



HAL
open science

Endomorphism Rings in Cryptography

Gaetan Bisson

► **To cite this version:**

Gaetan Bisson. Endomorphism Rings in Cryptography. Computer Science [cs]. Institut National Polytechnique de Lorraine - INPL; Technische Universiteit Eindhoven, 2011. English. NNT: 2011INPL047N . tel-01749554v2

HAL Id: tel-01749554

<https://theses.hal.science/tel-01749554v2>

Submitted on 18 Jul 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ENDOMORPHISM RINGS IN CRYPTOGRAPHY

Gaetan BISSON

COPYRIGHT © MMXI BY Gaetan BISSON

Unmodified copies of this document may be freely distributed.

A catalog record is available from the Eindhoven University of Technology Library.

ISBN: 90-386-2519-7

COVER PICTURE: 尾形月耕の「龍昇天」
(Ogata Gekkō's "dragon ascending" [the mount Fuji volcano])

Endomorphism Rings in Cryptography

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van de
rector magnificus, prof.dr.ir. C.J. van Duijn, voor een
commissie aangewezen door het College voor
Promoties in het openbaar te verdedigen
op donderdag 14 juli 2011 om 14.00 uur

door

Gaëtan Bisson

geboren te Les Ulis, Frankrijk

Dit proefschrift is goedgekeurd door de promotor:

prof.dr. Tanja Lange

Copromotor:
dr.habil. Pierrick Gaudry

ENDOMORPHISM RINGS IN CRYPTOGRAPHY

thèse préparée par

Gaëtan BISSON

LABORATOIRE LORRAIN DE RECHERCHE
EN INFORMATIQUE ET SES APPLICATIONS

&

ÉCOLE NATIONALE SUPÉRIEURE DES MINES DE NANCY

pour l'obtention du doctorat en informatique de l'

INSTITUT NATIONAL POLYTECHNIQUE DE LORRAINE
(ÉCOLE DOCTORALE IAEM LORRAINE)

*présentée et soutenue publiquement le 14 juillet 2011
devant un jury composé de*

PRÉSIDENT:

Arjeh M. COHEN, professeur

Technische Universiteit Eindhoven

DIRECTEURS:

Pierrick GAUDRY, directeur de recherche

Centre National de la Recherche Scientifique

Tanja LANGE, professeur

Technische Universiteit Eindhoven

RAPPORTEURS:

Steven D. GALBRAITH, professeur associé

University of Auckland

David R. KOHEL, professeur

Université de la Méditerranée

Henk C. A. VAN TILBORG, professeur

Technische Universiteit Eindhoven

EXAMINATEURS:

Jean-Marc COUVEIGNES, professeur

Université Toulouse II

Florian HESS, professeur

Universität Oldenburg

À grand-mère doudou

Foreword

Acknowledgments

LITTLE DID I KNOW when I signed up for a PhD project under the joint supervision of Pierrick GAUDRY and Tanja LANGE what great people they were. For the past three years, they have coped with me in shifts, and not only guided my work but aerated my brain through movies, beers, trolls, and smileys. My recent achievements have only been enabled by their ongoing support.

My research work is immensely indebted to Takakazu SATOH, who carefully selected a promising topic for my master's thesis, and to Andrew SUTHERLAND, who I have had the pleasure of working with twice. This work heavily builds upon that of David KOHEL and Steven GALBRAITH, and it is a honor to have them as reviewers for my thesis.

The infliction of reading this manuscript also fell upon Henk VAN TILBORG; he should be thanked twice as he was such a cheerful and diligent *daily boss* in Eindhoven. My defense committee is further completed by Arjeh COHEN, Jean-Marc COUVEIGNES, and Florian HESS, to whom I am most grateful for their enthusiasm about this event.

Throughout my PhD program, I have learned abundantly by osmosis from my colleagues Damien ROBERT and Romain COSSET, and through discussions with David GRUENEWALD in Marseille. Additionally, it was highly scientifically rewarding to be invited to present my work and exchange ideas with Jean-Luc BEUCHAT in Tsukuba, David LUBICZ in Rennes, Christophe RITZENTHALER in Marseille, Andreas ENGE in Bordeaux, Vanessa VITSE in Versailles, and Fabien LAGUILLAUMIE in Caen.

Conferences provide a rich experience where one travels, works, and relaxes all at once; on various occasions, I have had memorable times (combining all the above) with Nicolas BRISEBARRE, Peter SCHWABE, Jérémie DETREY, Anja BECKER, Laurent IMBERT, Peter BIRKNER, and Nicolas ESTIBALS. I extend my heartfelt thanks to all my coworkers as well, and especially to those I have repeatedly bothered with questions or favors, namely Christiane PETERS, Guillaume HANROT, Michael NAEHRIG, Paul ZIMMERMANN, Alexander KRUPPA, Emmanuel THOMÉ, and Dan BERNSTEIN.

For a Frenchman, living in Eindhoven may seem scarier than it actually is. Fortunately, many friends contributed to making my stays there enjoyable, most notably Peter VAN LIESDONK, Shona YU, Antonino SIMONE, and Mayla BRUSÒ; the weekly CASA poker games were greatly relaxing, and I particularly thank Patricio ROSEN and Mark VAN KRAAIJ for the organisational heavy lifting, and Jan-Willem KNOPPER for cooking on so many occasions. My office (and lunch) mates Daniel TRIVELLATO, Bruno ROCHA, Relinde JURRIUS, and Elisa COSTANTE never uttered a word about my irregular work schedule or cursing at the computer in various languages, for which I highly commend them.

My favorite escape opportunity from both of my workplaces was the practice of recreational and competitive sailing; I have had wonderful moments sailing with all my fellow crewmen, and if I may just thank three it would be Jean-François MORIZUR for providing more opportunities than I could accept, and François GARILLOT and Luc HABERT for putting up with the hard task of seconding me.

The École Normale Supérieure provided me with a profusion of social, scientific, and administrative experiences; I am indebted to its staff for providing such a great environment. Since then, most of the weekends I spent in transit in Paris were cheered up by old friends who did not escape the French capital: those people with a spare mattress, Marc SAGE, Pierre-Loïc MÉLIOT, and Mélanie JECKER, but also Pierre & Constance DESCOTES who hosted so many events. On the other hand, some dared going abroad too, and I really enjoyed visiting David DURLEMAN in London, and Jean-Dominique DELLE LUCHE and Pauline PUJO in Berlin.

Back in the “tough” days of *classes préparatoires*, I was lucky enough to befriend Sébastien ALAIWAN, Yannick AST, Jean FELDER, and Jean-Georges MOINEAU, who have always been a joy to see again since, although most of us are not living in the same continent anymore.

