frac{\sqrt{e_{21}}+\sqrt{e_{23}}}{\sqrt{e_{31}}}$ for some fourth roots of unity ζ, ζ'. In the formula (6.19) we are about to show, we can incorporate ζ' inside ε and obtain another fourth root of unity (which we still call ε). Equation (6.19) is equivalent to

$$\zeta\frac{\sqrt{e_{12}} - \sqrt{e_{13}}}{\sqrt{e_{32}}} = \frac{\sqrt{e_{31}} - \varepsilon\sqrt{e_{21}} - \varepsilon\sqrt{e_{23}}}{\sqrt{e_{31}} + \varepsilon\sqrt{e_{21}} + \varepsilon\sqrt{e_{23}}} ,$$

which is again equivalent to

$$\sqrt{e_{12}}\sqrt{e_{13}}(i\zeta - i\zeta\varepsilon) + \sqrt{e_{12}}\sqrt{e_{23}}(\zeta\varepsilon - \varepsilon) + \sqrt{e_{13}}\sqrt{e_{23}}(1 - \zeta\varepsilon) + i\zeta\varepsilon\sqrt{e_{12}} - i\zeta\sqrt{e_{13}} + i\varepsilon\sqrt{e_{23}} = 0 ,$$

where $i = \sqrt{-1}$. In view of the relation $e_{12} - e_{13} + e_{23} = 0$, this equation has the solution $(\zeta, \varepsilon) = (1, 1)$. The remaining cases are proved similarly: the case $(S, T) = (\pm P_1^-, \pm P_2^+)$ is direct by exchanging S and T; in the case $(S, T) = (\pm P_1^+, \pm P_2^+)$, Equation (6.19) becomes

$$\sqrt{e_{12}}\sqrt{e_{13}}(i\zeta - i\zeta\varepsilon) + \sqrt{e_{12}}\sqrt{e_{23}}(-\zeta\varepsilon - \varepsilon) + \sqrt{e_{13}}\sqrt{e_{23}}(1 + \zeta\varepsilon) + i\zeta\varepsilon\sqrt{e_{12}} - i\zeta\sqrt{e_{13}} - i\varepsilon\sqrt{e_{23}} = 0$$

which has the solution $(\zeta, \varepsilon) = (-1, 1)$. Finally, if $(S, T) = (\pm P_1^-, \pm P_2^-)$, Equation (6.19) becomes equivalent to

$$\sqrt{e_{12}}\sqrt{e_{13}}(i\zeta + i\zeta\varepsilon) + \sqrt{e_{12}}\sqrt{e_{23}}(\zeta\varepsilon - \varepsilon) + \sqrt{e_{13}}\sqrt{e_{23}}(1 + \zeta\varepsilon) + i\zeta\varepsilon\sqrt{e_{12}} + i\zeta\sqrt{e_{13}} + i\varepsilon\sqrt{e_{23}} = 0$$

which has the solution $(\zeta, \varepsilon) = (1, -1)$.

Let us now show the converse statements assuming in addition that $j(E) \notin \{0, 1728\}$. For the converse of *(i)*, assume $d_S = \sigma(d_T)$ and by contradiction that $S \neq \pm T^\sigma$. Then either $2S = 2T^\sigma$ and by *(ii)*, $d_S = \sigma(d_T) = 1/\sigma(d_T)$, or $2S \neq 2T^\sigma$ and by *(iii)*,

$$d_S = \sigma(d_T) = \left(\frac{1 - \epsilon\sigma(d_T)^{1/4}}{1 + \epsilon\sigma(d_T)^{1/4}}\right)^4 , \; \epsilon^4 = 1. \tag{6.20}$$

For the converse of *(ii)*, assume that $d_S = 1/\sigma(d_T)$ and by contradiction that $S = \pm T^\sigma$ or $2S \neq 2T^\sigma$. In the first case, we obtain by *(i)* that $d_S = 1/\sigma(d_T) = \sigma(d_T)$; in the second case,

we obtain by *(iii)* that

$$d_S = \frac{1}{\sigma(d_T)} = \left(\frac{1-\epsilon\sigma(d_T)^{1/4}}{1+\epsilon\sigma(d_T)^{1/4}}\right)^4 \ , \ \epsilon^4 = 1. \tag{6.21}$$

For the converse of *(iii)*, assume that $d_S = [(1-\epsilon\sigma(d_T)^{1/4})/(1+\epsilon\sigma(d_T)^{1/4})]^4$ for some fourth root of unity ϵ. By contradiction, if $2S = 2T^\sigma$, then either $S = \pm T^\sigma$ and by (i), we obtain the same equation as Equation (6.20); or $S \neq \pm T^\sigma$ and in this case by (ii), we obtain the same equation as (6.21).

Putting all together, it is enough to solve the equations $\sigma(d_T) = 1/\sigma(d_T)$, Equation (6.20) and (6.21) for $\sigma(d_T)$ and obtain contradictions when $j(E) \in \{0, 1728\}$. Clearly, $\sigma = 1/\sigma(d_T)$ gives $\sigma(d_T)^2 = 1$, i.e. $\sigma(d_T) = \pm 1$. The case $\sigma(d_T) = 1$ is excluded because E would not be elliptic; the case $\sigma(d_T) = -1$ is excluded because then $j(E) = 1728$. Changing ϵ to $-\epsilon$ and inverting (6.21) we see that it is equivalent to (6.20), whence it is enough to solve (6.20). By Equation (6.20) there is a fourth root of unity ϵ' such that $\epsilon'\sigma(d_T)^{1/4} = (1-\epsilon\sigma(d_T)^{1/4})/(1+\epsilon\sigma(d_T)^{1/4})$ which is equivalent to the quadratic equation

$$\epsilon\epsilon'\sigma(d_T)^{1/2} + (\epsilon+\epsilon')\sigma(d_T)^{1/4} - 1 = 0 \tag{6.22}$$

in $\sigma(d_T)^{1/4}$, with discriminant $\Delta(\epsilon,\epsilon') = \epsilon^2 + \epsilon'^2 + 6\epsilon\epsilon'$. Since Δ is symmetric in ϵ and ϵ', we reduce the number of cases for ϵ, ϵ' by two. We show that each case gives a contradiction. If $\epsilon = -\epsilon'$ we obtain from (6.22) that $\sigma(d_T)^{1/2} = 1/(\epsilon\epsilon')$ giving $\sigma(d_T) = 1$, contradicting that E is elliptic. If $\epsilon = \epsilon'$ then $\Delta = 8\epsilon^2 \in \{8, -8\}$, according to $\epsilon \in \{\pm 1\}$ or $\epsilon \in \{\pm i\}$. If $\epsilon = \epsilon' \in \{\pm 1\}$ then $\sigma(d_T)^{1/4} = (\mp 2 \pm 2\sqrt{2})/2 = \mp 1 \pm \sqrt{2}$ and $\sigma(d_T) = \mp 17 \pm 12\sqrt{12}$, which gives $j(E) = 1728$ (see Examples 6.5.1 and 6.6.10) and is excluded. Similarly, $\epsilon = \epsilon' \in \{\pm i\}$ with $\Delta = -8$ gives $j(E) = 1728$. Finally, if $\epsilon = i \cdot \epsilon'$ then $\Delta \in \{6i, -6i\}$ according to $\epsilon' = \pm i$ or $\epsilon' = \pm 1$. If $\epsilon' = \pm i$ then $\sigma(d_T)^{1/4} = (\mp 1 \mp i \pm \sqrt{6i})/(\pm 2i)$, and $\epsilon' = \pm 1$ gives $\sigma(d_T)^{1/4} = (\mp 1 \mp i \pm i\sqrt{6i})/(\pm 2i)$. Both cases give $\sigma(d_T) = -7 \pm 4\sqrt{3}$ which implies $j(E) = 0$, by Examples 6.5.1 and 6.6.10. □

Remark 6.7.6. *(i)* Theorem 6.7.5 is a stronger version of the result of Ahmadi-Granger [AG12, Theorem 4.2] recalled in Proposition 6.3.5. Theorem 6.7.5 explicitly describes how these formulae for d_P arise from the underlying geometry of the considered points on E. Moreover, the proof technique of Ahmadi-Granger relies on a completely different approach.

(ii) Let E/K be an elliptic curve and P, Q be generators for $E(\overline{K})[4] \simeq \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}$. Assume that $2P = (e_1, 0)$ and $2Q = (e_2, 0)$; then $2(P+Q) = (e_3, 0)$. For $i \in \{1, 2, 3\}$, let K_i/K be the extension considered in Theorem 6.6.5. By Theorem 6.7.5, the curves $\mathcal{E}_{d_P}$ and $\mathcal{E}_{1/d_P} = \mathcal{E}_{d_{P+2Q}}$ associated to $\pm\{P, P+2Q\}$ doubling to $(e_1, 0)$ are defined over K_1, those associated to $\pm\{Q, Q+2P\}$ doubling to $(e_2, 0)$ are defined over K_2, and those associated to $\pm\{P+Q, P-Q\}$ doubling to $(e_3, 0)$ are defined over K_3.

6.8 More on isogenies of Edwards curves

Isogenies between Edwards curves have been studied independently, see for example [MS16]. In this section we extend some known results.

6.8.1 Isogenies between curves in $\mathrm{Edw}_{\overline{K}}(j)$

Consider the set $\mathrm{Edw}_{\overline{K}}(j)$ for a fixed $j \in K$. We study isogenies between curves in this set. In [AG12, Section 3], Ahmadi and Granger describe 4-isogenies between Edwards curves (in fact, many results in [AG12] appear to extend to arbitrary fields instead of finite fields). For $d \in K \setminus \{0,1\}$, let $\mathcal{L}_d$ denote the Legendre curve given by $y^2 = x(x-1)(x-d)$. By [AG12, Theorem 3.1], there is an explicit 2-isogeny

$$\phi : \mathcal{E}_d \to \mathcal{L}_d \, , (x,y) \mapsto \left(\frac{1}{x^2}, \frac{y(d-1)}{x(1-y^2)} \right)$$

with dual isogeny

$$\phi^\vee : \mathcal{L}_d \to \mathcal{E}_d, (x,y) \mapsto \left(\frac{2y}{d-x^2}, \frac{y^2 - x^2(1-d)}{y^2 + x^2(1-d)} \right) \, .$$

Moreover, $\mathcal{L}_d \to \mathcal{L}_{1/d}$ given by $(x,y) \mapsto (x/d, y/d^{3/2})$ is an isomorphism, defined over $K(\sqrt{d})$ (see [AG12, Section 3.2] or [Sil09, Chapter III, §1, Proposition 1.7]). By composition, we obtain a 4-isogeny over $K(\sqrt{d})$ between $\mathcal{E}_d$ and $\mathcal{E}_{1/d}$. We remark that this is not true for all the curves in $\mathrm{Edw}_{\overline{K}}(j)$ but only for those corresponding to points of order 4 doubling to the same point of order 2 (see Theorem 6.7.5). Indeed, there are isomorphisms between $\mathcal{L}_d$ and $\mathcal{L}_{d'}$ for $d' \in \{d, 1-d, 1/d, 1-1/d, 1/(1-d), d/(d-1)\}$ (see [AG12]); the intersection of this set $\Sigma(d)$ is $\{d, 1/d\}$.

The curves $\mathcal{E}_d$ and $\mathcal{E}_{1/d}$ are quadratic twists of each other. An isomorphism $\mathcal{E}_d \to \mathcal{E}_{1/d}$ over $K(\sqrt{d})$ is defined by $(X,Y) \mapsto (X/\sqrt{d}, Y)$. In conclusion, the pairs of curves defined over K_1, K_2, K_3 in Theorem 6.6.5 are 4-isogenous and isomorphic over the quadratic extensions $K_1(\sqrt{d_1}), K_2(\sqrt{d_2})$ and $K_3(\sqrt{d_3})$.

6.8.2 General 2-isogenies

For an elliptic curve E/K, a finite field extension L/K, an integer $\ell \geq 2$ and $a, d \in L$, we say that E is *Edwards-ℓ-isogenous* over L if there is an elliptic curve E'/L, an ℓ-isogeny $\phi : E(L) \to E'(L)$ defined over L (or possibly an extension of L), and a birational map $\psi : E'(L) \to \mathcal{E}_d^a(L)$ where $\mathcal{E}_d^a$ is a twisted Edwards curve defined over L. Here we have denoted by $E(L)$ the group of L-points of E base-changed to L.

The following theorem generalizes [BBJ+08, Theorem 5.1], which is part of (1)(i) now. It gives explicit 2-isogenies between E and a twised Edwards curve when E has full K-rational 2-torsion (compare also with Remark 6.6.6 (ii)).

Theorem 6.8.1. *Let E/K be an elliptic curve and let $E(\overline{K})[2] = \{\mathcal{O}, (e_1,0), (e_2,0), (e_3,0)\}$ be its 2-torsion subgroup. Then*
(1) E is Edwards-2-isogenous over $K^{(1)} = K(e_{21}, e_{31})$, $K^{(2)} = K(e_{12}, e_{32})$, $K^{(3)} = K(e_{13}, e_{23})$, respectively, to the twisted Edwards curves

$$\begin{aligned}
\mathcal{E}^{(1)} &:= \mathcal{E}_{4e_{21}}^{4e_{31}} \, , \; \mathcal{E}'^{(1)} := \mathcal{E}_{4e_{31}}^{4e_{21}} \\
\mathcal{E}^{(2)} &:= \mathcal{E}_{4e_{32}}^{4e_{12}} \, , \; \mathcal{E}'^{(2)} := \mathcal{E}_{4e_{12}}^{4e_{32}} \\
\mathcal{E}^{(3)} &:= \mathcal{E}_{4e_{23}}^{4e_{13}} \, , \; \mathcal{E}'^{(3)} := \mathcal{E}_{4e_{13}}^{4e_{23}}
\end{aligned}$$

defined over $K^{(1)}, K^{(2)}, K^{(3)}$, respectively. In particular:

(i) if $E(K)[2] \simeq \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ then $K^{(i)} = K$ for all $i \in \{1,2,3\}$,

(ii) if $E(K)[2] \simeq \mathbb{Z}/2\mathbb{Z}$ then $K^{(i)} = K(\sqrt{\Delta(E)})$ for all $i \in \{1,2,3\}$,

(iii) if $E(K)[2] \simeq \{\mathcal{O}\}$ then $[K^{(i)} : K] \leq 6$ for all $i \in \{1,2,3\}$.

(2) The twisted Edwards curves $\{\mathcal{E}^{(i)}\}_{i\in\{1,2,3\}}$ and $\{\mathcal{E}'^{(i)}\}_{i\in\{1,2,3\}}$ are 4-isogenous over $K(E[2])$.

Proof. We start by proving the existence of the 2-isogenies following the proof of [BBJ+08, Theorem 5.1]. Given a 2-torsion point on E, which is taken to be $(0,0)$, the proof constructs from this point a twisted Edwards curve 2-isogenous to E. We repeat this for all the points $(e_1,0),(e_2,0),(e_3,0)$ in $E(\overline{K})[2]$.