My horizons were widely expanded by the free and open culture movement, and I would like to salute those who initiated me to it: in the software category, this ranges from the *salle S* to my fellow Arch Linux developers and, in the music category, this includes SomaFM’s eclectic music directors and talented artists.

Last but not least, my profound gratitude goes to the entire BISSON, LEAU, and JAUZEIN families for their continuous relief and kindness over the past quarter of a century.

Gaetan BISSON
Eindhoven, May 2011

Introduction

Suppose Mr. Athos wishes to write a private message to Mrs. Bonacieux while keeping its contents secret from his Eminence of Richelieu, to whom the courier is most certainly beholden; he could put the message in a safe box whose combination is only known to himself and to Bonacieux, and that would be very costly to break.

Rather than physical devices, cryptography rests on computational power to ensure data security and integrity. Athos and Bonacieux are each given a black box: Athos' is parametrized by a key and transforms messages into unintelligible data called ciphertexts; with the corresponding key, Bonacieux's reverses this operation. Ciphertexts can then be transmitted openly over any medium. Chapter I gives a brief overview of such techniques, with an emphasis on schemes allowing Athos' key to be public: they are only a few decades old and make extensive use of mathematical structures.

Abelian varieties are objects upon which such schemes can be built very efficiently and securely; they are formally introduced in Chapter II, which concisely presents certain of their theoretical aspects, focusing on computations over finite fields. Subsequent chapters, where the original contributions of this thesis are located, are concerned with algorithmic properties related to the endomorphism ring structure of abelian varieties; most of the theoretical background on this topic forms what is known as complex multiplication theory, which Chapter III covers.

An important application of endomorphism rings is the construction of abelian varieties with desirable properties. For instance, many featureful cryptographic schemes have recently been enabled by pairings; to make these schemes practical, abelian varieties endowed with efficient pairings must be generated. Chapter IV discusses this subject, including the work of B. and SATOH (2008) and related results.

The second half of this thesis addresses the problem of computing the endomorphism ring of a prescribed abelian variety, which can be seen as the inverse problem to variety generation. Chapter V recalls prior state-of-the-art methods, all of which have an exponential runtime in the size of the input. It also describes the general structure of isogeny graphs, which is later extensively relied on.

Our subexponential algorithms for computing endomorphism rings of ordinary abelian varieties are first described in Chapter VI in an idealized setting. They exploit complex multiplication theory in its relevance to the structure of isogeny graphs. When specialized to the case of dimension-one abelian varieties, this directly yields highly effective methods which are essentially equivalent to that of B. and SUTHERLAND (2009). Their complexity is rigorously analyzed in Chapter VII, as was done in B. (2011); this chapter ends with a discussion of the results of B. and SUTHERLAND (2011) in this context, from a different perspective than the original article.

Chapter VIII finally explains how our methods can be adapted to be effective in higher dimension, and reports on the implementation of B., COSSET, and ROBERT (2010) enabling the evaluation of general maps between abelian varieties (so-called isogenies), which is an important building block of our algorithms. We conclude by applying our technique to the computation of several illustrative and record examples.

Contributions

2008. Gaetan BISSON and Takakazu SATOH.
“More discriminants with the Brezing-Weng method”.
In: *Progress in Cryptology — INDOCRYPT '08*.
Edited by Dipanwita R. CHOWDHURY, Vincent RIJMEN, and Abhijit DAS.
Volume 5365. Lecture Notes in Computer Science. Springer. Pages 389–399.
DOI: 10.1007/978-3-540-89754-5_30.
2009. Gaetan BISSON and Andrew V. SUTHERLAND.
“Computing the endomorphism ring of an ordinary elliptic curve over a finite field”.
In: *Journal of Number Theory* 131.5. Edited by Neal KOBLITZ and Victor S. MILLER.
Special Issue on Elliptic Curve Cryptography. Pages 815–831.
DOI: 10.1016/j.jnt.2009.11.003.
2010. Gaetan BISSON, Romain COSSET, and Damien ROBERT.
AVIsogenies, a library for computing isogenies between abelian varieties.
Registered at the Agence pour la Protection des Programmes under reference
IDDN.FR.001.440011.000.R.P.2010.000.10000.
URL: <http://avisogenies.gforge.inria.fr/>.
2011. Gaetan BISSON.
Computing endomorphism rings of elliptic curves under the GRH.
arXiv.org: 1101.4323.
2011. Gaetan BISSON and Andrew V. SUTHERLAND.
“A low-memory algorithm for finding short product representations in finite groups”.
In: *Designs, Codes and Cryptography*. To appear.
DOI: 10.1007/s10623-011-9527-8.
2011. Gaetan BISSON.
Computing endomorphism rings of abelian varieties. In preparation.

Contents

Foreword	i
Acknowledgments	i
Introduction	iii
Contributions	iv

ABELIAN VARIETIES IN CRYPTOGRAPHY

I	Panorama of Cryptography	3
1	Symmetric Primitives	3
2	Asymmetric Primitives	7
3	Generic Methods	11
4	Cryptographic Groups	14
	References	18
II	Abelian Varieties	25
1	General Theory	25
2	Practical Settings	29
3	Pairings	33
4	Isogenies	37
	References	40
III	Complex Multiplication	45
1	Endomorphism Rings	45
2	Orders and Ideals	49
3	Plain Complex Multiplication	54
4	Polarized Complex Multiplication	57
	References	62

IV	Pairing-Friendly Varieties	65
1	Cryptographic Requirements	65
2	Complex Multiplication Method	69
3	Elliptic Curve Generation	73
4	Variety Generation	77
	References	80

COMPUTATION OF ENDOMORPHISM RINGS

V	Exponential Methods	87
1	Isogeny Volcanoes	87
2	Higher Dimension	93
3	General Methods	97
4	Supersingular Methods	102
	References	105
VI	Subexponential Method	109
1	Algorithm Overview	109
2	Finding Principal Ideals	115
3	Computing the Action of Ideals	120
4	Practical Computations	124
	References	128
VII	Complexity Analysis	133
1	Orders from Picard Groups	133
2	Picard Groups from Relations	138
3	Relations from Smooth Ideals	141
4	Relations from Thin Air	146
	References	151
VIII	Polarized Method	157
1	Algorithm	157
2	Computing Isogenies	162
3	Practical Computations	167
4	Isogeny Volcanoes	171
	References	178

Index	183
Concluding Remarks	187
Summary	187
Research Prospects	188
Curriculum Vitæ	189

ABELIAN VARIETIES
IN CRYPTOGRAPHY

Panorama of Cryptography

Historically, cryptography has prevalently been employed for secrecy, although over time it has come to provide other features, such as integrity protection and authentication. This chapter concisely presents standard techniques achieving such classical primitives; it serves as both a motivation and practical framework for computational number theory.