Let us start with $(e_1,0)$. In order to apply Theorem [BBJ+08, Theorem 5.1], we move $(e_1,0)$ to the origin by letting $x' = x - e_1; y' = y$. This transformation takes E to the elliptic curve with points (x',y') satisfying $y'^2 = x'(x'-e_{21})(x'-e_{31})$, i.e. $y'^2 = x'^3-(e_{21}+e_{31})x'^2+(e_{21}e_{31})x'$. By the proof of [BBJ+08, Theorem 5.1], there is a 2-isogeny to the elliptic curve $E^{(1)}$ given by $y'^2 = x'^3 + 2(e_{21}+e_{31})x'^2 + (e_{21}-e_{31})^2x'$, and using $e_{21}-e_{31} = e_{23}$, this is $y'^2 = x'^3 + 2(e_{21}+e_{31})x'^2 + e_{23}^2x'$, which is isomorphic to the Montgomery curve $\mathcal{M}^{(1)} := \mathcal{M}_{A_1,B_1}$ with $A_1 = 2(e_{21}+e_{31})/e_{23}$ and $B_1 = 1/e_{23}$, and the isogeny has kernel $\{\mathcal{O},(e_1,0)\}$ (and thereby we denote $\mathcal{M}^{(1)}$ as $E/\langle(e_1,0)\rangle$.) By [BBJ+08, Theorem 3.2], the latter is birationally equivalent to the twisted Edwards curve $\mathcal{E}^{(1)} := \mathcal{E}^{a_1}_{d_1}$ with $d_1 = (A_1-2)/B_1 = 4e_{21}$ and $a_1 = (A_1+2)/B_1 = 4e_{31}$. The curve $\mathcal{E}^{a_1}_{d_1}$ is defined over $K^{(1)}$.

The existence of the other 2-isogenies is obtained by repeating the construction for $(e_2,0)$ and $(e_3,0)$, respectively, and the isogenies are those having as kernel $\{\mathcal{O},(e_2,0)\}$ and $\{\mathcal{O},(e_3,0)\}$. Direct computations give $\mathcal{M}^{(2)} = \mathcal{M}_{A_2,B_2}$ with $A_2 = 2(e_{12}+e_{32})/e_{13}$ and $B_2 = 1/(e_{13})$, and $\mathcal{M}^{(3)} = \mathcal{M}_{A_3,B_3}$ with $A_3 = 2(e_{13}+e_{23})/e_{12}$ and $B_3 = 1/(e_{12})$, respectively. The twisted Edwards curves $\mathcal{E}^{(2)} = \mathcal{E}^{d_2}_{a_2}$ and $\mathcal{E}^{(3)} = \mathcal{E}^{d_3}_{a_3}$ then follow as previously with $d_k = (A_k-2)/B_k$ and $a_k = (A_k+2)/B_k$ for $k \in \{2,3\}$.

The second part of the statement on the degrees of $\{K^{(i)}\}_i$ is similar to the corresponding part in the proof of Theorem 6.6.5. To prove (2), it is enough to compose the 2-isogenies with their duals; let $i,j \in \{1,2,3\}$ be distinct; $\phi_i : E \to E^{(i)}$ and $\phi_j : E \to E^{(j)}$ the 2-isogenies from (1) inducing the 2-isogenies $\phi'_i : E \to \mathcal{E}^{(i)}$ and $\phi'_j : E \to \mathcal{E}^{(j)}$. Then $\phi'_j\phi'^{\vee}_i$ is a 4-isogeny between $\mathcal{E}^{(i)}$ and $\mathcal{E}^{(j)}$. □

In compact form the twisted Edwards curves in Theorem 6.8.1 are given by

$$\mathcal{E}^{(k)} = \mathcal{E}^{4e_{l'k}}_{4e_{lk}}\ ,\ \mathcal{E}'^{(k)} = \mathcal{E}^{4e_{lk}}_{4e_{l'k}}\ ,\ k \in \{1,2,3\}\ ,\ l \neq l' \in \{1,2,3\} \setminus \{k\}$$

defined over the fields $K^{(k)} = K(e_{lk}, e_{l'k})$, respectively. Combining with [AG12, Theorem 3.2], there are 4-isogenies between these twisted Edwards curves to the Legendre curves, i.e.

$$\mathcal{E}^{(k)} \to \mathcal{L}_{\lambda_k}\ ,\ \mathcal{E}'^{(k)} \to \mathcal{L}_{1/\lambda_k}\ ,\ \lambda_k = \frac{e_{lk}}{e_{l'k}}\ .$$

These isogenies are defined over quadratic extensions of $K^{(k)}$. Running with k through $\{1,2,3\}$ the set of $\{\lambda_k, 1/\lambda_k\}$ is exactly $\left\{\frac{e_{21}}{e_{31}}, \frac{e_{31}}{e_{21}}, \frac{e_{32}}{e_{12}}, \frac{e_{12}}{e_{32}}, \frac{e_{23}}{e_{13}}, \frac{e_{13}}{e_{23}}\right\}$, which is precisely the $\overline{K}$-isomorphism class of Legendre curves isomorphic to E.

6.9 Statistics for the rank in the Edwards family

In this section, we study properties of the group of rational points for the elliptic curves over $K = \mathbb{Q}$ in the family

$$E_{a,d} : y^2 = x^3 + 2(a+d)x^2 + (a-d)^2x \quad , \ a, d \in \mathbb{Z} \setminus \{0, 1\} \ , \ a \neq d \ . \tag{6.23}$$

We call this the *Edwards family* in view of the birational equivalence of $E_{a,d}$ to the twisted Edwards curve $\mathcal{E}_d^a$. The group $E_{a,d}(\mathbb{Q})$ of $\mathbb{Q}$-rational points is finitely generated by Mordell's Theorem, i.e. $E_{a,d}(\mathbb{Q}) \simeq \mathbb{Z}^{r(E_{a,d})} \oplus E_{a,d}(\mathbb{Q})_{\mathrm{tors}}$ for some non-negative integer $r(E_{a,d})$ (the rank of $E_{a,d}$) and a finite abelian (torsion) group $E_{a,d}(\mathbb{Q})_{\mathrm{tors}}$. In Section 6.9.1, we explicitly describe the torsion subgroup for the one-parameter family $\{E_{1,d}\}_d$. As for the rank of $E_{1,d}$, we provide related statistical data in Sections 6.9.4 and 6.9.5.

6.9.1 Torsion subgroup of E_d

We consider the case $a = 1$ and write E_d instead of $E_{1,d}$ for the curve in Equation (6.23). The following theorem describes the torsion subgroup of E_d. For the first part, K is an arbitrary field of characteristic not 2, for the second part we specialize to $K = \mathbb{Q}$.

Theorem 6.9.1. *Let K be a field of characteristic not 2. Let $d \in K \setminus \{0, 1\}$. Let E_d be defined over K by $y^2 = x^3 + 2(1+d)x^2 + (1-d)^2$.*

(i) The points $(0, 0) \in E_d(K)$ and $\pm P_2 := (-(1 \pm \sqrt{d})^2, 0)$ in $E_d(K(\sqrt{d}))$ have exact order 2. In particular, E_d has full K-rational 2-torsion if and only if d is a square in K.
The points $\pm P_4^+ := (1-d, \pm 2(1-d))$ in $E_d(K)$ and $\pm P_4^- := (d-1, \pm 2(1-d)\sqrt{d})$ in $E_d(K(\sqrt{d}))$ have exact order 4, and $2(\pm P_4^+) = 2(\pm P_4^-) = (0, 0)$.
Let $t = \sqrt{1 + \sqrt{1-d}}$. The points $\pm P_8^+ := ((t-1)(t+1)^3, \pm 2t(t-1)(t+1)^3)$ in $E_d(K(t))$ and $\pm P_8^- := ((t+1)(t-1)^3, \pm 2t(t+1)(t-1)^3)$ in $E_d(K(t))$ have exact order 8, and $2(\pm P_8^+) = 2(\pm P_8^-) = \pm P_4^+$.

(ii) For $K = \mathbb{Q}$ and $d \in \mathbb{Z} \setminus \{0, 1\}$, it holds

$$E_d(\mathbb{Q})_{\mathrm{tors}} \simeq \begin{cases} \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z} & \text{if } d \text{ is a square} \\ \mathbb{Z}/8\mathbb{Z} & \text{if } d \text{ is not a square and } d = 1 - (t^2-1)^2, \text{ for some } t \in \mathbb{Z} \\ \mathbb{Z}/4\mathbb{Z} & \text{if } d \text{ is not a square and } d \neq 1 - (t^2-1)^2, \text{ for all } t \in \mathbb{Z} \end{cases}$$

Proof. (i) The cubic in x defining E_d factors as $x(x^2 + 2(1+d)x + (1-d)^2)$, from which it is immediate to determine the complete 2-torsion of E_d; all roots lie in K if and only if d is a square in K. Moreover, it is readily checked that $(\pm 1, 0) \in \mathcal{E}_d(K)$ of order 4 are mapped to $\pm P_4^+ \in E_d(K)$ under the birational equivalence (6.2), and the torsion is preserved under birational equivalence. The points $\pm P_4^-$ are constructed easily (see e.g. Lemma 6.6.7).

For the points of order 8 doubling to $\pm P_4^+$, we make use of the duplication map for E_d. For $Q \in E_d(K)$, we have:

$$x(2Q) = \frac{x(Q)^4 - b_4 x(Q)^2 - 2b_6 x(Q) - b_8}{4x(Q)^3 + b_2 x(Q)^2 + 2b_4 x(Q) + b_6} \ ,$$

with b_2, b_4, b_6, b_8 the b-invariants given by $b_2 = 8(d+1), b_4 = 2(1-d)^2, b_6 = 0, b_8 = -(1-d)^4$. Let $t = \sqrt{1+\sqrt{1-d}}$ and Q be a point of order 8 on E_d doubling to $\pm P_4^+$. Writing out the equation $x(2Q) = x(\pm P_4) = 1-d$ amounts to compute $x(Q)$ as the roots of the quartic polynomial $x^4 - 2(1-d)^2x^2 + (1-d)^4 - (1-d)[4x^3 + 8(d+1)x^2 + 4(1-d)^2x]$. Putting $u = 1-d$, its four roots can be written as

$$x(Q) \in \left\{ u + 2\sqrt{u} \pm 2\sqrt{u(1+\sqrt{u})}, u - 2\sqrt{u} \pm 2\sqrt{u(1-\sqrt{u})} \right\} .$$

Letting $s = \sqrt{u}$ gives $t = \sqrt{1+s}$, i.e. $s = t^2 - 1$. The first pair of roots [2] $x(Q)$ is then equal to

$$\begin{aligned} s^2 + 2s \pm 2st &= s(s + 2 \pm 2t) = s(t^2 \pm 2t + 1) \\ &= s(t \pm 1)^2 = (t-1)(t+1)(t \pm 1)^2 = (t \mp 1)(t \pm 1)^3 . \end{aligned}$$

The corresponding y-coordinates $y(Q)$ are computed as solutions to $y(Q)^2 = x(Q)^3 + 2(1+d)x(Q)^2 + (1-d)^2x(Q) = x(Q)^3 + 2(2-u)x(Q)^2 + u^2x(Q)$, or equivalently, using $s^2 = u = (t^2-1)^2$,

$$y(Q)^2 = x(Q)^3 + 4x(Q)^2 - 2(t^2-1)^2x(Q)^2 + (t^2-1)^4x(Q) .$$

Plugging $x(Q) = (t \mp 1)(t \pm 1)^3$ gives the two equations

$$\begin{aligned} y(Q)^2 &= (t \mp 1)^3(t \pm 1)^9 + 4(t \mp 1)^2(t \pm 1)^6 - 2(t \mp 1)^4(t \pm 1)^8 + (t \mp 1)^5(t \pm 1)^7 \\ &= (t \mp 1)^2(t \pm 1)^6 \cdot \big((t \mp 1)(t \pm 1)^3 + 4 - 2(t \mp 1)^2(t \pm 1)^2 + (t \mp 1)^3(t \pm 1)\big) . \end{aligned}$$

A direct calculation gives that the last factor equals $(t \mp 1)(t \pm 1)^3 + 4 - 2(t \mp 1)^2(t \pm 1)^2 + (t \mp 1)^3(t \pm 1) = 4t^2$. Therefore, $y(Q)^2 = 4t^2(t \mp 1)^2(t \pm 1)^6$ and the two solutions for $y(Q)$ are $y(Q) = \pm 2t(t \mp 1)(t \pm 1)^3$. Thus the points $P_8^{\pm} = ((t \mp 1)(t \pm 1)^3, 2t(t \mp 1)(t \pm 1)^3)$ and their opposites $-P_8^{\pm} = ((t \mp 1)(t \pm 1)^3, -2t(t \mp 1)(t \pm 1)^3)$ have order 8 on $E_d(K(t))$ and $2(\pm P_8^{\pm}) = \pm P_4^+$.

(ii) Let $K = \mathbb{Q}$ and $d \in \mathbb{Z}$. The group $E_d(\mathbb{Q})_{\mathrm{tors}}$ injects into $E_d(\mathbb{F}_p)$ for every prime $p \geq 3$ of good reduction for E_d. Let $p \geq 3$ be a prime number such that $d \equiv -1 \pmod{p}$, which exists by Dirichlet's Theorem on primes in arithmetic progression. Then $E_d/\mathbb{F}_p$ is non-singular and given by $y^2 = x^3 + 4x$. We count the number of points of $E_d(\mathbb{F}_p)$ explicitly, using [Was03, Theorem 4.23].

Assume first that $d \in \mathbb{Z}$ is a square different from $0, 1$. Then d is not of the form $1-(t^2-1)^2$ with $t \in \mathbb{Z}$ (as otherwise with $d = e^2$, say, one has $e^2 + (t^2-1)^2 = 1$ giving $d = e^2 = 0$ or $d = e^2 = 1$, which are excluded). In particular, the points of order 8 from (i) are not $\mathbb{Q}$-rational. Since $d \equiv -1 \pmod{p}$ and d is a square, it is a square modulo p and thus -1 is a square modulo p, which implies $p \equiv 1 \pmod{4}$. We write $p = a^2 + b^2$ with $a, b \in \mathbb{Z}$ such that b is even and $a + b \equiv 1 \pmod{4}$. Applying [Was03, Theorem 4.23] and noting that -4 is a fourth power modulo p (the rational quartic residue symbol of -4 equals 1), one has $\#E_d(\mathbb{F}_p) = p + 1 - 2a$. We now show that this quantity is always a multiple of 8. In the given case, p can only be 1 or 5 modulo 8. We work out the details for $p \equiv 1 \pmod{8}$; the case $p \equiv 5 \pmod{8}$ is very

[2] the last pair of roots corresponds to points of order 8 doubling to $\pm P_4^-$, which are not part of the statement. Their determination goes similarly.

similar. Since b is even, we must have $a^2 \equiv 1 \pmod 8$ or $a^2 \equiv 5 \pmod 8$, the latter being excluded as 5 is not a square modulo 8. But $a^2 \equiv 1 \pmod 8$ implies a is one of $\{\pm 1, \pm 3\}$ modulo 8, and the cases $a \in \{-1, 3\}$ are excluded in view of $a + b \equiv 1 \pmod 4$ and $b \equiv 0 \pmod 2$. In both cases left, one therefore has $p + 1 - 2a \equiv 1 + 1 - 2 \equiv 0 \pmod 8$. Since by (i), $E_d(\mathbb{Q})_{\mathrm{tors}}$ contains $\{\mathcal{O}, (0,0), \pm P_2, \pm P_4^+, \pm P_4^-\} \simeq \langle P_2 \rangle \times \langle P_4^+ \rangle \simeq \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}$, the order is exactly 8 and $E_d(\mathbb{Q})_{\mathrm{tors}} \simeq \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}$.

Assume now that $d \in \mathbb{Z}$ is not a square and that there is $t \in \mathbb{Z}$ such that $d = 1 - (t^2 - 1)^2$. Then the points of order 8 from (i) are $\mathbb{Q}$-rational. Let $p \geq 3$ be a prime such that $d \equiv -1 \pmod p$. It follows that $(t^2 - 1)^2 \equiv 2 \pmod p$; in particular 2 is a square modulo p and thus $p \equiv \pm 1 \pmod 8$. If $p \equiv 1 \pmod 8$ then $p \equiv 1 \pmod 4$, and by [Was03, Theorem 4.23], $\#E_d(\mathbb{F}_p) = p + 1 - 2a$ with $p = a^2 + b^2$ and $a, b \in \mathbb{Z}$ such that $a + b \equiv 1 \pmod 4$ and b even, as before. The same calculations as above show that $\#E_d(\mathbb{F}_p) \equiv 0 \pmod 8$. If $p \equiv -1 \pmod 8$ then $p \geq 7$ and $p \equiv 3 \pmod 4$ and $E_d/\mathbb{F}_p$ is supersingular and has order $p + 1 \equiv 0 \pmod 8$. By (i), $E_d(\mathbb{Q})_{\mathrm{tors}}$ contains the cyclic subgroup $\{\mathcal{O}, (0,0), \pm P_4^+, \pm P_8^+, \pm P_8^-\} = \langle P_8^+ \rangle \simeq \mathbb{Z}/8\mathbb{Z}$, therefore there is an isomorphism $E_d(\mathbb{Q})_{\mathrm{tors}} \simeq \mathbb{Z}/8\mathbb{Z}$.

Assume now that $d \in \mathbb{Z}$ is not a square and not of the form $1 - (t^2 - 1)^2$ with $t \in \mathbb{Z}$. Let $p \geq 3$ be a prime such that $d \equiv -1 \pmod p$. If $p \equiv 1 \pmod 4$, then write $p = a^2 + b^2$ with $a, b \in \mathbb{Z}$ and $a + b \equiv 1 \pmod 4$ and b even, and as above, $\#E_d(\mathbb{F}_p) = p + 1 - 2a \equiv 2 - 2a \pmod 4$. Since $b \equiv 0, 2 \pmod 4$ it follows $a \equiv 1 - b \equiv \pm 1 \pmod 4$, and $\#E_d(\mathbb{F}_p) \equiv 2 - 2a \equiv 0 \pmod 4$. On the other hand, if $p \equiv 3 \pmod 4$ then $E_d/\mathbb{F}_p$ is supersingular and has order $p + 1 \equiv 0 \pmod 4$. In both cases, $E_d(\mathbb{F}_p)$ has order dividing 4 for all but finitely many primes p. By (i), since $E_d(\mathbb{Q})_{\mathrm{tors}}$ contains the cyclic subgroup $\{\mathcal{O}, (0,0), \pm P_4\} = \langle P_4 \rangle$, there must be an isomorphism $E_d(\mathbb{Q})_{\mathrm{tors}} \simeq \mathbb{Z}/4\mathbb{Z}$. □

In particular, if d is not a square then E_d has a unique K-rational point of order 2, in which case, the corresponding Edwards curve has a *complete* addition law (see [BL07, Theorem 2.1(2)]). Note that the second part of the statement is only true for $K = \mathbb{Q}$; the torsion over number fields is well-studied (see e.g. [KM88, KN12] for quadratic fields). In [BBLP13, Section 6], the authors study Edwards curves with large torsion subgroup.

Example 6.9.2. We illustrate three cases of Theorem 6.9.1 *(ii)*. Let $d = 25$. Then $E_d/\mathbb{Q}$ has torsion subgroup $\mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ with generators $P_4^+ = (-24, 48)$ and $P_2 = (16, 0)$. For $d = -8 = 1 - (t^2 - 1)^2$ with $t = 2$, $E_d/\mathbb{Q}$ has torsion subgroup $\mathbb{Z}/8\mathbb{Z}$ with generator $P_8^- = (3, 12)$. For $d = -48$ which is not of the form $1 - (t^2 - 1)^2$ for $t \in \mathbb{Z}$, $E_d/\mathbb{Q}$ has torsion subgroup $\mathbb{Z}/4\mathbb{Z}$ generated by $P_4^+ = (49, 98)$.

6.9.2 Explicit invariants for the Edwards family

The study of ranks of elliptic curves requires an ordering of the curves. Natural such orderings are for example by discriminant or by conductor. In this section, we therefore establish formulae for the j-invariant, discriminant, and conductor of the elliptic curve $E_{a,d}/\mathbb{Q}$.

6.9.2.1 The j-invariant and discriminant

By Equation (6.6), the j-invariant of $E_{a,d}$ is

$$j(E_{a,d}) = 16\frac{a^2+14ad+d^2)^3}{ad(a-d)^4} .$$

The discriminant $\Delta(E_{a,d})$ of $E_{a,d}$ is given by

$$\Delta(E_{a,d}) = 2^8 ad(a-d)^4 , \tag{6.24}$$

which is seen by either using the defining cubic of $E_{a,d}$, or the formula relating the discriminant and the j-invariant. This is coherent with the fact that $a, d, a-d$ are non-zero as otherwise the curve defining $E_{a,d}$ is singular.

6.9.2.2 The conductor

We describe a formula for the conductor $N(E_{a,d})$ of $E_{a,d}$ in terms of a, d. We refer to [Sil94, Chapter IV, §10, §11] for a description of general results, on which we rely.
Let $E_{a,d}^{\min}$ denote a minimal model for $E_{a,d}$ and $\Delta_{\min}(E_{a,d})$ the minimal discriminant of $E_{a,d}$. There exists $u \in \mathbb{Z}$ such that a Weierstrass equation for $E_{a,d}^{\min}$ is obtained after a change of variables of the form $(x, y) \mapsto (u^{-2}x, u^{-3}y)$ and $\Delta_{\min}(E_{a,d}) = u^{-12}\Delta(E_{a,d})$. We write $N(E_{a,d})$ as

$$N(E_{a,d}) = \prod_{p|\Delta_{\min}(E_{a,d})} p^{f_p(E_{a,d})} ,$$

where the product runs over the prime numbers at which $E_{a,d}$ has bad reduction, i.e. the prime divisors of $\Delta_{\min}(E_{a,d})$, and $f_p(E_{a,d}) \in \mathbb{N}$ is the associated *conductor exponent*. We write $f_p(E_{a,d}) = \varepsilon_p(E_{a,d}) + \delta_p(E_{a,d})$ where $\varepsilon_p(E_{a,d}) \in \{0, 1, 2\}$ is the tame part and $\delta_p(E_{a,d})$ is the wild part. If $E_{a,d}$ has (bad) multiplicative reduction at p, then $\varepsilon_p(E_{a,d}) = 1$, if $E_{a,d}$ has (bad) additive reduction at p, then $\varepsilon_p(E_{a,d}) = 2$. We can reformulate this in terms of the j-invariant of $E_{a,d}$: let v_p be the p-adic valuation; then:

$$\begin{cases} \varepsilon_p(E_{a,d}) = 1 , \text{ if } v_p(j(E_{a,d})) < 0 \\ \varepsilon_p(E_{a,d}) = 2 , \text{ if } v_p(j(E_{a,d})) \geq 0 \end{cases} .$$

The wild part is often zero by the following criterion: $\delta_p(E_{a,d})$ is zero if $p \geq 5$ or if $E_{a,d}$ has good or split multiplicative reduction at p.

Let p be a prime divisor of $\Delta_{\min}(E_{a,d}) = u^{-12}2^8 ad(a-d)^4$. We compute the exponent $f_p(E_{a,d})$. Assume $p \geq 5$; we see that $v_p(j(E_{a,d})) = v_p(16) + 3v_p(a^2+14ad+d^2) - v_p(a) - v_p(d) - 4v_p(a-d)$ and thus

$$\varepsilon_p(E_{a,d}) = \begin{cases} 1 , \text{ if } 3v_p(a^2+14ad+d^2) < v_p(a) + v_p(d) + 4v_p(a-d) \\ 2 , \text{ if } 3v_p(a^2+14ad+d^2) \geq v_p(a) + v_p(d) + 4v_p(a-d) \end{cases}$$

If $p \geq 5$, $\delta_p(E_{a,d}) = 0$; thus $f_p(E_{a,d})$ is given by $\varepsilon_p(E_{a,d})$. We deduce for $p \geq 5$:

$$f_p(E_{a,d}) = \begin{cases} 1 , \text{ if } 3v_p(a^2+14ad+d^2) < v_p(a) + v_p(d) + 4v_p(a-d) \\ 2 , \text{ if } 3v_p(a^2+14ad+d^2) \geq v_p(a) + v_p(d) + 4v_p(a-d) \end{cases} \tag{6.25}$$

	κ	I_0	I_n	II	III	IV	I_0^*	I_n^*	IV*	III*	II*
	m_κ	1	n	1	2	3	5	$n+5$	7	8	9
$p=2$	$v_2(\Delta_{\min}(E_{a,d}))$	0	n			4			8		
	$f_2^{(\kappa)}(E_{a,d})$	0	1			2			2		
$p=3$	$v_3(\Delta_{\min}(E_{a,d}))$	0	n		3		6	$n+6$		9	
	$f_3^{(\kappa)}(E_{a,d})$	0	1		2		2	2		2	

Table 6.2: Conductor exponents for $E_{a,d}$ in characteristic 2 and 3

It remains to treat the cases $p = 2, 3$. We rely on Ogg's Formula and Kodaira's symbol for reduction types of $E_{a,d}$. Let $\mathcal{K}$ be the set of Kodaira symbols describing the reduction types of $E_{a,d}$. For $\kappa \in \mathcal{K}$, let m_κ be the number of components (counted with multiplicity 1) on the special fiber of $E_{a,d}$. Then Ogg's formula for E and reduction type κ, states:

$$f_p^{(\kappa)}(E_{a,d}) = v_p(\Delta_{\min}(E_{a,d})) - m_\kappa + 1 \; . \tag{6.26}$$

By Ogg's formula (6.26), for $\kappa \in \mathcal{K}$ we can thus write

$$f_p^{(\kappa)}(E_{a,d}) = -12v_p(u) + v_p(\Delta(E_{a,d})) - m_\kappa + 1 \; ,$$

where, by Equation (6.24):

$$v_p(\Delta(E_{a,d})) = \begin{cases} 8 + v_2(a) + v_2(d) + 4v_2(a-d) \; , \text{ if } p = 2 \\ v_3(a) + v_3(d) + 4v_3(a-d) \; , \text{ if } p = 3 \end{cases} \; .$$

We explicitly rely on Table 4.1 in [Sil94, Chapter IV, §9] to compute $f_p^{(\kappa)}(E_{a,d})$. The missing grey cells in Table 6.2 are not directly defined for the given characteristic, and thus computed by Ogg's formula and listed below, with

$$c_p := -12v_p(u) + v_p(a) + v_p(d) + 4v_p(a-d) \quad , \quad p \in \{2, 3\} \; . \tag{6.27}$$

$$f_2^{(\kappa)}(E_{a,d}) = \begin{cases} c_2 + 8 & \text{if } \kappa = \mathrm{II} \\ c_2 + 7 & \text{if } \kappa = \mathrm{III} \\ c_2 + 4 & \text{if } \kappa = \mathrm{I}_0^* \\ c_2 + 4 - n & \text{if } \kappa = \mathrm{I}_n^*, n \geq 1 \\ c_2 + 1 & \text{if } \kappa = \mathrm{III}^* \\ c_2 & \text{if } \kappa = \mathrm{II}^* \end{cases} \qquad f_3^{(\kappa)}(E_{a,d}) = \begin{cases} c_3 & \text{if } \kappa = \mathrm{II} \\ c_3 - 2 & \text{if } \kappa = \mathrm{IV} \\ c_3 - 6 & \text{if } \kappa = \mathrm{IV}^* \\ c_3 - 8 & \text{if } \kappa = \mathrm{II}^* \end{cases} \tag{6.28}$$

Edwards Family with $a = 1$

For the rest of this section we fix $a = 1$, and let $E_d/\mathbb{Q}$ be the curve $E_{1,d}/\mathbb{Q}$. Let $N(E_d)$ be the conductor of E_d. The next proposition says that the conductor exponent at bad primes larger than 5 is always 1, and the exponent at 3 is 1 exactly when $d \not\equiv 2 \pmod 3$, and 0 otherwise. In particular, up to a power of 2, $N(E_d)$ is the radical $\mathrm{rad}(d(d-1)) = \prod_{p \mid d(d-1)} p$ of $d(d-1)$, that is, the product of its prime divisors.

Proposition 6.9.3. *(i) For every $d \in \mathbb{Z} \setminus \{0,1\}$, the conductor exponent of E_d at 3 is given by*

$$f_3(E_d) = \begin{cases} 0 & \text{if } d \equiv 2 \pmod 3 \\ 1 & \text{if } d \equiv 0,1 \pmod 3 \end{cases}.$$

(ii) For every $d \in \mathbb{Z} \setminus \{0,1\}$ and every prime $p \geq 5$ with $p \mid N(E_d)$, one has $f_p(E_d) = 1$. In particular, the formula for the conductor is

$$N(E_d) = 2^{f_2} 3^{f_3} \prod_{\substack{p \mid d(d-1) \\ p \neq 2,3}} p = 2^{f_2 - 1} \cdot \mathrm{rad}(d(d-1)),$$

where f_2 is described above, and $f_3 \in \{0,1\}$ is as in (i).

Proof. (i) If $d \equiv 0,1 \pmod 3$, then the discriminant of E_d is zero modulo 3, i.e. 3 is a bad prime. To see that 3 divides $N(E_d)$ exactly once, it is enough to show that the Kodaira symbol at 3 is I_n with $n = v_3(\Delta_{\min}(E_d))$ (see Table 6.2). We rely on Tate's Algorithm for computing the Kodaira symbol (see [Sil94, Chapter IV, §9, Algorithm 9]). By this algorithm, it is enough to show that 3 does not divide the first b-invariant b_2 of the singular Weierstrass equation (see Step 2) and in that case, the Kodaira symbol is I_n with $n = v_3(\Delta(E_d))$ (the algorithm does not necessarily require a minimal model as input). Assume first that $d \equiv 0 \pmod 3$. Then the equation for E_d is $y^2 = x^3 + 2x^2 + x = x(x+1)^2$ over $\mathbb{F}_3$, which has a singularity at $(-1, 0)$. Moving the singularity to $(0,0)$ via $x' := x + 1$ leads to the equation $y^2 = x'^3 - x'^2$. This Weierstrass model has $b_2 = a_1^2 + 4a_2 = -4$, not divisible by 3. Thus, the Kodaira symbol at 3 is I_n with $n = v_3(\Delta) = v_3(d)$. If $d \equiv 1 \pmod 3$ then E_d is given by $y^2 = x^3 + 4x^2$ over $\mathbb{F}_3$, which has a singularity at $(0,0)$. It has $b_2 = 16$ which is not divisible by 3, thus the Kodaira symbol at 3 is I_n with $n = v_3(\Delta) = 4v_3(d-1)$.

(ii) Let $p \geq 5$ be a prime divisor of $N(E_d)$. Then p divides $d(d-1)$, and thus either d or $d-1$. We show that $f_p(E_d) = 1$ by using Equation (6.25). It is enough to show that $3v_p(1 + 14d + d^2) < v_p(d) + 4v_p(1-d)$. The right-hand side is strictly positive since p divides either d or $1-d$. We next show that $v_p(1 + 14d + d^2) = 0$, thus proving the proposition. Since d and $d-1$ are coprime, it is enough to show that $1 + 14d + d^2$ is not divisible by p in both cases $d \equiv 0 \pmod p$ and $d - 1 \equiv 0 \pmod p$. In the first case, $1 + 14d + d^2 \equiv 1 \not\equiv 0 \pmod p$; in the second case, $1 + 14d + d^2 \equiv 16 \not\equiv 0 \pmod p$, as $p \neq 2$. □

The conductor plot. Let N be the function $\mathbb{Z} \setminus \{0,1\} \to \mathbb{N}, d \mapsto N(E_d)$. We remark that N globally presents a certain structure, although $N(E_d)$ locally presents "jumps" when varying d. As depicted in Figure 6.1, the points $(d, N(E_d))$ describe a pencil of parabolas. The middle parabola can be seen as a *limiting* parabola for the sequence of all parabolas.