1.1 Symmetric Primitives

Early cryptography necessitated a secret, called the *key*, to be shared between the parties involved. Primitives of that lineage are said to be *symmetric*; they are in widespread use and development today, mostly due to their flexible and fast implementations.

CIPHERS

Denote by $\mathbb{S} = \{0, 1\}^{(\mathbb{N})}$ the set of all *strings*, that is, finite sequences of bits.

Definition 1.1.1. Symmetric encryption schemes consist of two families E and D of functions, not necessarily everywhere defined, from \mathbb{S} to \mathbb{S} such that $D_k \circ E_k = \text{Id}_{\text{dom}(E_k)}$ for all strings k .

Intuitively, E and D are the black boxes to provide Athos and Bonacieux: the *cipher* E is parametrized by a *key* k , takes *plaintexts* m as input, and returns *ciphertexts* $E_k(m)$, while the *decipher* D does the converse. His Eminence should be unable to gain any insight on the message m from the sole knowledge of the ciphertext $E_k(m)$; in the strictest sense, this is formalized as *perfect secrecy*, which requires that, for all finite sets of strings M and M' ,

$$\text{Prob}_{k,m}[m \in M \mid E_k(m) \in M'] = \text{Prob}_m[m \in M].$$

Early ciphers, going back to several centuries BC, simply swapped or shifted bytes of the plaintext in a regular fashion derived from the key; for instance, splitting strings as sequences of bytes that encode letters A–Z as integers 0–25, the cipher

$$E_k : (m_i) \mapsto (m_i + k \bmod 26)$$

is still in limited use today with $k = 13$. Similar schemes not obviously as weak have also been designed using larger keys; virtually all have since been broken by the development of frequency analysis.

SHANNON (1949) established the existence and essential uniqueness of a cryptosystem achieving perfect secrecy: the *one-time pad* — it requires a key to be drawn independently and uniformly at random from $\{0, 1\}^n$ for each n -bit plaintext, and returns as ciphertext the bit-by-bit xor of the plaintext and the key. Its practical use is only limited by the ability to carry suitcases full of pads around, prior to doing any encryption.

To mimic its behavior while overcoming the need for lengthy keys transmission, *stream ciphers* (also known as *pseudorandom number generators*), on input a small key called the *seed*, deterministically generate pads to be xored with the plaintext; as before, measurable statistical deviations of such pads from random strings should be avoided. Nowadays, *block ciphers*, which encrypt fixed-length blocks of bits, are the most widely used, and particularly that of DAEMEN and RIJMEN (1999) later standardized as the AES. Procedures for encrypting sequences of blocks, known as *modes of operations*, prevent additional information leakage when handling messages of arbitrary length.

CONCRETE SECURITY

The above overview calls for a more down-to-earth discussion of security aspects: the result of SHANNON (1949) concerns whether the key *can theoretically* be recovered from a certain amount of ciphertext, not how resource-demanding that process is.

One of the cheapest ways of effectively compromising the key is to peek at Athos' notebook, or simply to ask him about it over a nice glass of wine; such *side-channel attacks* will not be discussed here, as we focus on cryptosystems themselves, not their implementations.

Definition 1.1.2. *A cipher E is computationally secure if, for most keys k , it is computationally infeasible to derive plaintexts m from ciphertexts $E_k(m)$.*

“Computationally infeasible” means that, with today's state-of-the-art machines, this computation would take more time than is available, say, billions of years.

Other conditions might be desirable as well; for instance, that the output of E_k cannot feasibly be told apart from that of a random function. However, as our interest will shift to the mathematical building blocks of cryptosystems, this distinction will bear little relevance.

Most cryptosystems do not achieve perfect secrecy, and are thus susceptible to *brute-force attacks*, which decrypt given ciphertexts by trying all possible keys in turn. For “ideal ciphers,” this is the best attack, and for “ideal keys,” which have no special property that reduces the search range, it takes $2^n/2$ runs on average to find an n -bit key.

With today’s technology, the total number of elementary arithmetic operations realistically achievable can be bounded from above by 2^{128} ; keys bearing (at least) 128 bits of entropy are thus recommended. Naturally, this should be tempered by several factors:

- the gravity of the encrypted information;
- the desired lifetime of the cryptosystem;
- the available processing power.

For instance, a news agency broadcasting encrypted live reports to its paying subscribers with different keys each day might only need to withstand limited-resources attacks for 24 hours.

Summing up the above, assessing the security of a cryptosystem calls for a deep understanding of the ways and costs to attack it. MOORE (1965) predicted an exponential growth in available computing power which has been verified for the past four decades; as a consequence, the costs should be considered for increasing key-sizes.

Rather than relying on a rigorous computing model such as the multi-tape universal machines of TURING (1937), we will simply analyze algorithms by looking at both their actual runtime on practical computations, and their long-term behavior embodied in asymptotic bit-complexity estimates. In particular, we disregard quantum-computing models.

To emphasize the need for an asymptotic analysis, denote by $c_E(n)$ the operation count of the best method for attacking a cipher E with n -bit keys: if c_E grows subexponentially, key-sizes are required to increase more than linearly in time to provide a constant level of security, which may eventually prove to be quite cumbersome.

HASH FUNCTIONS

One-way functions formalize the behavior which is expected of ciphers parametrized by unknown keys; they have countless applications, far beyond cryptography, such as hash tables. Like ciphers, they can be defined in a complexity-theoretic way, as *functions that can be evaluated by polynomial-time algorithms, but for which no polynomial-time algorithm can successfully find preimages on more than an exponentially small fraction of the image*.

Since the existence of such functions implies $P \neq NP$, we look for a more practical stance.

Definition I.1.3. *A function $h : \mathbb{S} \rightarrow \mathbb{S}$ is one-way if it is computationally infeasible to find preimages of most of its image. It is also a hash function if its image is contained in $\{0, 1\}^n$ for some n and it is computationally infeasible to find two strings $x \neq x'$ verifying $h(x) = h(x')$.*

Again, additional conditions might be required for specific applications. The *random oracle* is a convenient ideal encompassing most expectations: it is nothing but the Cartesian power by \mathbb{S} of the uniform distribution on n -bit strings, or, more pragmatically, a “map” whose images are drawn uniformly at random from $\{0, 1\}^n$.

Since there typically are at least a few functions (such as constant ones) that are unsuitable, designs using hash functions h are often analyzed by assuming that h has the uniform distribution, and proving that the desired properties hold with overwhelming probability.

Traditionally, hash functions are crafted as a mix of logic gates, but some have also been built on top of mathematical structures, which allows to analyze their behavior much more rigorously. For instance, the construction of CHARLES, LAUTER, and GOREN (2009) involves isogeny graphs of supersingular elliptic curves, a structure that we will investigate later (for completely independent reasons).