The limiting parabola. Let us denote the limiting parabola by $\mathcal{P}_0$. By quadratic interpolation, we easily derive that points $(d, N(E_d))$ on $\mathcal{P}_0$ satisfy $N(E_d) = 16d(d-1)$. Define:

$$\mathcal{D}_0 = \{d \in \mathbb{Z} \setminus \{0,1\} : N(E_d) = 16d(d-1)\}.$$

In other words, $\mathcal{D}_0$ is the set of integers d such that $(d, N(E_d))$ lies on $\mathcal{P}_0$.

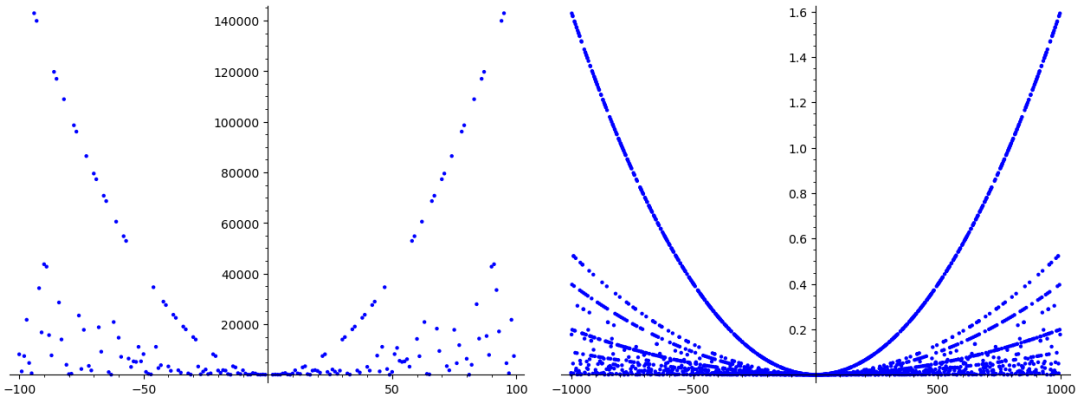

Figure 6.1: $N(E_d)$ as a function of d; on the left-hand side $-100 \leq d \leq 100$, and on the right-hand side $-1000 \leq d \leq 1000$

Below, we shall prove an interesting *density* result for $\mathcal{D}_0$. For a positive integer $X \geq 2$, let $\mathcal{D}_0^{(X)} \subseteq \mathcal{D}_0$ be defined by $\mathcal{D}_0 \cap [-X, X]$. Clearly, $\#\mathcal{D}_0^{(X)} \leq 2X - 1$. The study of $\mathcal{D}_0$ and $\mathcal{P}_0$ is motivated by Figure 6.1 in two ways. First, $\mathcal{P}_0$ describes the largest quadratic growth among all the conductors within a fixed range $[-X, X]$. Second, $\mathcal{P}_0$ seemingly is the densest parabola (i.e. the parabola containing most points within a fixed range $[-X, X]$). We experimentally compute the quantity $\#\mathcal{D}_0^{(X)}/(2X-1)$ for increasingly large values of $X \geq 2$ in SageMath [S$^+$20]. Our computations give:

$$\begin{aligned}
\#\mathcal{D}_0^{(X)}/(2X-1) &= 0.323016150807540... \quad , \quad X = 10^4 \\
\#\mathcal{D}_0^{(X)}/(2X-1) &= 0.322691613458067... \quad , \quad X = 10^5 \\
\#\mathcal{D}_0^{(X)}/(2X-1) &= 0.322619161309581... \quad , \quad X = 10^6
\end{aligned}$$

This suggests that about one third of all curves E_d in our family have conductor $16d(d-1)$. We are interested in computing the natural density of $\mathcal{D}_0$. We say that a subset $A \subseteq \mathbb{N}$ has *natural density* $0 \leq \alpha \leq 1$ if the proportion of elements in A among all natural numbers up to X tends to α, as $X \to +\infty$. This definition extends to $A \subseteq \mathbb{Z}$ by considering the centrally symmertric interval $[-X, X] \cap \mathbb{Z}$, that is,

$$\alpha = \lim_{X \to +\infty} \#(A \cap [-X, X])/(2X-1) \, ,$$

provided that the limit exists.

Under an assumption on the conductor exponent at $p = 2$, the following proposition shows that $\mathcal{D}_0$ is equal to the set of integers d such that $d(d-1)$ is squarefree. In particular, this gives a formula for its natural density in the integers, in terms of the Feller-Tornier constant [FT33], defined by

$$C_{\mathrm{F,T}} = \frac{1}{2} + \frac{1}{2} \prod_{p \text{ prime}} (1 - 2/p^2) = 0.66131704946... \, ,$$

which is the density of integers which have an even number of prime factors p^m with $m > 1$ in their prime factorization, [FT33].

Proposition 6.9.4. *Let $T = \{x \in \mathbb{Z} : x(x-1)$ is squarefree$\}$. Assume that for every $d \in \mathcal{D}_0 \cup T$, one has $f_2(E_d) = 5$. Then, it holds $\mathcal{D}_0 = T$. In particular, the natural density of $\mathcal{D}_0$ in the integers is*

$$\lim_{X \to +\infty} \frac{\#\mathcal{D}_0^{(X)}}{2X-1} = 2 \cdot C_{\mathrm{F,T}} - 1 = 0.322634098920...$$

Proof. To see that $\mathcal{D}_0 \subseteq T$, note that, by Proposition 6.9.3:

$$d \in \mathcal{D}_0 \iff \exists\, (f_2, f_3) \in \mathbb{N} \times \{0,1\} : 2^{f_2} 3^{f_3} \prod_{\substack{p \mid d(d-1) \\ p \neq 2,3}} p = 16d(d-1) = N(E_d)\,,$$

where f_2, f_3 are described in Proposition 6.9.3. By assumption, $f_2(E_d) = 5$. Dividing by 16 gives $d(d-1) = 2 \cdot 3^{f_3} \prod_{p \mid d(d-1), p \neq 2,3} p$, thus $d \in T$.

Conversely, let $d \in T$. As $v_2(d(d-1)) = 1$, the prime factorization of $d(d-1)$ has the form

$$d(d-1) = 2 \cdot 3^{a_3} \prod_{\substack{p \mid d(d-1) \\ p \neq 2,3}} p\,,$$

with $a_3 = 1$, if $d \equiv 0, 1 \pmod 3$, and $a_3 = 0$, if $d \equiv 2 \pmod 3$. Thus, by Proposition 6.9.3, $a_3 = f_3(E_d) =: f_3$. Multiplying by 16 gives $16d(d-1) = 2^5 3^{f_3} \prod_{p \mid d(d-1), p \neq 2,3} p$. By the assumption $f_2(E_d) = 5$, unique factorization, and Proposition 6.9.3, it follows $N(E_d) = 16d(d-1)$.

To compute the density of $\mathcal{D}_0$, we compute the density of T. Now, $d \in T$ if and only if $d \not\equiv 0 \pmod{p^2}$ and $d \not\equiv 1 \pmod{p^2}$ for every prime p. Thereby, there are left only $p^2 - 2$ classes in $\mathbb{Z}/p^2\mathbb{Z}$ for the choice of d modulo p^2. Consequently, the natural density of T is

$$\prod_{p \text{ prime}} \frac{p^2-2}{p^2} = \prod_{p \text{ prime}} \left(1 - \frac{2}{p^2}\right)\,.$$

The result follows from the definition of $C_{\mathrm{F,T}}$. □

Remark 6.9.5. We make some remarks on the assumption in Proposition 6.9.4.

(i) Note that for every $d \in \mathcal{D}_0$, one has $f_2(E_d) = v_2(16d(d-1)) = 4 + v_2(d(d-1)) \geq 5$. Therefore, the assumption $f_2(E_d) = 5$ says that this bound is *tight*. If moreover $d \in T$, then $f_2(E_d) = 5$. In particular, our assumption is trivially satisfied on $\mathcal{D}_0 \cap T$.

(ii) Conjecturally, Proposition 6.9.4 holds true *without* the assumption $f_2(E_d) = 5$ for $d \in \mathcal{D}_0 \cup T$. By a direct SageMath verification, it is verified that $f_2(E_d) = 5$ for $d \in \mathcal{D}_0^{(X)}$ for $X = 10^6$. Our work in progress includes the removal of the assumption and thus an unconditional result: roughly 32% of the curves in $\{E_d\}_d$ have conductor $16d(d-1)$.

One way of removing the assumption is to establish an explicit formula for $f_2(E_d)$ in terms of d, similarly as done for $f_3(E_d)$ in Proposition 6.9.4 *(i)*. This could for example be done by relying on Tate's Algorithm for finding a minimal model ([Sil94, Chapter IV, §9, Algorithm 9]), and combining with Equation (6.28). By computational observation, we conjecture that for every $d \in \mathbb{Z} \setminus \{0,1\}$, $f_2(E_d) = 5$ if and only if $v_2(d(d-1)) = 1$. In that it is explicit in d, this is a more handy condition. It holds true for $-10^6 \leq d \leq 10^6$. Using SageMath [S$^+$20], we see that for every $d \in \mathcal{D}_0^{(X)}$ with $X = 10^6$, the Kodaira symbol at $p = 2$ is I_0^*, if d is even, and I_3^*, if d is odd. Moreover, the discriminant for E_d is a minimal discriminant for E_d, implying $u = 1$ in Equation (6.27), and thus $c_2 = v_2(d) + 4v_2(1-d)$ (remember that $a = 1$).

6.9.3 Descent via isogenies

Since the curve E_d has a point of order 2, the descent method via 2-isogenies presented in Theorem 2.3.4 can be applied explicitly to E_d. For more background, we refer to Section 2.3.4. Via this method one can obtain the following upper bound on the rank of an elliptic curve $E/\mathbb{Q}$ given by $y^2 = x^3+ax^2+bx$ with $a, b \in \mathbb{Z}$ such that $(a^2-4b)b \neq 0$ (see e.g. [ALRP08, Proposition 1.1]):

$$r(E) \leq \omega(a^2 - 4b) + \omega(b) - 1 \,, \tag{6.29}$$

where for a non-zero integer N, $\omega(N)$ denotes the number of its distinct prime divisors. Elliptic curves achieving this bound are said to have *maximal* Mordell-Weil rank in [ALRP08]. Applied to the curve E_d, this gives the following lemma.

Lemma 6.9.6. *Let $d \in \mathbb{Z} \setminus \{0, 1\}$ and E_d the elliptic curve $y^2 = x^3 + 2(1 + d)x^2 + (1 - d)^2x$ over $\mathbb{Q}$. Then $r(E_d) \leq \omega(2d) + \omega(1 - d) - 1$.*

Proof. Equation (6.29) with $a = 2(1 + d)$ and $b = (1 - d)^2$ gives $a^2 - 4b = 16d$, and $r(E_d) \leq \omega(16d) + \omega(1 - d) - 1$, and the bound follows. □

This upper bound is tight and cannot be improved for general d: for example $d = 194 = 2 \cdot 97$ yields the upper bound $\omega(2 \cdot 2 \cdot 97) + \omega(-193) - 1 = 2$ on the rank of E_{194}. A direct SageMath calculation shows that its rank is 2: the points

$$(-169, 130) \,, \left(\frac{1550154384}{27889}, \frac{61246789681740}{4657463}\right) \in E_{194}(\mathbb{Q})$$

are not torsion points (see Theorem 6.9.1), and are generators for the free part of $E_{194}(\mathbb{Q})$.

Remark 6.9.7. Inspired by the upper bound in Lemma 6.9.6, one can ask whether the family $\{E_d\}_d$ contains curves of large rank. In [EK20, Section 11] the authors describe fibrations of elliptic curves with torsion $\mathbb{Z}/4\mathbb{Z}$. Therefore, these curves admit a rational Edwards model. For the three rank record breaking rank-13 curves resulting from these fibrations (see [EK20, Appendix B.4]) the Edwards curve has parameter $d \in \mathbb{Q} \setminus \mathbb{Z}$ (namely $d \approx 1.374, 1.060, 3.332$), as can be seen by applying Proposition 6.6.1 with $P = (0, 0)$, and thus does not lie in our family. Explicitly, for the Edwards model corresponding to the fibration $\mathcal{E}_1$ (this is the notation used in [EK20]; not to confuse with our notation for Edwards curves) (see [EK20, Equation (5)]) we find, with $P = (0, 0)$ and $2P = (-8(31t - 7)(15t - 8)(t + 1), 0)$, that $d_P = 1 - \frac{128(31t-7)(15t-8)(t+1)}{(32t+7)^2(8t-1)^2}$. The only integral solution for d_P is $t \in \{-1, 7/31, 8/15\}$ and $d_P = 1$, which leads to a singular curve. Therefore, curves of large ranks in the family $\{E_d\}_d$ are not obtained by the fibrations from [EK20].

We now apply the descent via 2-isogenies to E_d. More precisely, we apply Theorem 2.3.4 to the curve $E_d : y^2 = x^3 + 2(1 + d)x^2 + (1 - d)^2x$ over $\mathbb{Q}$. Let ϕ be the 2-isogeny between E_d and $E'_d : y^2 = x^3 - 4(1 + d)x^2 + 16dx$, with kernel $\{\mathcal{O}, (0, 0)\}$. Following the notation of Theorem 2.3.4, let $\mathcal{S} = \mathcal{S}(d)$ be the set of places of $\mathbb{Q}$

$$\mathcal{S} = \{\infty, 2\} \cup \{p \in \mathbb{N} \text{ prime} : p \mid d(d - 1)\} \,.$$

If $\{p \in \mathbb{N} \text{ prime} : p \mid d(d - 1)\} \cup \{2\} = \{p_1, \ldots, p_l\}$ (where $l = \omega(d(d - 1))$) then $\mathbb{Q}(\mathcal{S})$ contains the representatives $\{\pm 1, \pm 2\}$ and those coming from the divisors of $d(d-1)$, i.e. $\mathbb{Q}(\mathcal{S})$

is identified with the set $\{\pm p_1^{s_1}\cdots p_l^{s_l} : s_i = 0,1, i = 1,\dots,l\}$ modulo $(\mathbb{Q}^\times)^2$. In particular, $\#\mathbb{Q}(\mathcal{S}) = 2^{l+1} = 2^{\omega(d(d-1))+1}$. Since $\mathrm{Sel}^{(\phi)}(E_d)$ is a subgroup of $\mathbb{Q}(\mathcal{S})$, we deduce that

$$\#\mathrm{Sel}^{(\phi)}(E_d) \leq 2^{\omega(d(d-1))+1} ,$$

and similarly for $\#\mathrm{Sel}^{(\phi^\vee)}(E'_d)$. The descriptions of $\mathrm{Sel}^{(\phi)}(E_d)$ and $\#\mathrm{Sel}^{(\phi^\vee)}(E'_d)$ are obtained by checking whether the principal homogenous spaces (for E_d, resp. E'_d)

$$\begin{cases} C_\lambda : \lambda W^2 = \lambda^2 - 4(1+d)\lambda Z^2 + 16dZ^4 \\ C'_\lambda : \lambda W^2 = \lambda^2 + 8(1+d)\lambda Z^2 + 16(d-1)^2 Z^4 , \end{cases} \tag{6.30}$$

obtained by interchanging the roles of E_d and E'_d (and considering $\phi^\vee : E'_d \to E_d$), locally have rational points for every $\lambda \in \mathbb{Q}(\mathcal{S})$.