PROVABLE SECURITY

Confidently evaluating the complexity c_E of the best attack on a cryptosystem E is a difficult task. Provable cryptography aims at designing cryptosystems on which successful attacks can be *reduced* into disproofs of certain ideal properties of the underlying blocks. However, since many traditional blocks feature components specifically designed to obscure their behavior, assessing the veracity of these ideal properties is not always possible.

Alternatively, the machinery of mathematics provides well-studied building blocks, bundled with tools adapted to rigorous analyses, although this often comes at the expense of slower implementations.

As a prominent example, let us give a result of SHoup (1997) regarding the *discrete logarithm problem*, which is that of inverting the function $\exp_g : n \in \mathbb{Z} \mapsto g^n \in G$, where g is a fixed element of a group G .

Theorem 1.1.4. *In prime-order groups G , no generic algorithm can solve random instances of the discrete logarithm problem in time $o(\sqrt{\#G})$.*

Later, we will rigorously define generic algorithms and explain how they can invert discrete logarithms in time $O(\sqrt{\#G})$; in essence, this theorem states that no attacker using the group as a black box (thus unable to exploit any “special” property) can do better than that.

Assuming that a cryptosystem E builds upon the discrete logarithm problem on a group where generic attacks are the best available, we can often, after some calibration, estimate the value of c_E at finite parameters by its asymptotic behavior: if a key k has about the same size as the group G that E_k uses, then it must be roughly 256-bit long in order to provide an expected 128 bits of symmetric security.

Researches have built cryptographic blocks upon mathematical objects of various kinds: DIFFIE and HELLMAN (1976) used discrete logarithms, MERKLE and HELLMAN (1978) relied on knapsacks, RIVEST, SHAMIR, and ADLEMAN (1978) suggested using integer factorization, MCELIECE (1978) made the case for error-correcting codes, MATSUMOTO and IMAI (1988) employed certain multivariate polynomials, ZÉMOR (1994) exploited Cayley graphs, AJTAI (1996) proposed using lattices, etc.

This thesis is concerned with some of the underlying mathematical aspects of discrete-logarithm-based systems. The groups G with which they are concerned will be presented in the next chapter — for now, let us keep motivating their introduction.

I.2 Asymmetric Primitives

Although ciphers can be implemented efficiently, the need for a shared key to be secretly transmitted prior to any two-party communication is inconvenient. Most often today, a shared key is first established using asymmetric techniques (which overcome this problem) *over the insecure channel*, and then used to encrypt the data via a stream or block cipher.

PUBLIC-KEY PARADIGM

DIFFIE and HELLMAN (1976) introduced the key exchange below, which solves precisely this problem: making two individuals agree, over an open channel, on a shared secret key (to be subsequently used for encryption); it proceeds as follows:

1. Athos chooses an element g of some group G and sends it to Bonacieux.
2. Athos picks an integer a and sends g^a to Bonacieux.
3. Bonacieux picks an integer b and sends g^b to Athos.
4. Athos and Bonacieux compute the *shared secret* g^{ab} as $(g^a)^b$ and $(g^b)^a$ respectively.

When a passive observer breaks this scheme, they have solved the following.

Definition I.2.1. *The Diffie–Hellman problem is that of computing g^{ab} from g , g^a , and g^b .*

It is obviously no harder than the discrete logarithm problem, and is believed to neither be weaker. This key-exchange is hence considered secure in well-chosen groups of order 2^{256} .

The problem of authentication remains, since Milady de Winter could bribe the courier so as to intercept and forge messages: she would pick her own integer c and impersonate Bonacieux to Athos (with secret g^{ac}) and Athos to Bonacieux (with secret g^{bc}), thus spying on (and actively interfering with) the whole communication.

Definition 1.2.2. Asymmetric encryption schemes consist of two families E and D of functions, not necessarily everywhere defined, from \mathbb{S} to \mathbb{S} and a one-way function w such that $D_k \circ E_{w(k)} = \text{Id}_{\text{dom} E_{w(k)}}$ for all strings k . It is a signing scheme provided $E_{w(k)} \circ D_k = \text{Id}_{\text{dom} D_k}$ also holds.

The map w is the *key-generation function*: it takes a *private key* k as input and returns the corresponding *public key* $w(k)$, to be publicly distributed along with E , making anybody able to encrypt messages that only the holder of k can decrypt. Conversely, if the key holder of a signing scheme broadcasts $D_k(m)$ for some message m , everyone can evaluate $E_{w(k)}(D_k(m))$ and be assured that the *signature* $D_k(m)$ originates from the holder of k .

In practice, signing schemes are designed independently from encryption schemes; however, for our brief presentation, this naïve framework encompassing both will suffice.

Asymmetric schemes rarely deal with large amounts of data: for encryption, ciphers are used and only their keys are encrypted asymmetrically; for authentication, it suffices to sign a hash of the message. Without loss of generality, we will therefore now describe primitives dealing with subsets of \mathbb{S} whose coding as bits will be understood.

EARLY CONSTRUCTIONS

Definition 1.2.3. In a group G noted multiplicatively, the short product problem is that of finding a subsequence of a given sequence $S \in G^{(\mathbb{N})}$ whose product is a prescribed element z .

Products of subsequences of S are called short products; in addition, when S has no repeated elements, this problem is known as the subset sum problem in additive groups and as the knapsack problem for $G = \mathbb{Z}$.

Some of its instances are equivalent to discrete logarithm problems: if S' is a subsequence of $S = (g^{2^0}, g^{2^1}, \dots, g^{2^{\lfloor \log_2 \#G \rfloor}})$ with product z , then $z = g^n$ where the i^{th} bit of n is one if $g^{2^i} \in S'$ and zero otherwise. From a cryptographic standpoint, this means that the map

$$E_S : (x_i) \in \{0, 1\}^{\lfloor \log_2 \#G \rfloor} \mapsto \prod_{i=1}^{\lfloor \log_2 \#G \rfloor} s_i^{x_i} \in G$$

is a tentative one-way function for certain groups G and sequences S of length about $\log_2 \#G$.

MERKLE and HELLMAN (1978) proposed an asymmetric scheme which scrambles easy knapsacks (the private keys) into seemingly harder ones (the public keys): let $(s_i) \in \mathbb{N}^n$ be a sequence such that $\sum_{i < j} s_i < s_j$ for $j \in \{1, \dots, n\}$, put $v = \sum s_i$ and define S as the projection of (s_i) to \mathbb{Z}/v ; the map E_S can then be inverted in polynomial time by a greedy algorithm. Now, choose an integer u coprime to v , and publish the sequence $T = (t_i) = (us_i \bmod v)$. In the formalism above, we have $k = (S, u, v)$ as the private key, $w : k \mapsto T$ as the key-generation

map, and $E_{w(k)} : (m_i) \in \{0, 1\}^n \mapsto \sum m_i \cdot t_i$ as the encryption function; the greedy algorithm decrypts a ciphertext m' by finding a subsequence of S with sum $u^{-1}m' \bmod v$. SHAMIR (1982) later broke this scheme due to the simplicity of its scrambling process.