Example 6.9.8. We illustrate Theorem 2.3.4 for E_d by a few computations. To control the prime divisors of $d(d-1)$, we let p and q be odd prime numbers such that such that

$$d = p = 2q + 1 ,$$

and consider the curve E_p over $\mathbb{Q}$. In this case, q is called a *Sophie-Germain prime*, and p a *safe prime*. It is not known whether there are infinitely many pairs (p,q); this is a conjecture. Therefore, we do not know whether $\{E_p\}_p$ is an infinite family. By Lemma 6.9.6, these curves have rank at most 3. Let $\mathcal{S} = \{\infty, 2, p, q\}$ and $\mathbb{Q}(\mathcal{S}) = \{\pm 1, \pm 2, \pm p, \pm q, \pm 2p, \pm 2q, \pm pq, \pm 2pq\}$, which is of cardinality 16 (note: $16 = 2^{\omega(p(p-1))+1}$).

(i) One has

$$\mathrm{Sel}^{(\phi)}(E_p) \simeq \mathbb{Z}/2\mathbb{Z} .$$

To compute $\mathrm{Sel}^{(\phi)}(E_p)$ we look for the existence of local points on the quartic C_λ from Equation (6.30) for every $\lambda \in \mathbb{Q}(\mathcal{S})$. Using Theorem 2.3.4, we see that the connecting homomorphism δ maps $(0,0)$ to $a^2 - 4b = 16p = p \pmod{(\mathbb{Q}^\times)^2}$. Therefore $p \in \mathrm{Sel}^{(\phi)}(E_p)$.

Further, it is easy to see by a sign argument that if $\lambda < 0$, then $\lambda \notin \mathrm{Sel}^{(\phi)}(E_p)$ because $C_\lambda(\mathbb{R}) = \emptyset$. Indeed, if $(Z,W) \in C_\lambda(\mathbb{R})$ then $\lambda W^2 = \lambda^2 - 4(1+p)\lambda Z^2 + 16pZ^4 = (\lambda - 4Z^2)(\lambda - 4pZ^2)$, where the left-hand-side is < 0 and the right-hand-side is > 0.

Next, we establish that $\lambda = 2 \notin \mathrm{Sel}^{(\phi)}(E_p)$ by showing that $C_2(\mathbb{Q}_2) = \emptyset$. Dividing the equation for $C_2 : 2W^2 = 4 - 8(1+p)Z^2 + 16pZ^4$ by 2 gives $W^2 = 2 - 4(1+p)Z^2 + 8pZ^4$. Let $(Z,W) \in C_2(\mathbb{Q}_2)$. Because $v_2(W^2)$ and $v_2(8pZ^4)$ have different parity, we can assume that (Z,W) is 2-adically integral, i.e. $(Z,W) \in C_2(\mathbb{Z}_2)$. Reducing modulo 2 gives $W \equiv 0 \pmod 2$. The case $Z \equiv 0 \pmod 2$ gives $W^2 \equiv 2 \pmod{32}$, a contradiction, therefore $Z \equiv 1 \pmod 2$. But this gives $W^2 \equiv 2 \pmod 8$, a contradiction. Hence $2 \notin \mathrm{Sel}^{(\phi)}(E_p)$, which implies (since $\mathrm{Sel}^{(\phi)}(E_p) \subseteq \mathbb{Q}(\mathcal{S})$ is a subgroup) that $\#\mathrm{Sel}^{(\phi)}(E_p) \in \{2,4\}$.

We show that $C_q(\mathbb{Q}_q) = \emptyset$, which gives $q \notin \mathrm{Sel}^{(\phi)}(E_p)$. Similarly to before, the equation of C_q over $\mathbb{Q}_q$ implies that $v_q(qW^2)$ and $v_q(16pZ^4)$ have different parity, thus we look for $\mathbb{Z}_q$-integral points. The equation gives $Z \equiv 0 \pmod q$, which in turn implies $qW^2 \equiv q^2 \pmod q$, thus $W \equiv 0 \pmod q$. But this gives the contradiction $q^2 \equiv 0 \pmod{q^3}$.

Our arguments show that $\mathrm{Sel}^{\phi}(E_p) = \{1,p\} \simeq \mathbb{Z}/2\mathbb{Z}$.

(ii) The computations for $\mathrm{Sel}^{\phi^\vee}(E'_p)$ are similar, using C'_λ in Equation (2.13). We will not do the details. Using the SageMath [S+20] command

```
sage: E.simon_two_descent(verbose=4)
```

we have verified, for $3 \leq q \leq 2000$ (giving 294 choices for p and thus 294 curves E_p), that $\#\,\mathrm{Sel}^{(\phi^\vee)}(E'_p/\mathbb{Q}) = 4$, $\#\,\mathrm{Sel}^{(\phi)}(E_p/\mathbb{Q}) = 2$ (as proved in *(i)*), $\#\,Ш(E_p/\mathbb{Q})[2] = 1$, and that $\mathrm{rk}_{\mathbb{Z}}(E_p(\mathbb{Q})) = 1$.

6.9.4 Our computations of the ranks in the Edwards family

Motivation. Unlike for the torsion subgroup of elliptic curves, computing the rank is often much harder, and often a more complicated and rather mysterious task. The literature on the ranks of elliptic curves is vast and combines a variety of methods from number theory, such as their relation to analytic number theory via the Birch and Swinnerton-Dyer Conjecture [BSD65, Wil06]. Analyzing statistical distributions of the rank of elliptic curves (in a given family) is a common approach to make observations and state conjectures.

For example, Zagier and Kramarz [ZK87] studied the ranks in the family $x^3 + y^3 = m$ for cubefree integers m. These curves are birationally equivalent to the elliptic curves $E_{(m)}$: $y^2 = x^3 - 432m^2$. Extensive computations of the rank of $E_{(m)}$ for m up to 7000 support their conjecture, claiming that curves with rank at least 2 occur with positive density. Watkins [Wat07] subsequently extends the data from [ZK87] to $m \leq 10^7$ and shows that the density is more likely to tend to zero. This is somewhat surprising in view of Goldfeld's Conjecture [Gol79]: on average, the rank of elliptic curves is $1/2$, i.e. half of the curves have rank 0, and half of the curves have rank 1. The remaining curves should be 0%, thus not leaving room for *larger* ranks. Although originally formulated for elliptic curves in families of *twists*, this conjecture is believed to hold more generally.

In this section, we provide statistics for the rank of elliptic curves in the family $\{E_{a,d}\}_{a,d}$. In this thesis we include current results for 20000 curves[3], obtained by varying $-10^4 \leq d \leq 10^4$ and fixing $a = 1$. More computations for larger ranges of d and choices of $a \neq 1$ are in progress and will be included in the preprint.

Ordering of curves and average rank. For producing statistics on elliptic curves it is important to choose a suitable way to *order* them. Possible orderings are typically by classical invariants, such as, the discriminant or the conductor. These orderings are sometimes unsuited, as asymptotically it is not known how many elliptic curves there are up to a certain bound on the discriminant or conductor. Instead, one often works with the naive height of elliptic curves, as it is known that there are on the order of $h^{5/6}$ curves with height at most h.

In this work however, we choose a different ordering of our curves, namely the *natural* order induced by $\mathbb{Z}$. For $X \in \mathbb{N}$, we define $\mathbb{E} = \{E_d : d \in \mathbb{Z} \setminus \{0,1\}\}$ and

$$\mathbb{E}_X = \{E_d : d \in [-X, X] \cap \mathbb{Z} \setminus \{0,1\}\}\ .$$

Remark 6.9.9. The set $\mathbb{E}$ does not contain isomorphic curves, except $E_{16} \simeq E_{81}$. Namely, for $d \in \mathbb{Z} \setminus \{0,1\}$, we obtain isomorphic curves for $d' \in \Sigma(d) \cap (\mathbb{Z} \setminus \{0,1\})$. The only integer solution d' is obtained by taking $d' = ((1 \pm d^{1/4})/(1 \mp d^{1/4}))^4$. To have $d' \in \mathbb{Z} \setminus \{0,1\}$, d must be a fourth power, say $d = s^4$ with $s \in \mathbb{Z} \setminus \{0,1\}$, and either $(1-s) \mid (1+s)$ or $(1+s) \mid (1-s)$. If

[3] 19998 to be precise, in view of choosing $d \neq 0, 1$

$s > 0$ then $1+s > 1-s$, so the second case cannot occur, and we must have $(1+s) = \ell(1-s)$, for some $\ell \in \mathbb{Z} \setminus \{0,1\}$. Equivalently, $s = (\ell-1)/(\ell+1) = 1 - 2/(\ell+1)$. Thus, $\ell+1$ must divide 2, giving the possibilities $\ell \in \{-3,-2\}$, and the corresponding solutions for s are $\{2,3\}$. Similarly, when $s < 0$, the only possibilities for $d' \in \mathbb{Z} \setminus \{0,1\}$ are $s \in \{-3,-2\}$. This gives the only solutions $d = s^4 \in \{16, 81\}$.

Our results. We describe statistics for the algebraic rank of the curves in $\mathbb{E}_X$ for increasingly large values of X, obtained in SageMath [S+20] and Magma [BCP97]. In Table 6.3, we describe the statistics for $\mathbb{E}_X$ for various bounds X, including the ranks of the curves, the number of curves (num. of curves) of given rank, their proportion, and the average rank (av. rank).

X	rank	num. of curves	proportion (%)	av. rank
1000	0	886	44.32	0.62
	1	988	49.42	
	2	122	6.10	
	3	3	0.15	
5000	0	4219	42.19	0.66
	1	4988	49.88	
	2	764	7.64	
	3	28	0.28	
10000	0	8363	41.81	0.67
	1	9933	49.66	
	2	1633	8.16	
	3	69	0.34	

Table 6.3: Distribution of the ranks in the family $\{E_d\}_d$

The database LMFDB provides statistics for rational elliptic curves; see the link provided in [LMF20b]. A major difference with our statistics is that the curves are ordered by conductor in [LMF20b]. The database contains all curves of conductor up to 299996953, which are in total 3824372 curves. We give a short comparison with our Table 6.3, even though their number of curves largely exceeds ours. In [LMF20b], 12.90% of the curves have rank 2 and 0.98% have rank 3. In Table 6.3, the corresponding proportions (i.e. 8.16% and 0.34% for $X = 10000$) are already rather large, given the much smaller number of curves. This indicates that rank-2 curves potentially accumulate much more rapidly in the family $\{E_d\}_d$.

Note also that the conductors of the curves $\{E_d\}_d$ grow drastically (see Section 6.9.2.2, i.e. $N(E_d) \sim 16d^2$ for about 32% of the curves, see Proposition 6.9.4). The largest conductor in $\mathbb{E}_{10^4}$ is

$$N(E_d) = 1599200096 = 2^5 \cdot 13 \cdot 769 \cdot 4999 \,,$$

occuring for $d = -9997$ and $d = 9998$. This value exceeds the largest conductor (i.e. 299996953) in LMFDB by a factor roughly 5. Therefore, our family of curves is not included in the family of curves for the LMFDB-database.

6.9.5 Computation of the analytic order of the Shafarevich-Tate group

We now provide a simple algorithm for computing the analytic order of the Shafarevich-Tate group. We refer to Section 2.3.4 for the definition and background for this group. Our technique relies on *truncations* of L-functions and assumes the validity of the Birch and Swinnerton-Dyer Conjecture. To our knowledge, our algorithm extends, from a theoretical point of view, special cases in the literature, as we will discuss below.

Our algorithm works for general elliptic curves, not necessarily in Edwards form. For our exposition, we, therefore, let $E/\mathbb{Q}$ be an arbitrary elliptic curve and $Ш(E)$ its Shafarevich-Tate group (see 2.3.2). Later we will provide results for the family $\{E_d\}_d$.

Recall from Section 2.3.3 that the BSD Conjecture (Conjecture 2.3.1) states that $Ш(E)$ is a finite group and

$$\frac{L^{(r)}(E,1)}{r!} = \frac{\#Ш(E)\cdot\Omega_E\cdot R_E\cdot c_E}{(\#E(\mathbb{Q})_{\mathrm{tors}})^2}\,, \tag{6.31}$$

where r is the analytic rank of E, that is, the order of vanishing of the L-series of E, at $s=1$. Henceforth, we will assume the BSD Conjecture as true, and rely on Equation (6.31) to "define" $\#Ш(E)$. Therefore, we call this value the *analytic* order of $Ш(E)$.

We rely on Equation (6.31) to compute an approximation of $\#Ш(E)$ from an approximation of $L^{(r)}(E,1)$. When $r=0$, we obviously interpret $L^{(r)}(E,1)$ as $L(E,1)$. It is given by the series (see [Cre97, Proposition 2.11.1]):

$$L(E,1) = 2\sum_{n\geq 1}\frac{a_n}{n}\exp\left(-\frac{2\pi n}{\sqrt{N(E)}}\right)\,. \tag{6.32}$$

Here the sequence $(a_n)_n$ denotes the sequence of Hecke eigenvalues for E, and are implemented using the recurrence relation from Equation (2.8). When $r\geq 1$, the formula generalizes as (see [Cre97, Proposition 2.13.1]):

$$L^{(r)}(E,1) = 2r!\sum_{n\geq 1}\frac{a_n}{n}G_r\left(\frac{2\pi n}{\sqrt{N(E)}}\right)\,, \tag{6.33}$$

where

$$G_r(x) = \frac{1}{(r-1)!}\int_1^\infty \exp(-tx)(\log(t))^{r-1}\frac{dt}{t}\,,\quad r\geq 1$$

is the generalized exponential integral.

The convergence speed of $L^{(r)}(E,1)$ typically depends on the size of the conductor $N(E)$. Our technique builds on a *truncation* of the series in Equations (6.32) and (6.33). For an integer $n_0\geq 1$, denote by

$$[L^{(r)}(E,1)]_{n_0} := 2r!\sum_{n=1}^{n_0}\frac{a_n}{n}G_r\left(\frac{2\pi n}{\sqrt{N(E)}}\right)\,,\quad r\geq 1\,, \tag{6.34}$$

and similarly when $r=0$, the sum of the first n_0 terms of $L^{(r)}(E,1)$. Once n_0 is fixed, this finite sum can easily be computed. A too small choice of n_0 makes the error too gross to

reconstruct $\#\,Ш(E)$ from the truncation. A large enough value of n_0 guarantees that the approximation error between $[L^{(r)}(E,1)]_{n_0}$ and $L^{(r)}(E,1)$ is negligible. We define the approximation $[\#\,Ш(E)]_{n_0}$ of $\#\,Ш(E)$, by combining Equation (2.9) (BSD) with Equation (6.34). We define:

$$[\#\,Ш(E)]_{n_0} = \frac{[L^{(r)}(E,1)]_{n_0}}{r!} \cdot \frac{(\#E(\mathbb{Q})_{\mathrm{tors}})^2}{\Omega_E \cdot R_E \cdot c_E}\,. \tag{6.35}$$

If our approximation $[L^{(r)}(E,1)]_{n_0}$ is good enough, that is, $|L^{(r)}(E,1) - [L^{(r)}(E,1)]_{n_0}| \leq \varepsilon$ for a small enough ε, then, assuming Equation (2.9) (BSD), also the inequality

$$|\#\,Ш(E) - [\#\,Ш(E)]_{n_0}| \leq \varepsilon$$

holds, from which we can read the integer $\#\,Ш(E)$.

We have implemented our algorithm in [S+20]. The choice of n_0 depends on the conductor $N(E)$. Computationally, we have observed that $n_0 = O(\sqrt{N(E)})$ is sufficient for guaranteeing a good approximation and revealing $\#\,Ш(E)$.

Algorithm 14 Algorithm to compute $\#\,Ш(E)$ assuming the BSD Conjecture

Parameters: An integer $s \geq 1$ (for a desired approximation precision)
Input: An elliptic curve E defined over $\mathbb{Q}$ of conductor $N(E)$
Output: $|\#\,Ш(E)|$

1: $n_0 := \lceil \sqrt{N(E)} \rceil$
2: $r := \mathrm{rk}_{\mathbb{Z}}(E(\mathbb{Q}))$
3: Compute the list $\mathcal{A} = \{a_n(E) : 1 \leq n \leq n_0\}$ using the recurrence relation (2.8)
4: **if** $r = 0$ **then**
5: Compute the list $\mathcal{C} = \{\frac{a_n}{n} \exp(-2\pi n/\sqrt{N(E)}) : 1 \leq n \leq n_0\}$ from the list $\mathcal{A}$
6: **end if**
7: **if** $r \geq 1$ **then**
8: Compute the list $\mathcal{C} = \{\frac{a_n}{n} G_r(2\pi n/\sqrt{N(E)}) : 1 \leq n \leq n_0\}$ from the list $\mathcal{A}$
9: **end if**
10: Compute $L := 2r! \sum_{x \in \mathcal{C}} x$
11: Compute BSD invariants for E: the regulator R_E of E; the real period Ω_E of E; the Tamagawa number c_E of E; the torsion order T_E of E
12: Compute $\#\,Ш := \frac{L}{r!} \cdot \frac{T_E^2}{\Omega_E \cdot R_E \cdot c_E}$
13: **if** $|\#\,Ш - \lfloor \#\,Ш \rceil| < 10^{-s}$ **then**
14: Return $\#\,Ш$
15: **else**
16: Increase n_0 and restart
17: **end if**

Remark 6.9.10. Let us make some remarks about Algorithm 14.

(i) The algorithm is known in the case $r = 0$, see [HY15]. Their algorithm also uses $O(\sqrt{N(E)})$ terms in the truncated finite sum in order to recover $\#\,Ш(E)$. Our Algorithm 14 extends to all ranks with the same bound on the number of terms in the truncated finite sum.

(ii) Note that in order to perform the approximation of the L-series, the rank needs to be known. This is computed in first place by our algorithm. Since our algorithm is conditioned on the BSD Conjecture, r can either be the Mordell-Weil rank $\mathrm{rk}_{\mathbb{Z}}(E(\mathbb{Q}))$ of E, or the analytic rank of E.

(iii) The BSD invariants for E are computable by built-in functions from SageMath [S$^+$20]. We obtain R_E, Ω_E, c_E and T_E by the commands (see Section 2.3.3 for a brief definition of these objects):

```
sage: R_E=E.regulator()
sage: Omega_E=E.period_lattice().omega()
sage: c_E=E.tamagawa_product_bsd()
sage: T_E=E.torsion_order()
```

(iv) A more direct method to compute the special L-value of an elliptic curve E of rank r is to use Dokchitser's SageMath-function:

```
sage: E.lseries().dokchitser().derivative(1,r)/r.factorial()
```

From this value, the analytic order of $Ш(E)$ is computed using Equation (6.31). We note that our outputs of Algorithm 14 are coherent with the outputs of this function (see Table 6.5).

Remark 6.9.11. Our choice $n_0 = \lceil \sqrt{N(E)} \rceil$ is based on computational evidence and is not proven. Note that the algorithm in [HY15] for $r = 0$ is also heuristic on the choice of n_0. This part is currently in progress and we point out a proof direction here. To obtain a proven upper bound, we want to rigorously upper bound the error tail $|L^{(r)}(E,1) - [L^{(r)}(E,1)]_{n_0}|$, resp. the quantity

$$t := \left| \sum_{n \geq n_0} \frac{a_n}{n} G_r\left(\frac{2\pi n}{\sqrt{N(E)}}\right) \right| .$$

By the Deligne bound one has $|a_n| \leq d(n)\sqrt{n}$ (see e.g. [Ken04, Theorem 2.32] applied with $k = 2$) where $d(n)$ is the number of divisors of n. Using the trivial bound $d(n) \leq 2\sqrt{n}$ (because the divisors of n come in pairs $(m, n/m)$ and there are at most $\sqrt{n}$ choices for m), this gives $|a_n| \leq 2n$. Moreover, $G_r(x) \sim x^{-r}\exp(-x)$ as $x \to +\infty$ (see [Mac02]), which gives for all $\eta > 0$ and large enough x, $G_r(x) \leq x^{-r}\exp(-x)\eta$. Applied with $x = 2\pi n/\sqrt{N(E)}$, we obtain

$$t \leq \sum_{n \geq n_0} \frac{|a_n|}{n} G_r\left(\frac{2\pi n}{\sqrt{N(E)}}\right) \leq 2\eta \sum_{n \geq n_0} \left(\frac{2\pi n}{\sqrt{N(E)}}\right)^{-r} \exp\left(-\frac{2\pi n}{\sqrt{N(E)}}\right) .$$

Bounding $n \geq n_0$ in the first factor of the summand, and putting $\alpha = \exp(-2\pi/\sqrt{N(E)})$ (note: $|\alpha| < 1$) gives:

$$\begin{aligned} t &\leq 2\eta \left(\frac{\sqrt{N(E)}}{2\pi n_0}\right)^r \sum_{n \geq n_0} \alpha^n \leq 2\eta \left(\frac{\sqrt{N(E)}}{2\pi n_0}\right)^r \sum_{n \geq 1} \alpha^n \\ &= 2\eta \left(\frac{\sqrt{N(E)}}{2\pi n_0}\right)^r \frac{\alpha}{1-\alpha} = C \cdot n_0^{-r} , \end{aligned}$$

with $C := \frac{2\alpha\eta}{1-\alpha}\left(\frac{\sqrt{N(E)}}{2\pi}\right)^r$. It then remains to resolve the inequality $Cn_0^{-r} \leq \varepsilon$ for n_0, for a targeted error bound $\varepsilon > 0$. This gives $n_0 \geq (C/\varepsilon)^{1/r}$, which is of the order $O(\sqrt{N(E)})$.

Our results for the Edwards family. We provide the results of our SageMath implementation [S+20] of Algorithm 14 for the Edwards family $\{E_d\}_d$. By Theorem 6.9.1, we have

$$\#E_d(\mathbb{Q})_{\text{tors}} = \begin{cases} 8 & \text{if } d \text{ is a square} \\ 8 & \text{if } d \text{ is not a square and } d = 1-(t^2-1)^2, \text{ for some } t \in \mathbb{Z} \\ 4 & \text{if } d \text{ is not a square and } d \neq 1-(t^2-1)^2, \text{ for all } t \in \mathbb{Z} \end{cases},$$

thereby, the torsion subgroup does not need to be precomputed first in order to write the right-hand side of Equation (6.35).

Recall from Section 2.3.4, that the order of $Ш(E)$ is a square, if it is finite. Table 6.4 shows the distribution of $\#Ш(E_d)$ for $-10^4 \leq d \leq 10^4$. Among these curves, the largest Shafarevich-Tate group has order $1369 = 37^2$. Most of the curves in our family (about 70%) have a trivial Shafarevich-Tate group. That this group clearly dominates in our data, is in line with the statistics from LMFBD [LMF20b], where 91.59% of curves have trivial Shafarevich-Tate group. For $\#Ш(E_d) = 4$, our set of curves contains 13.15%, while LMFDB only contains 5.55%. While we are considering much fewer curves than the LMFDB-statistics (compare also with Section 6.9.4), it is interesting to point out the resemblence with [LMF20a].

Table 6.5 compares our Algorithm 14 with the analytic Shafarevich-Tate order computed using the SageMath-function by Dockchitser (compare with Remark 6.9.10). For the seven curves with $\#Ш(E_d) = 22^2$ from Table 6.4, the approximation of $\#Ш(E_d)$, with $r = 1$ (the analytic rank ("an. rank") is 1), computed by our algorithm (column "Algorithm 14") is very close to the approximation of $Ш(E_d)$ computed from Dockchitser's function (column "Dockchitser"). Since the order of $Ш(E_d)$ is a square, the technique relies on finding the closest square integer to these real numbers output by the algorithm. Both algorithms imply that the analytic order of $Ш(E_d)$ is 484 (note: $21^2 = 441, 22^2 = 484, 23^2 = 529$). Notice again the large conductors, exceeding the conductor sizes of [LMF20a] (compare with Section 6.9.2.2). Our Algorithm 14 precomputes a list $\mathcal{A}$ with $O(\sqrt{N(E_d)})$ coefficients.

6.10 Appendix Section

6.10.1 Special families of elliptic curves

Let K be a field of characteristic different from 2. We describe a family of elliptic curves over K with certain properties with respect to their 4-torsion points, and their corresponding set of birationally equivalent Edwards curves. Let $r, s, t \in K$ with s, t non-zero and $s \neq t$. Let $E_{r,s,t}$ be the elliptic curve defined by

$$y^2 = (x-r)(x-r+s^2)(x-r+t^2)\,.$$

$\#Ш(E_d)$	num. of curves	prop. (%)	$Ш(E_d)$	num. of curves	prop. (%)
1^2	14193	70.96	20^2	16	0.08
2^2	2630	13.15	21^2	7	0.03
3^2	1003	5.01	22^2	7	0.03
4^2	716	3.58	23^2	5	0.02
5^2	217	1.08	24^2	4	0.02
6^2	157	0.78	25^2	1	0.005
7^2	48	0.24	26^2	3	0.015
8^2	105	0.52	27^2	0	0.00
9^2	40	0.20	28^2	1	0.005
10^2	22	0.11	29^2	0	0.00
11^2	192	0.96	30^2	0	0.00
12^2	211	1.05	31^2	1	0.005
13^2	142	0.71	32^2	0	0.00
14^2	94	0.47	33^2	0	0.00
15^2	69	0.34	34^2	1	0.005
16^2	44	0.22	35^2	0	0.00
17^2	26	0.13	36^2	0	0.00
18^2	18	0.09	37^2	1	0.005
19^2	25	0.12	38^2	0	0.00

Table 6.4: Distribution of the order of the Shafarevich-Tate group in the family $\{E_d\}_d$ in $\mathbb{E}_X$ with $X = 10^4$

Proposition 6.10.1. *(i) The points of order 2 in $E_{r,s,t}(K)$ are $(r,0)$, $(r-s^2,0)$ and $(r-t^2,0)$. The four points $\pm(r \pm st, st(s \pm t)) \in E_{r,s,t}(K)$ where the second and third $\pm$ signs agree, have order exactly 4 and double to $(r,0)$. In particular, $E_{r,s,t}(K)$ contains a subgroup isomorphic to $\mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$.*

(ii) Define

$$d(s,t) = \left(\frac{s \pm t}{s \mp t}\right)^2 \in K \setminus \{0,1\} .$$

For every $r \in K$, $E_{r,s,t}$ is birationally equivalent over K to the Edwards curve $\mathcal{E}_{d(s,t)}$.

(iii) If r is a square in K and $r + s^2 = t^2$ (i.e. $(\sqrt{r}, s, t)$ is a Pythagorean triple), then all the Edwards curves in $\mathrm{Edw}_{\overline{K}}(j(E_{r,s,t}))$ are defined over $K(\sqrt{-1})$.

Proof. (i) The statement about the 2-torsion points is clear. By Lemma 6.6.7, points of order 4 lying over $(r,0)$ are those mentioned in the statement, and are clearly defined over K.

(ii) Setting $e_1 = r$, $e_2 = r - s^2$ and $e_3 = r - t^2$ we see that $e_{12} = e_1 - e_2 = s^2$ and $e_{13} = e_1 - e_3 = t^2$ are squares in K. More precisely, by Lemma 6.6.9, the Edwards parameter $d(s,t)$ assigned to these points is exactly as indicated in the statement. Clearly $d(s,t) \neq 0, 1$.

(iii) Take e_1, e_2 and e_3 as in *(ii)*. Then by assumption, also $e_{23} = e_2 - e_3 = t^2 - s^2 = r$ is a square; thus $K(\sqrt{-1}, \sqrt{e_{12}}, \sqrt{e_{13}}, \sqrt{e_{23}}) = K(\sqrt{-1})$ and we conclude by Remark 6.6.8. □

Remark 6.10.2. *(i)* Under the map $K^2 \to K \setminus \{0,1\}, (s,t) \mapsto d(s,t)$, the elliptic curves $\{E_{r,s,t}\}_{r,s,t}$ become isomorphic to the elliptic curves $\{E_{1,d}\}_d$, defined in (6.2) with d a square in K. In this case, the reader can compare with Theorem 6.9.1. More generally,

d	$N(E_d)$	prime fact. of $N(E_d)$	an. rank	Algorithm 14	Dockchitser
-9930	1577837280	$2^5 \cdot 3 \cdot 5 \cdot 331 \cdot 9931$	1	501.959869932160	501.959726754146
3678	216384096	$2^5 \cdot 3 \cdot 613 \cdot 3677$	1	502.401396998320	502.398279446343
4938	390062496	$2^5 \cdot 3 \cdot 823 \cdot 4937$	1	487.393905081071	487.395541243312
5928	70270512	$2^4 \cdot 3 \cdot 13 \cdot 19 \cdot 5927$	1	463.825041040073	463.825507014968
7994	1022336672	$2^5 \cdot 7 \cdot 571 \cdot 7993$	1	484.921298063484	484.922884220681
8423	1135016096	$2^5 \cdot 4211 \cdot 8423$	1	472.556689785723	472.554902645102
9008	81135056	$2^4 \cdot 563 \cdot 9007$	1	498.696434325650	498.693078391666

Table 6.5: The curves E_d with $\#\text{Ш}(E_d) = 22^2 = 484$ and comparison of Algorithm 14 with the build-in Dockchitser function from [S+20]

remark that $E_{d(s,t)}$ is birationally equivalent over K to the curve defined by $sx(y^2-1) = ty(x^2-1)$ in Huff's form (see [JTV10, §2.2]).

(ii) Proposition 6.10.1 *(iii)* gives a way to parametrize, via Pythagorean triples, elliptic curves together with a minimal extension of K (i.e. $K(\sqrt{-1})$) over which all their birationally equivalent Edwards curves are defined.

6.10.2 Computations for Example 6.6.10

(i) A representative for the class $\mathfrak{Ell}(1728)$ is the curve E_a defined by $y^2 = x^3 + ax$ with $a \in K$. The cubic defining E_a factors as $x(x-\sqrt{-a})(x+\sqrt{-a})$ over $\overline{K}$ and Remark 6.6.8 gives $K(E[4]) = K(\sqrt{-1}, \sqrt{2}, \sqrt[4]{-a}) = K(\zeta_8, \sqrt[4]{-a})$ where ζ_8 is an 8-th root of unity. Denoting $e_1 = 0, e_2 = \sqrt{-a}, e_3 = -\sqrt{-a}$, we obtain

$$\begin{aligned} e_{12} = -\sqrt{-a} \quad &, \quad \sqrt{e_{12}} = i\sqrt{i}\sqrt[4]{a} \\ e_{13} = \sqrt{-a} \quad &, \quad \sqrt{e_{13}} = \sqrt{i}\sqrt[4]{a} \\ e_{23} = 2\sqrt{-a} \quad &, \quad \sqrt{e_{23}} = \sqrt{2}\sqrt{i}\sqrt[4]{a}\,. \end{aligned}$$

By Lemma 6.6.9, we associate to points of order 4 doubling to $(e_1, 0)$, the following d:

$$d = \left(\frac{i\sqrt{i} \mp \sqrt{i}}{i\sqrt{i} \pm \sqrt{i}}\right)^2 = -1\,.$$

For points doubling to $(e_2, 0)$, resp. $(e_3, 0)$, we compute d by

$$\left(\frac{\sqrt{i} \mp \sqrt{2}\sqrt{i}}{\sqrt{i} \pm \sqrt{2}\sqrt{i}}\right)^2 = 17 \mp 12\sqrt{2} \quad , \quad \left(\frac{i\sqrt{i} \mp \sqrt{2}i\sqrt{i}}{i\sqrt{i} \pm \sqrt{2}i\sqrt{i}}\right)^2 = 17 \mp 12\sqrt{2}\,,$$

respectively, and finally deduce that $\mathrm{Edw}_{\overline{K}}(1728) = \{\mathcal{E}_{-1}, \mathcal{E}_{17\mp 12\sqrt{2}}\}$.