MERKLE (1979) constructed a much more conservative signature scheme, built entirely from a hash function h , and certified its security assuming that of h . This was achieved by developing an original idea of LAMPORT (1979): if one selects private strings x and y and publishes their images $h(x)$ and $h(y)$ by a hash function, he may later sign a bit of data by releasing either x (if the bit is zero) or y (if it is one).

MODERN CONSTRUCTIONS

The RSA cryptosystem of RIVEST, SHAMIR, and ADLEMAN (1978) rests on the problem of integer factoring, although subexponential factoring algorithms were already known at the time. Nevertheless, it has become widely used despite the large keys and *a fortiori* computing resources required by reasonable levels of security.

Let $n = pq$ be a product of two primes, and pick an integer r coprime to $(p-1)(q-1)$; this ensures that the map $m \mapsto m^r$ is an automorphism of $(\mathbb{Z}/n)^\times$. Let the private key be (p, q, r) , and publish (n, r) as the public key and $E_{(n,r)} : m \mapsto m^r \bmod n$ as the encryption function; decrypting then consists in applying the inverse automorphism $D : m \mapsto m^s$ where s can be computed from p and q (and conversely) since $s = r^{-1} \bmod (p-1)(q-1)$.

The key-length of an RSA cryptosystem is the bit-size of n . The following table shows, at various levels of security, the key-lengths recommended by ECRYPT II (2010) for RSA, ElGamal (see below), and equivalently secure symmetric schemes *in the best case*, that is, assuming well-chosen parameters. The superlinear growth of RSA keys is due to the aforementioned subexponential factoring techniques.

SYMMETRIC	RSA	ElGamal
80	1248	160
128	3248	256
256	15424	512

ELGAMAL (1985) designed a cryptosystem based on the Diffie–Hellman problem: let g be a generator of some group G , and pick an integer x . The public key is (g, b) where $b = g^x$, and x is the secret key. The ciphertext of a message m (encoded as an element of G) is $(g^y, m \cdot b^y)$ where y is a random integer; to decrypt it, simply put g^y to the power x and divide it out from $m \cdot b^y$.

Compared to many other cryptosystems, the ElGamal scheme stands out for its elegance and flexibility: since the group G it uses is not restricted to a certain class (such as RSA which

uses $G = (\mathbb{Z}/n)^\times$, it has more latitude to find one that has both an effective group law, and in which no attack is faster than generic ones.

ADVANCED PRIMITIVES

Beyond encrypting and signing, many advanced and/or exotic cryptographic schemes exist, most of which are enabled by the computability of certain mathematical objects.

Zero-knowledge proofs are protocols where Athos is to convince Bonacieux that he knows some secret without revealing anything about it. For instance, the secret could be a (dedicated) private key; to be convinced of his knowledge of the private key, Bonacieux could send Athos a random message encrypted with the associate public key and challenge him to reveal the plaintext — she would learn nothing regarding the private key but that Athos knows it. Many other constructions exist, notably that of GOLDREICH, MICALI, and WIGDERSON (1986) which demonstrated the power of a graph-based approach.

Homomorphic encryption aims at performing operations on plaintexts seamlessly via ciphertexts. For instance, in the ElGamal scheme, the term-by-term product of ciphertexts for m and m' is a valid ciphertext for mm' since

$$(g^y, mb^y) \cdot (g^{y'}, m'b^{y'}) = (g^{y+y'}, mm'b^{y+y'}).$$

Fully homomorphic systems feature two such algebraic operations; they are far more powerful as they enable the encrypted evaluation of any circuit. GENTRY (2009) described such a scheme using lattices but its practicality is still a topic of active research.

The past decade also saw a plethora of novel cryptographic schemes exploiting the richness of *pairings*, that is, non-degenerate bilinear maps $\Psi : G_1 \times G_2 \rightarrow H$ where the groups G_i are noted additively, and H is noted multiplicatively. The first was a one-round tripartite Diffie–Hellman key-exchange: assume Athos, Bonacieux, and Chevreuse are to derive a shared secret key over an insecure channel; the protocol of JOUX (2000) goes as follows:

1. Athos chooses and broadcasts a pairing Ψ and a pair $(x, y) \in G_1 \times G_2$.
2. Athos picks an integer a and broadcasts ax and ay .
3. Bonacieux picks an integer b and broadcasts bx and by .
4. Chevreuse picks an integer c and broadcasts cx and cy .
5. Everybody computes $\Psi(ax, by)^c = \Psi(bx, cy)^a = \Psi(cx, ay)^b$.

1.3 Generic Methods

The security of a cryptographic scheme based on a group does not depend on its isomorphism type alone, since an explicit isomorphism might be very costly to compute; it depends on how the group problem is *encoded* by the function E . For instance, discrete logarithm problems are much easier to solve in $\mathbb{Z}/(p-1)$ than in $(\mathbb{Z}/p)^\times$ although their underlying groups are isomorphic.

This section considers algorithms which apply to any group G regardless of its coding; later, we will come back to which specific codings make which problems easier.

GENERIC ALGORITHMS

The framework of generic algorithms abstracts group problems (such as the discrete logarithm problem) from specific codings which might render it “artificially” easier. Beware that our definition is not strictly-speaking the most classical one, as we assume that elements are uniquely identified and can be drawn uniformly at random.

Definition 1.3.1. *A coding of a group G is an injective map $\gamma : G \rightarrow \mathbb{S}$.*

A generic group is a black-box interface to a group G which can output $\gamma(z)$ for a random z and evaluate $(x, y) \mapsto \gamma(\gamma^{-1}(x) \cdot \gamma^{-1}(y))$ and $x \mapsto \gamma(1/\gamma^{-1}(x))$, where the coding γ is unknown.

A generic algorithm takes as input a sequence of encoded group elements $\gamma(x_i)$ and is allowed calls to the black box; its complexity is measured by the number of such calls.

Intuitively, a generic group is a group with shuffled elements, so that nothing is left to exploit in their representation: generic algorithms can only compute the group law.

We will see that many hard problems can be solved by generic algorithms in time $O(\sqrt{\#G})$ but not less. However, determining the order of an element (a special case of discrete logarithm) and, as a consequence, computing the group structure of abelian groups were recently proved by SUTHERLAND (2007) to require far fewer operations. Nevertheless, for the specific problems we are concerned with, namely the discrete logarithm problem and the short product problem, the generic algorithms described below are believed to be the best known to date.