(ii) A representative for the class $\mathfrak{Ell}(0)$ is the curve E_b given by $y^2 = x^3 + b$ with $b \in K$. The cubic defining E_b factors as $(x-e_1)(x-e_2)(x-e_3)$ over $\overline{K}$ with $e_1 = -\sqrt[3]{b}$, $e_2 = \frac{\sqrt[3]{b}}{2}(1+i\sqrt{3})$

and $e_3 = \frac{\sqrt[3]{b}}{2}(1 - i\sqrt{3})$, where $i^2 = -1$. We obtain

$$\begin{aligned} e_{12} &= -\sqrt[3]{b} \cdot \frac{3 + i\sqrt{3}}{2} \quad , \quad \sqrt{e_{12}} = i\sqrt[6]{b}\sqrt{\frac{3 + i\sqrt{3}}{2}} \\ e_{13} &= -\sqrt[3]{b} \cdot \frac{3 - i\sqrt{3}}{2} \quad , \quad \sqrt{e_{13}} = i\sqrt[6]{b}\sqrt{\frac{3 - i\sqrt{3}}{2}} \\ e_{23} &= \sqrt[3]{b} \cdot i\sqrt{3} \quad , \quad \sqrt{e_{23}} = \sqrt{i}\sqrt[6]{b}\sqrt[4]{3}\,. \end{aligned}$$

and Remark 6.6.8 gives $K(E[4]) = K(i, \sqrt[4]{3}, \sqrt[3]{b}, \sqrt{(1 - i\sqrt{3})/2})$. According to Lemma 6.6.9 we associate to points of order 4 doubling to $(e_i, 0)$ for $i = 1, 2, 3$ the following coefficient d, in this order:

$$\left(\frac{\sqrt{3 + i\sqrt{3}} \mp \sqrt{3 - i\sqrt{3}}}{\sqrt{3 + i\sqrt{3}} \pm \sqrt{3 - i\sqrt{3}}}\right)^2 = -7 \pm 4\sqrt{3}\,, \quad \left(\frac{\sqrt{\frac{1}{2}(3 + i\sqrt{3})} \mp \sqrt{i}\sqrt[4]{3}}{\sqrt{\frac{1}{2}(3 + i\sqrt{3})} \pm \sqrt{i}\sqrt[4]{3}}\right)^2 = -7 \pm 4\sqrt{3}\,,$$

$$\left(\frac{\sqrt{\frac{1}{2}(3 - i\sqrt{3})} \mp i\sqrt{i}\sqrt[4]{3}}{\sqrt{\frac{1}{2}(3 - i\sqrt{3})} \pm i\sqrt{i}\sqrt[4]{3}}\right)^2 = -7 \mp 4\sqrt{3}\,.$$

It follows that $\mathrm{Edw}_{\overline{K}}(0) = \{\mathcal{E}_{-7 \pm 4\sqrt{3}}\}$.

List of Algorithms

List of Tables

Bibliography

[AB15] Benny Applebaum and Zvika Brakerski. *Obfuscating Circuits via Composite-Order Graded Encoding*, pages 528–556. Springer Berlin Heidelberg, Berlin, Heidelberg, 2015.

[AG12] Omran Ahmadi and Robert Granger. On isogeny classes of Edwards curves over finite fields. *J. Number Theory*, 132(6):1337–1358, 2012.

[Ajt96] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 99–108, New York, NY, USA, 1996. Association for Computing Machinery.

[Ajt06] M. Ajtai. Generating random lattices according to the invariant distribution. Draft, 2006.

[AKS01] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, STOC '01, page 601–610, New York, NY, USA, 2001. Association for Computing Machinery.

[Alb17] Martin R. Albrecht. On dual lattice attacks against small-secret LWE and parameter choices in helib and SEAL. In *Advances in Cryptology - EUROCRYPT 2017, Proceedings, Part II*, volume 10211 of *Lecture Notes in Computer Science*, pages 103–129, 2017.

[ALNR11] Christophe Arène, Tanja Lange, Michael Naehrig, and Christophe Ritzenthaler. Faster computation of the Tate pairing. *J. Number Theory*, 131(5):842–857, 2011. With supplementary material available online.

[ALRP08] Julián Aguirre, Álvaro Lozano-Robledo, and Juan Carlos Peral. Elliptic curves of maximal rank. In *Proceedings of the "Segundas Jornadas de Teoría de Números"*, Bibl. Rev. Mat. Iberoamericana, pages 1–28. Rev. Mat. Iberoamericana, Madrid, 2008.

[Ban93] Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(1):625–635, 1993.

[BBJ+08] Daniel J. Bernstein, Peter Birkner, Marc Joye, Tanja Lange, and Christiane Peters. Twisted edwards curves. In *Proceedings of the Cryptology in Africa 1st International Conference on Progress in Cryptology*, AFRICACRYPT'08, pages 389–405, Berlin, Heidelberg, 2008. Springer-Verlag.

[BBLP13] Daniel J. Bernstein, Peter Birkner, Tanja Lange, and Christiane Peters. ECM using Edwards curves. *Math. Comp.*, 82(282):1139–1179, 2013.

[BCP97] Wieb Bosma, John Cannon, and Catherine Playoust. The Magma algebra system. I. The user language. *J. Symbolic Comput.*, 24(3-4):235–265, 1997. Computational algebra and number theory (London, 1993).

[BDL+11] Daniel J. Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang. High-speed high-security signatures. In Bart Preneel and Tsuyoshi Takagi, editors, *Cryptographic Hardware and Embedded Systems – CHES 2011*, pages 124–142, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[Ber06] Daniel J. Bernstein. Curve25519: New Diffie-Hellman Speed Records. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *Public Key Cryptography - PKC 2006, 9th International Conference on Theory and Practice of Public-Key Cryptography, New York, NY, USA, April 24-26, 2006, Proceedings*, volume 3958 of *Lecture Notes in Computer Science*, pages 207–228. Springer, 2006.

[BF01] Dan Boneh and Matthew K. Franklin. Identity-based encryption from the weil pairing. In *Proceedings of the 21st Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '01, page 213–229, Berlin, Heidelberg, 2001. Springer-Verlag.

[BL07] Daniel J. Bernstein and Tanja Lange. Faster addition and doubling on elliptic curves. In *Advances in Cryptology - ASIACRYPT 2007, 13th International Conference on the Theory and Application of Cryptology and Information Security, Kuching, Malaysia, December 2-6, 2007, Proceedings*, volume 4833 of *Lecture Notes in Computer Science*, pages 29–50. Springer, 2007.

[BLS01] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the weil pairing. In Colin Boyd, editor, *Advances in Cryptology — ASIACRYPT 2001*, pages 514–532, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.

[BNNT11] Éric Brier, David Naccache, Phong Q. Nguyen, and Mehdi Tibouchi. Modulus fault attacks against rsa-crt signatures. In Bart Preneel and Tsuyoshi Takagi, editors, *Cryptographic Hardware and Embedded Systems – CHES 2011*, pages 192–206, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[BPV98] Victor Boyko, Marcus Peinado, and Ramarathnam Venkatesan. Speeding up discrete log and factoring based schemes via precomputations", booktitle="advances in cryptology — eurocrypt'98. pages 221–235, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.

[BS03] Dan Boneh and Alice Silverberg. Applications of multilinear forms to cryptography. In *Topics in algebraic and noncommutative geometry (Luminy/Annapolis, MD, 2001)*, volume 324 of *Contemp. Math.*, pages 71–90. Amer. Math. Soc., Providence, RI, 2003.

[BSD65] B. J. Birch and H. P. F. Swinnerton-Dyer. Notes on elliptic curves. II. *J. Reine Angew. Math.*, 218:79–108, 1965.

[BSZ14] Manjul Bhargava, Christopher Skinner, and Wei Zhang. A majority of elliptic curves over q satisfy the birch and swinnerton-dyer conjecture, 2014.

[Cas71] J.W.S. Cassels. *An Introduction to the Geometry of Numbers*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 1971.

[CFL+16] Jung Hee Cheon, Pierre-Alain Fouque, Changmin Lee, Brice Minaud, and Hansol Ryu. Cryptanalysis of the new clt multilinear map over the integers. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016*, pages 509–536, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.

[CG20] Jean-Sébastien Coron and Agnese Gini. A polynomial-time algorithm for solving the hidden subset sum problem. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part II*, volume 12171 of *Lecture Notes in Computer Science*, pages 3–31. Springer, 2020.

[CGH+15] Jean-Sébastien Coron, Craig Gentry, Shai Halevi, Tancrède Lepoint, Hemanta K. Maji, Eric Miles, Mariana Raykova, Amit Sahai, and Mehdi Tibouchi. Zeroizing without low-level zeroes: New MMAP attacks and their limitations. In *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 247–266. Springer, 2015.

[CGH17] Yilei Chen, Craig Gentry, and Shai Halevi. Cryptanalyses of candidate branching program obfuscators. In *Advances in Cryptology - EUROCRYPT 2017 - Proceedings, Part III*, pages 278–307, 2017.

[CH13] Henry Cohn and Nadia Heninger. Approximate common divisors via lattices. *The Open Book Series*, 1(1):271–293, 2013.

[Cha85] Komaravolu Chandrasekharan. *Elliptic functions*, volume 281. Springer-Verlag, 1985.

[CHL+15] Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. In *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 3–12. Springer, 2015.

[CLLT16] Jean-Sébastien Coron, Moon Sung Lee, Tancrède Lepoint, and Mehdi Tibouchi. Cryptanalysis of GGH15 multilinear maps. In *Advances in Cryptology - CRYPTO 2016 - Proceedings, Part II*, pages 607–628, 2016.

[CLLT17] Jean-Sébastien Coron, Moon Sung Lee, Tancrède Lepoint, and Mehdi Tibouchi. Zeroizing attacks on indistinguishability obfuscation over CLT13. In *Public-Key Cryptography - PKC 2017 - Proceedings, Part I*, pages 41–58, 2017.

[CLT13] Jean-Sébastien Coron, Tancrède Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In Ran Canetti and Juan A. Garay, editors, *CRYPTO*, volume 8042 of *LNCS*, pages 476–493. Springer, 2013.

[CLT15] Jean-Sébastien Coron, Tancrède Lepoint, and Mehdi Tibouchi. New multilinear maps over the integers. In Rosario Gennaro and Matthew Robshaw, editors, *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I*, volume 9215 of *Lecture Notes in Computer Science*, pages 267–286. Springer, 2015.

[CN11] Yuanmi Chen and Phong Q. Nguyen. BKZ 2.0: Better lattice security estimates. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*, volume 7073 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2011.

[CN12] Yuanmi Chen and Phong Q. Nguyen. Faster algorithms for approximate common divisors: Breaking fully-homomorphic-encryption challenges over the integers. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, pages 502–519, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

[CN19a] Jean-Sébastien Coron and Luca Notarnicola. Cryptanalysis of CLT13 Multilinear Maps with Independent Slots. In *Advances in Cryptology - ASIACRYPT 2019 - Proceedings, Part II*, pages 356–385, 2019.

[CN19b] Jean-Sébastien Coron and Luca Notarnicola. Source code for 'Cryptanalysis of CLT13 Multilinear Maps with Independent Slots'. *GitHub*, 2019. https://github.com/lucanotarnicola/Cryptanalysis-CLT13.

[CNT10] Jean-Sébastien Coron, David Naccache, and Mehdi Tibouchi. Fault attacks against emv signatures. In Josef Pieprzyk, editor, *Topics in Cryptology - CT-RSA 2010, The Cryptographers' Track at the RSA Conference 2010, San Francisco, CA, USA, March 1-5, 2010. Proceedings*, volume 5985 of *Lecture Notes in Computer Science*, pages 208–220. Springer, 2010.

[CNT12] Jean-Sébastien Coron, David Naccache, and Mehdi Tibouchi. Public key compression and modulus switching for fully homomorphic encryption over the integers. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, pages 446–464, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

[CNW20a] Jean-Sébastien Coron, Luca Notarnicola, and Gabor Wiese. Simultaneous diagonalization of incomplete matrices and applications. *Proceedings of the Fourteenth Algorithmic Number Theory Symposium (ANTS-XIV), edited by Steven Galbraith, Open Book Series 4, Mathematical Sciences Publishers, Berkeley*, pages 127–142, 2020.

[CNW20b] Jean-Sébastien Coron, Luca Notarnicola, and Gabor Wiese. Source code for 'Simultaneous Diagonalization of Incomplete Matrices and Applications'. *GitHub*, 2020. https://github.com/lucanotarnicola/Simultaneous-Diagonalization-Incomplete-Matrices.

[Con] Keith Conrad. Dual modules.

[CP19] Jean-Sébastien Coron and Hilder V. L. Pereira. On Kilian's Randomization of Multilinear Map Encodings. In *Advances in Cryptology - ASIACRYPT 2019 - Proceedings, Part II*, pages 325–355, 2019.

[Cre97] J. E. Cremona. *Algorithms for modular elliptic curves*. Cambridge University Press, Cambridge, second edition, 1997.

[CS18] Craig Costello and Benjamin Smith. Montgomery curves and their arithmetic. *Journal of Cryptographic Engineering*, 8(3):227–240, 2018.

[CSV18] Jingwei Chen, Damien Stehlé, and Gilles Villard. Computing an LLL-reduced Basis of the Orthogonal Lattice. In *Proceedings of the 2018 ACM International Symposium on Symbolic and Algebraic Computation*, pages 127–133, 2018.

[CVW18] Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. GGH15 beyond permutation branching programs: Proofs, attacks, and candidates. In *Advances in Cryptology - CRYPTO 2018 - Proceedings, Part II*, pages 577–607, 2018.

[DH82] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. In *Secure communications and asymmetric cryptosystems*, volume 69 of *AAAS Sel. Sympos. Ser.*, pages 143–180. Westview, Boulder, CO, 1982.

[DT14] Jintai Ding and Chengdong Tao. A new algorithm for solving the approximate common divisor problem and cryptanalysis of the FHE based on GACD. *IACR Cryptol. ePrint Arch.*, 2014:42, 2014.

[Edw07] Harold Edwards. A normal form for elliptic curves. *Bulletin of the American Mathematical Society*, 44(3):393–422, 2007.

[EK20] Noam Elkies and Zev Klagsbrun. *Proceedings of the Fourteenth Algorithmic Number Theory Symposium (ANTS-XIV), Mathematical Sciences Publishers, Berkeley*, 2020.

[FRS17] Rex Fernando, Peter M. R. Rasmussen, and Amit Sahai. Preventing CLT attacks on obfuscation with linear overhead. In *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part III*, pages 242–271, 2017.

[FT33] Willy Feller and Erhard Tornier. Mengentheoretische Untersuchung von Eigenschaften der Zahlenreihe. *Math. Ann.*, 107(1):188–232, 1933.

[Gal05] S. Galbraith. *Pairings*, page 183–214. London Mathematical Society Lecture Note Series. Cambridge University Press, 2005.

[Gal12] Steven D. Galbraith. *Mathematics of Public Key Cryptography*. Cambridge University Press, USA, 1st edition, 2012.

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 169–178. ACM, 2009.

[GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT*, volume 7881 of *LNCS*, pages 1–17. Springer, 2013.

[GGH+13b] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS*, pages 40–49. IEEE Computer Society, 2013.

[GGH+13c] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 40–49. IEEE Computer Society, 2013.

[GGH+13d] Sanjam Garg, Craig Gentry, Shai Halevi, Amit Sahai, and Brent Waters. Attribute-based encryption for circuits from multilinear maps. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 479–499. Springer, 2013.

[GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC*, volume 9015 of *LNCS*, pages 498–527, 2015.

[GGM16] Steven D Galbraith, Shishay W Gebregiyorgis, and Sean Murphy. Algorithms for the approximate common divisor problem. *LMS Journal of Computation and Mathematics*, 19(A):58–72, 2016.

[GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 467–476. ACM, 2013.

[GHGKN06a] Nicolas Gama, Nick Howgrave-Graham, Henrik Koy, and Phong Q. Nguyen. Rankin's constant and blockwise lattice reduction. In Cynthia Dwork, editor, *Advances in Cryptology - CRYPTO 2006*, pages 112–130, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

[GHGKN06b] Nicolas Gama, Nick Howgrave-Graham, Henrik Koy, and Phong Q. Nguyen. Rankin's constant and blockwise lattice reduction. In Cynthia Dwork, ed-

itor, *Advances in Cryptology - CRYPTO 2006*, pages 112–130, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

[GJLR16] Enrique González-Jiménez and Álvaro Lozano-Robledo. Elliptic curves with abelian division fields. *Math. Z.*, 283(3-4):835–859, 2016.

[GLSW15] Craig Gentry, Allison Bishop Lewko, Amit Sahai, and Brent Waters. Indistinguishability obfuscation from the multilinear subgroup elimination assumption. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 151–170, 2015.

[GLW14] Craig Gentry, Allison B. Lewko, and Brent Waters. Witness encryption from instance independent assumptions. In *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, pages 426–443, 2014. `https://eprint.iacr.org/2014/273`.

[GN08a] Nicolas Gama and Phong Q. Nguyen. Finding short lattice vectors within mordell's inequality. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, page 207–216, New York, NY, USA, 2008. Association for Computing Machinery.

[GN08b] Nicolas Gama and Phong Q. Nguyen. Predicting Lattice Reduction. In *Advances in Cryptology - EUROCRYPT 2008, 27th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Istanbul, Turkey, April 13-17, 2008. Proceedings*, volume 4965 of *Lecture Notes in Computer Science*, pages 31–51. Springer, 2008.

[Gol79] Dorian Goldfeld. Conjectures on elliptic curves over quadratic fields. In *Number theory, Carbondale 1979 (Proc. Southern Illinois Conf., Southern Illinois Univ., Carbondale, Ill., 1979)*, volume 751 of *Lecture Notes in Math.*, pages 108–118. Springer, Berlin, 1979.

[Gou94] Fernando Q. Gouvêa. "A marvelous proof". *Amer. Math. Monthly*, 101(3):203–222, 1994.

[HHGP+03] Jeffrey Hoffstein, Nick Howgrave-Graham, Jill Pipher, Joseph H. Silverman, and William Whyte. Ntrusign: Digital signatures using the ntru lattice. In Marc Joye, editor, *Topics in Cryptology — CT-RSA 2003*, pages 122–140, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.

[HJ16] Yupu Hu and Huiwen Jia. Cryptanalysis of GGH map. In *Advances in Cryptology - EUROCRYPT 2016 - Proceedings, Part I*, pages 537–565, 2016.

[HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. Ntru: A ring-based public key cryptosystem. In Joe P. Buhler, editor, *Algorithmic Number Theory*, pages 267–288, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.

[HPS08] Jeffrey Hoffstein, Jill Pipher, and J.H. Silverman. *An Introduction to Mathematical Cryptography*. Springer Publishing Company, Incorporated, 1 edition, 2008.

[HPS11a] Guillaume Hanrot, Xavier Pujol, and Damien Stehlé. Analyzing blockwise lattice algorithms using dynamical systems. In Phillip Rogaway, editor, *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*, pages 447–464. Springer, 2011.

[HPS11b] Guillaume Hanrot, Xavier Pujol, and Damien Stehlé. Analyzing blockwise lattice algorithms using dynamical systems. In *Proceedings of the 31st Annual Conference on Advances in Cryptology*, CRYPTO'11, pages 447–464, Berlin, Heidelberg, 2011. Springer-Verlag.

[Hus04] Dale Husemöller. *Elliptic curves*, volume 111 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 2004. With appendices by Otto Forster, Ruth Lawrence and Stefan Theisen.

[HWCD08] Huseyin Hisil, Kenneth Koon-Ho Wong, Gary Carter, and Ed Dawson. Twisted Edwards curves revisited. 5350:326–343, 2008.

[HY15] Dustin Hinkel and Matthew P. Young. The distribution of central values of elliptic curve l-functions. *Journal of Number Theory*, 156:15–20, 2015.

[Jou04] Antoine Joux. A one round protocol for tripartite Diffie-Hellman. *J. Cryptology*, 17(4):263–276, 2004.

[JQNP01] Marc Joye, editor="Koç Çetin K. Quisquater, Jean-Jacques", David Naccache, and Christof Paar. Hessian elliptic curves and side-channel attacks. In *Cryptographic Hardware and Embedded Systems — CHES 2001*, pages 402–410, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.

[JTV10] Marc Joye, Mehdi Tibouchi, and Damien Vergnaud. Huff's model for elliptic curves. In *Algorithmic number theory*, volume 6197 of *Lecture Notes in Comput. Sci.*, pages 234–250. Springer, Berlin, 2010.

[Ken04] Ono Ken. *The web of modularity : arithmetic of the coefficients of modular forms and q-series / Ken Ono*. Regional conference series in mathematics. Published for the Conference Board of the Mathematical Sciences by the American Mathematical Society, Providence, Rhodes Island, right 2004.

[KK98] Chang Heon Kim and Ja Kyung Koo. Arithmetic of the modular function $j_{1,4}$. *Acta Arith.*, 84(2):129–143, 1998.

[KM88] M. A. Kenku and F. Momose. Torsion points on elliptic curves defined over quadratic fields. *Nagoya Math. J.*, 109:125–149, 1988.

[KN12] Sheldon Kamienny and Filip Najman. Torsion groups of elliptic curves over quadratic fields. *Acta Arith.*, 152(3):291–305, 2012.

[Kna92] Anthony W. Knapp. *Elliptic curves*, volume 40 of *Mathematical Notes*. Princeton University Press, Princeton, NJ, 1992.

[Kob87] Neal Koblitz. Elliptic curve cryptosystems. *Mathematics of Computation*, 48(177):203–209, January 1987.

[Kol88] V. A. Kolyvagin. Finiteness of $E(\mathbf{Q})$ and Sha$(E, \mathbf{Q})$ for a subclass of Weil curves. *Izv. Akad. Nauk SSSR Ser. Mat.*, 52(3):522–540, 670–671, 1988.

[Lan05] Tanja Lange. Mathematical countermeasures against side-channel attacks. In Henri Cohen, Gerhard Frey, Roberto Avanzi, Christophe Doche, Tanja Lange, Kim Nguyen, and Frederik Vercauteren, editors, *Handbook of Elliptic and Hyperelliptic Curve Cryptography*, pages 687–714. Chapman and Hall/CRC, 2005.

[Len87] H. W. Lenstra, Jr. Factoring integers with elliptic curves. *Ann. of Math. (2)*, 126(3):649–673, 1987.

[Liu09] Zhi-Guo Liu. Addition formulas for Jacobi theta functions, Dedekind's eta function, and Ramanujan's congruences. *Pacific J. Math.*, 240(1):135–150, 2009.

[LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261(4):515–534, 1982.

[LMF20a] The LMFDB Collaboration. The L-functions and modular forms database. `http://www.lmfdb.org`, 2020.

[LMF20b] The LMFDB Collaboration. The L-functions and modular forms database. `https://www.lmfdb.org/EllipticCurve/Q/stats`, 2020. Elliptic curves over $\mathbb{Q}$: Statistics.

[LN20] Jianwei Li and Phong Q. Nguyen. A complete analysis of the bkz lattice reduction algorithm. Cryptology ePrint Archive, Report 2020/1237, 2020. `https://eprint.iacr.org/2020/1237`.

[LO85] J. C. Lagarias and A. M. Odlyzko. Solving low-density subset sum problems. *J. Assoc. Comput. Mach.*, 32(1):229–246, 1985.

[LWXZ11] Mingjie Liu, Xiaoyun Wang, Guangwu Xu, and Xuexin Zheng. Shortest Lattice Vectors in the Presence of Gaps. *IACR Cryptology ePrint Archive*, 2011:139, 2011.

[LWZ13] Liangze Li, Hongfeng Wu, and Fan Zhang. Pairing computation on edwards curves with high-degree twists. In Dongdai Lin, Shouhuai Xu, and Moti Yung, editors, *Information Security and Cryptology - 9th International Conference, Inscrypt 2013, Guangzhou, China, November 27-30, 2013, Revised Selected Papers*, volume 8567 of *Lecture Notes in Computer Science*, pages 185–200. Springer, 2013.

[Mac02] Allan J. MacLeod. The efficient computation of some generalised exponential integrals. *Journal of Computational and Applied Mathematics*, 148(2):363–374, 2002.

[Mic11] Daniele Micciancio. *The Geometry of Lattice Cryptography*, pages 185–210. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

[Mil86] Victor S. Miller. Use of elliptic curves in cryptography. In Hugh C. Williams, editor, *Advances in Cryptology — CRYPTO '85 Proceedings*, pages 417–426, Berlin, Heidelberg, 1986. Springer Berlin Heidelberg.

[MO90] J. E. Mazo and A. M. Odlyzko. Lattice points in high-dimensional spheres. *Monatsh. Math.*, 110(1):47–61, 1990.

[Mon87] Peter L. Montgomery. Speeding the Pollard and elliptic curve methods of factorization. *Math. Comp.*, 48(177):243–264, 1987.

[MS16] Dustin Moody and Daniel Shumow. Analogues of Vélu's formulas for isogenies on alternate models of elliptic curves. *Math. Comp.*, 85(300):1929–1951, 2016.

[MSZ16] Eric Miles, Amit Sahai, and Mark Zhandry. Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over GGH13. In *Advances in Cryptology - CRYPTO 2016 - Proceedings, Part II*, pages 629–658, 2016.

[NS97] Phong Q. Nguyen and Jacques Stern. Merkle-Hellman Revisited: A Cryptanalysis of the Qu-Vanstone Cryptosystem Based on Group Factorizations. In *Advances in Cryptology - CRYPTO '97, 17th Annual International Cryptology Conference*, pages 198–212, 1997.

[NS98] Phong Q. Nguyen and Jacques Stern. Cryptanalysis of a fast public key cryptosystem presented at SAC '97. In Stafford E. Tavares and Henk Meijer, editors, *Selected Areas in Cryptography '98, SAC'98, Kingston, Ontario, Canada, August 17-18, 1998, Proceedings*, volume 1556 of *Lecture Notes in Computer Science*, pages 213–218. Springer, 1998.

[NS99] Phong Q. Nguyen and Jacques Stern. The hardness of the hidden subset sum problem and its cryptographic implications. In *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference*, pages 31–46, 1999.

[NS01] Phong Q. Nguyen and Jacques Stern. The two faces of lattices in cryptology. In *Revised Papers from the International Conference on Cryptography and Lattices*, CaLC, pages 146–180, London, UK, UK, 2001. Springer-Verlag.

[NS05] Phong Q. Nguyen and Jacques Stern. Adapting density attacks to low-weight knapsacks. In *Advances in Cryptology - ASIACRYPT 2005, 11th International Conference on the Theory and Application of Cryptology and Information Security*, pages 41–58, 2005.

[NS09] Phong Q. Nguyen and Damien Stehlé. An LLL algorithm with quadratic complexity. *SIAM J. Comput.*, 39(3):874–903, 2009.

[NSV11] Andrew Novocin, Damien Stehlé, and Gilles Villard. An LLL-reduction algorithm with quasi-linear time complexity: extended abstract. In Lance Fortnow and Salil P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on*

Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011, pages 403–412. ACM, 2011.

[NV10] Phong Q. Nguyen and Brigitte Vallée, editors. *The LLL Algorithm - Survey and Applications*. Information Security and Cryptography. Springer, 2010.

[NW21] Luca Notarnicola and Gabor Wiese. Source code for 'The Hidden Lattice Problem'. *GitHub*, 2021. https://github.com/lucanotarnicola/HLP.

[PZ11] Yanbin Pan and Feng Zhang. A note on the density of the multiple subset sum problems. *IACR Cryptology ePrint Archive*, 2011:525, 2011.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 84–93. ACM, 2005.

[S+17] W. A. Stein et al. *Sage Mathematics Software (Version 8.0)*. The Sage Development Team, 2017. `http://www.sagemath.org`.

[S+20] W. A. Stein et al. *Sage Mathematics Software (Version 9.2)*. The Sage Development Team, 2020. `http://www.sagemath.org`.

[Sch85] René Schoof. Elliptic curves over finite fields and the computation of square roots mod p. *Math. Comp.*, 44(170):483–494, 1985.

[Sch87] Claus-Peter Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.

[Sch14] Edward F. Schaefer. Corrigendum to "Class groups and Selmer groups" [J. Number Theory 56 (1) (1996) 79–114] [mr1370197]. *J. Number Theory*, 135:390, 2014.

[SE94] C. P. Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathematical Programming*, 66(1):181–199, 1994.

[Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. 26(5), 1997.

[Sil94] Joseph H. Silverman. *Advanced topics in the arithmetic of elliptic curves*, volume 151 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994.

[Sil09] Joseph H. Silverman. *The Arithmetic of Elliptic Curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, Dordrecht, second edition, 2009.

[Str15] Marco Streng. Generators of the group of modular units for $\Gamma_1(N)$ over $\mathbb{Q}$. *arXiv preprint arXiv:1503.08127*, 2015.

[TW95] Richard Taylor and Andrew Wiles. Ring-theoretic properties of certain Hecke algebras. *Ann. of Math. (2)*, 141(3):553–572, 1995.

[VDGHV10] Marten Van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully homomorphic encryption over the integers. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 24–43. Springer, 2010.

[Was03] Lawrence C. Washington. *Elliptic curves*. Discrete Mathematics and its Applications (Boca Raton). Chapman & Hall/CRC, Boca Raton, FL, 2003. Number theory and cryptography.

[Wat07] Mark Watkins. Rank distribution in a family of cubic twists. In *Ranks of elliptic curves and random matrix theory*, volume 341 of *London Math. Soc. Lecture Note Ser.*, pages 237–246. Cambridge Univ. Press, Cambridge, 2007.

[Wil95] Andrew Wiles. Modular elliptic curves and Fermat's Last Theorem. *Ann. of Math. (2)*, 141(3):443–551, 1995.

[Wil06] Andrew Wiles. The Birch and Swinnerton-Dyer conjecture. In *The millennium prize problems*, pages 31–41. Clay Math. Inst., Cambridge, MA, 2006.

[Zim15] Joe Zimmerman. How to obfuscate programs directly. In *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II*, pages 439–467, 2015.

[ZK87] D. Zagier and G. Kramarz. Numerical investigations related to the L-series of certain elliptic curves. *J. Indian Math. Soc. (N.S.)*, 52:51–69 (1988), 1987.