REDUCTION TO PRIME GROUPS

The method of POHLIG and HELLMAN (1978) was originally directed at computing discrete logarithms in $(\mathbb{Z}/p)^\times$ but, more generally, it reduces many problems on abelian groups G into smaller prime groups. It combines two ingredients, the first of which is the following consequence of the Chinese remainder theorem.

Theorem 1.3.2. *Let G be an abelian group of order $n = \prod p^{\alpha_p}$ for some primes p and positive integers α_p . The map*

$$x \in G \mapsto \left(x^{n/p^{\alpha_p}}\right)_p \in \prod_{p|n} G[p^{\infty}]$$

is an isomorphism where the p -Sylow subgroup $G[p^{\infty}]$ denotes the subgroup of all elements whose order is a power of p . Its inverse is effectively given by the Chinese remainder theorem.

Once the order of G is factored, this reduces any instance of a problem compatible with the group law to several instances, one in each group $G[p^{\infty}]$ of prime-power order.

To get down to prime-order groups, the second ingredient is a lifting approach: assuming that G has order p^α , a subgroup series $G = G_0 \rightarrow G_1 \rightarrow \cdots \rightarrow G_\alpha = \{1\}$ where each arrow has index p is used to reduce problems into the quotient groups G_i/G_{i-1} . This technique applies to many problems, such as computing square roots modulo n as TONELLI (1891) showed, but its specifics depend on the particular problem considered.

For instance, suppose that $g \in G$ has order p^α , and write the discrete logarithm of a certain $h = g^x$ as $x = \sum_{i=0}^{\alpha-1} x_i p^i$ for some $x_i \in \{0, \dots, p-1\}$; the integers x_i can be recursively computed by

$$x_i = \log_{g^{(p^{\alpha-1})}} \left(g^{-\sum_{j=0}^{i-1} x_j p^j} h^{(p^{\alpha-1-i})} \right)$$

which amounts to projecting discrete logarithms from G_i/G_{i-1} to $G_{\alpha-1}$.

Here, we have assumed that the group order was known; in many cryptographic settings, this is actually the case. Although generic algorithms require exponential time to compute the group structure, we believe that it is questionable to base the security of a scheme on hiding the structure of a group (as RSA does), and that almost exclusively groups of prime (or near-prime) orders should be used in cryptography.

BABY-STEP GIANT-STEP

SHANKS (1971) developed the *baby-step giant-step method* for computing discrete logarithms, although it applies to a broad range of problems. Our presentation here uses the formalism of B. and SUTHERLAND (2011), the generality of which we will later exploit.

The general idea is to design sets A and B so that *collisions*, that is, common elements to A and B , yield solutions to the problem. Specifically, we construct A and B as the respective images of two maps ϕ and ψ with values in G and seek collisions of the form $\phi(x) = \psi(y)$.

For instance, to compute the logarithm of h in base g , put $\phi : i \mapsto g^i$ and $\psi : j \mapsto hg^{-Nj}$ for $i, j \in \{0, \dots, N\}$ where $N = \lceil \sqrt{\#G} \rceil$; collisions of the form $\phi(i) = \psi(j)$ yield $\log_g h = i + Nj$, and there must exist at least one such collision due to the existence of the discrete logarithm.

To quickly search for elements of $A \cap B$, a data structure allowing fast lookups is required; fast insertions are also a must. We therefore typically use hash tables or red black trees. The cost of computing $A \cap B$ is then $(\#A + \#B)O(\log n)$ for $n = \#G$, where the last term denotes the complexity of the searching and inserting.

When A and B are not as explicit as above, it might not be possible to prove the existence of a collision. The algorithm can then be randomized to rely on the *birthday paradox*:

Proposition 1.3.3. *Let A and B be uniformly distributed subsets of cardinality $a\sqrt{n}$ and $b\sqrt{n}$ in a set G of cardinality n ; then*

$$\text{Prob}[A \cap B = \emptyset] \xrightarrow{n \rightarrow \infty} e^{-ab}.$$

Assuming ϕ and ψ are random, \sqrt{n} images of each thus suffice to have a $1 - 1/e$ chance of finding a collision. In the unlucky event there is none, we can repeat this process m times, adding more images to our red-black tree; this increases the likelihood of success to $1 - 1/e^{m^2}$.

From now on, we say that a *probabilistic algorithm* has complexity X , or that an algorithm has *probabilistic complexity* X , to mean that it always returns the correct answer (this is known as a *Las Vegas algorithm*) and that, with probability at least $1/2$, its runtime is bounded by X . By the discussion above, up to a constant, it is equivalent to the notion of average complexity.

POLLARD'S RHO

The baby-step giant-step method requires storing $O(\sqrt{n})$ elements; an algorithm emulating its behavior with minimal space storage was developed by POLLARD (1975) for integer factoring, and later applied to discrete logarithms by POLLARD (1978).

Let us first unify things in a map $\pi : \mathcal{C} \rightarrow G$ equal to ϕ and ψ on their respective domains, where \mathcal{C} denotes their disjoint union. The rho method involves a *pseudorandom function* $\rho : \mathcal{C} \rightarrow \mathcal{C}$, that is, an effective map for which the distribution of $\rho^{(i)}(w)$ (the composition of i copies of ρ) is seemingly uniform as $w \in \mathcal{C}$ is fixed and the integer i varies. It is required to preserve collisions, that is, $\pi(x) = \pi(y) \Rightarrow \pi(\rho(x)) = \pi(\rho(y))$.

The map ρ is thought of as generating A and B under π , and the crucial step is to find collisions $\pi\rho^{(i)}(w) = \pi\rho^{(j)}(w)$ without storing many values; when $\rho^{(i)}(w) \neq \rho^{(j)}(w)$ collide through π , we expect that one is an image of ϕ and the other is one of ψ , which gives a *proper collision* — when their sizes are equal, this happens with probability a half.

Avoiding storage requires a *cycle-detection* method on the graph of iterates of ρ evaluated at w . The simplest such method is due to FLOYD who observed that, whenever $\rho^{(i)}(w)$ and $\rho^{(j)}(w)$ collide for some integers i and j satisfying $i > 2j$, then $\rho^{(2(i-j))}(w)$ and $\rho^{(i-j)}(w)$ also

collide. Thus, it suffices to compute $\rho^{(2^i)}(w)$ alongside $\rho^{(i)}(w)$ for increasing i 's and wait for them to collide; then, ρ maps are unstacked until the original collision is found. Better cycle-detection methods improve the runtime by a constant factor using more memory.

The difficulty lies in designing a function ρ suited to a given problem; more details will be given on that later, especially for the short product problem. To factor an integer n , POLLARD (1975) put $\mathcal{C} = \mathbb{Z}/n$ and chose ρ to be a polynomial function; the map π can then be the projection to any subgroup of \mathbb{Z}/n which need not be known: by computing $\gcd(\rho^{(i)}(w) - \rho^{(j)}(w), n)$, we can detect when a collision occurs and hopefully find a factor of n . This method is nowadays mostly used for small integers n , as asymptotically faster factoring algorithms have since been developed.

A current international effort (2009) aims at solving a discrete logarithm problem challenge in a group of 129-bit order (this group is an elliptic curve where generic algorithms are the best available); when completed, it will likely be the record rho algorithm run.

1.4 Cryptographic Groups

Let us now review the cryptographic security of various groups, mostly focusing on the discrete logarithm problem.

FINDING PRIMES

We advocated for prime-order groups; now let us mention how prime numbers can be found. The best method for this is simply to draw numbers at random until a prime is found; for numbers of n bits, this requires an expected $O(n)$ operations by the theorem below.

Assuming the generalized Riemann hypothesis, MILLER (1975) first derived a fast (polynomial time) deterministic primality test, later turned into an unconditional but probabilistic method by RABIN (1980). Although AGRAWAL, KAYAL, and SAXENA (2004) have since proved that deterministic primality proving need not rely on unproven assumptions, the dependency on the generalized Riemann hypothesis is interesting: this conjecture predicts the behavior of primes in various fields. First recall the celebrated prime number theorem of HADAMARD (1896) and DE LA VALLÉE-POUSSIN (1896).

Theorem 1.4.1. *The number of prime integers less than x is asymptotically equivalent to*

$$\int_2^x \frac{dt}{\log t} \sim \frac{x}{\log x}.$$

Proofs of this theorem involve establishing certain properties of analytic functions related to integers; more generally, if K is any number field, define, for $s \in \mathbb{C}$ with $\Re(s) > 1$,

$$\zeta_K(s) = \sum_{\mathfrak{a} \in \mathfrak{I}} N(\mathfrak{a})^{-s}$$

where \mathfrak{I} is the set of ideals of the ring of integers of K , and extend ζ_K to \mathbb{C} by analytic continuation. This function encodes the behavior of prime ideals of K ; to obtain precise results on their distribution, one often assumes the *extended Riemann hypothesis* which states that all zeroes s of ζ_K in the strip $0 < \Re(s) < 1$ lie on the line $\Re(s) = 1/2$. The extended Riemann hypothesis follows from the stronger *generalized Riemann hypothesis*, and we often assume the latter when only the former is needed.

MILLER (1975) actually exploited the following result of ANKENY (1952), where the label “(GRH)” denotes that the statement holds under the generalized Riemann hypothesis.

Theorem 1.4.2 (GRH). *Let p and q be integers such that q divides $p - 1$. The least integer x which cannot be written as $y^q \pmod p$ for some $y \in \mathbb{N}$ is asymptotically $O(\log^2 p)$.*

We conclude with a conjecture of BATEMAN and HORN (1965) generalizing the prime number theorem; it is useful for generating elliptic curves as we will see later. Essentially, it asserts that distinct irreducible polynomials take prime values almost independently, and that this “almost” is quantified by their values modulo primes p .

Conjecture 1.4.3. *Let F be a set of distinct irreducible non-constant polynomials of $\mathbb{Z}[X]$. The number of integers less than x at which all its polynomials simultaneously take prime values is asymptotically equivalent to*

$$\frac{C}{\prod_{f \in F} \deg f} \int_2^x \frac{dt}{(\log t)^{\#F}}$$

where $C = \prod_p \left(1 - \frac{1}{p} \# \left\{ z \in \mathbb{F}_p : \prod_{f \in F} f(z) = 0 \right\} \right) \left/ \left(1 - \frac{1}{p} \right)^{\#F} \right.$.

INDEX CALCULUS

Since the baby-step giant-step or rho method use $O(\sqrt{p})$ operations to find a factor p of an integer n , factors of n can always be found in $O(n^{1/4})$ time. By iterating this search for factors and testing the primality of the factors obtained, an integer n can be factored in probabilistic time $O(n^{1/4})$. When the RSA cryptosystem was proposed, much faster algorithms already existed and they were substantively improved subsequently.

The simplest such method is due to KRAITCHIK (1926). To split an integer n , it crafts a nontrivial relation $x^2 = y^2 \pmod n$ by combining many easier relations so as to eliminate non-square factors; the easier relations are of the form $z^2 \pmod n = \prod p^{a_p}$ for primes p less than some bound $L(n)$. To bound the probability that such a factorization exists, we rely on this result of CANFIELD, ERDŐS, and POMERANCE (1983).

Theorem 1.4.4. *For any $c > 0$, the probability for a random number of $\{1, \dots, x\}$ to have no prime factor larger than $L(x)^c$ is equivalent to $L(x)^{-1/2c+o(1)}$ as $x \rightarrow \infty$, where we used the function*

$$L_a(x) = \exp\left((\log x)^a (\log \log x)^{1-a}\right)$$

with the convention that omitting the parameter $a \in (0, 1)$ means $a = 1/2$.

Assuming Gaussian elimination takes cubic time in the number of variables, we set $c = 1/2$ and obtain a nontrivial splitting of n in time $L(n)^{3/2+o(1)}$.

The broad family of *combining congruences algorithms* encompasses methods using *factor bases* (as the primes up to $L(n)$); they apply to many integer-based problems such as discrete logarithms in finite fields and integer factoring. Under unproven assumptions, the asymptotically fastest such method is the *number field sieve* of COPPERSMITH (1993), which builds up on the work of many including LENSTRA and LENSTRA (1993), with heuristic complexity

$$L_{1/3}^{c_{\text{NFS}}}(n) \quad \text{where} \quad c_{\text{NFS}} = 2\sqrt[3]{\frac{46 + 13\sqrt{13}}{108}} \approx 1.902$$

Recently, KLEINJUNG *et alii* (2010) used a similar method to factor a 768-bit RSA modulus, thereby deprecating smaller RSA keys; the effectiveness of this attack is blatant when compared to elliptic curves whose discrete logarithms can only be attacked up to 130 bits.

Unconditionally proven factoring algorithms are slightly slower, with the state-of-the-art method of LENSTRA and POMERANCE (1992) using an expected $L(n)^{1+o(1)}$ operations; it exploits a similar factor base paradigm in certain class groups. Since these objects are built from ideals it is not surprising that subexponential methods should apply to them as well, and we will elaborate on that later as class groups become a building block of our own algorithms.

ABELIAN VARIETIES

Cryptosystems based on the discrete logarithm problem in finite fields have been proposed as alternatives to RSA; however, up to certain modifications, modern integer factoring algorithms also apply to this problem, so it provides no additional security.

Shortly after LENSTRA (1987) introduced a novel factoring algorithm based on elliptic curves, MILLER (1986) and KOBLITZ (1987) suggested their use in cryptography; subsequently, KOBLITZ (1989) further proposed using the broader class of abelian varieties. This has motivated tremendous developments in computational number theory, and has enabled a wide spectrum of possibilities in cryptography.

These applications are motivated by two facts: first, that the group law of abelian varieties can be computed efficiently, and second, that no algorithm better than generic ones is currently known to attack the discrete logarithm problem on most abelian varieties of dimension one and two. Before formally defining abelian varieties, we briefly give loose statements highlighting their applicability to cryptography.

Abelian varieties are objects endowed with two compatible structures:

- a *geometric* structure: it is the zero locus of multivariate polynomials over a field k ;
- a *group* structure: it admits a group law given by rational functions.

When the defining polynomials have certain forms, the group law can be evaluated efficiently using short rational functions. This can be done for all varieties of dimension one and two (the *dimension* is roughly the number of variables minus the number of polynomials).

Cryptography uses finite fields k and such forms, allowing fast arithmetic; for instance, BERNSTEIN and LANGE (2007) suggested defining G as the set of points $(x, y) \in k^2$ verifying

$$x^2 + y^2 = 1 + dx^2y^2$$

for some non-square parameter $d \in k$, endowed with the addition law defined by

$$(x, y) + (x', y') = \left(\frac{xy' + x'y}{1 + dx'x'yy'}, \frac{yy' - xx'}{1 - dx'x'yy'} \right).$$

Since the number of points of an abelian variety of dimension g defined over k (that is, the order of the underlying group) is roughly $(\#k)^g$ and otherwise behaves quite randomly, a prime-order one can be sought by drawing varieties at random while their orders are composite. Alternatively, we will later discuss the theory of complex multiplication which provides means to generate abelian varieties with a prescribed order.

SPECIAL ATTACKS

We stated that attacks on the discrete logarithm problem of most elliptic curves are not known to be faster than generic ones. To conclude this chapter, we give an exhaustive list of classes of abelian varieties for which this does not hold, so remaining ones can *a priori* be considered secure. Details on these attacks can be found in AVANZI, COHEN, DOCHE, FREY, LANGE, NGUYEN, and VERCAUTEREN (2005).

Index-calculus with subspace as factor base. Gröbner basis algorithms can decompose points of abelian varieties into sums of points in certain subspaces (such as having certain coordinates equal to zero, or defined over some strict subfield); this enables index-calculus attacks effective on varieties of dimension $g > 2$ or defined over non-prime base fields.

Reduction to finite fields via pairings. The Weil pairing maps pairs of points of order ℓ from an abelian variety to the multiplicative group of an extension of degree $e(\ell)$ of the base field k . It transports the discrete logarithm problem, so the value of $e(\ell)$ must be large enough to prevent attacks in the extension field from being feasible.

Lift to characteristic zero. Certain abelian varieties with special properties (such as the infamous *anomalous curves*, whose cardinality is that of their base field) can be lifted to p -adic fields, from where discrete logarithm problems can be transferred to \mathbb{Z}/p .

Isogenies. Isogenies are morphisms between abelian varieties; they can transport the discrete logarithm from a variety \mathcal{A} to about ℓ^g other varieties in time $\ell^{O(g^2)}$ for most primes ℓ ; if any of those varieties have one of the above weaknesses, then so does \mathcal{A} .

Since no attack faster than generic algorithms is known to affect randomly chosen, prime-order abelian varieties of dimension one or two defined over finite fields with p or 2^p elements where p is a prime, we conclude that these are currently the best choice for public-key cryptography in a cryptosystem of ElGamal type.

References

1891. Alberto TONELLI.
 “Bemerkung über die Auflösung quadratischer Congruenzen”.
 In: *Nachrichten von der Königlichen Gesellschaft der Wissenschaften und der Georg-Augusts-Universität zu Göttingen* 1891.10. Pages 344–346.
1896. Charles-Jean DE LA VALLÉE-POUSSIN.
 “Recherches analytiques sur la théorie des nombres premiers”.
 In: *Annales de la Société Scientifique de Bruxelles* 20.2. Pages 183–256.
1896. Jacques HADAMARD.
 “Sur la distribution des zéros de la fonction $\zeta(s)$ et ses conséquences arithmétiques”.
 In: *Bulletin de la Société Mathématique de France* 24. Pages 199–220.

1926. Maurice KRAITCHIK.
Théorie des Nombres. Volume 2.
Analyse indéterminée du second degré et factorisation. Gauthier-Villars.
1937. Alan TURING.
“On computable numbers, with an application to the Entscheidungsproblem”.
In: *Proceedings of the London Mathematical Society* 42.2. Pages 230–265.
DOI: 10.1112/plms/s2-42.1.230.
1949. Claude SHANNON.
“Communication theory of secrecy systems”.
In: *Bell System Technical Journal* 28.4. Pages 656–715.
1952. Nesmith C. ANKENY.
“The least quadratic non residue”. In: *Annals of Mathematics* 55.1. Pages 65–72.
DOI: 10.2307/1969420.
1965. Paul T. BATEMAN and Roger A. HORN.
“Primes represented by irreducible polynomials in one variable”.
In: *Theory of Numbers*. Edited by Albert L. WHITEMAN. Volume 8.
Proceedings of Symposia in Pure Mathematics. American Mathematical Society.
Pages 119–132.
1965. Gordon E. MOORE.
“Cramming more components onto integrated circuits”.
In: *Electronics Magazine* 38.8. Pages 114–117.
1971. Daniel SHANKS.
“Class number, a theory of factorization, and genera”.
In: *1969 Number Theory Institute*. Edited by Donald J. LEWIS. Volume 20.
Proceedings of Symposia in Pure Mathematics. American Mathematical Society.
Pages 415–440.
1975. Gary L. MILLER.
“Riemann’s hypothesis and tests for primality”.
In: *Symposium on Theory of Computing — STOC ’75*.
Edited by William C. ROUNDS, Nancy MARTIN, Jack W. CARLYLE, and
Michael A. HARRISON. Association for Computing Machinery. Pages 234–239.
DOI: 10.1145/800116.803773.
1975. John M. POLLARD.
“A Monte Carlo method for factorization”.

- In: *BIT Numerical Mathematics* 15.3. Pages 331–334.
DOI: 10.1007/BF01933667.
1976. Whitfield DIFFIE and Martin E. HELLMAN.
“New directions in cryptography”.
In: *IEEE Transactions on Information Theory* 22.6. Pages 644–654.
DOI: 10.1109/TIT.1976.1055638.
1978. Robert J. McELIECE.
“A public-key cryptosystem based on algebraic coding theory”.
In: *DSN Progress Report*. Volume 42–44. Jet Propulsion Laboratory. Pages 114–116.
1978. Ralph C. MERKLE and Martin E. HELLMAN.
“Hiding information and signatures in trapdoor knapsacks”.
In: *IEEE Transactions on Information Theory* 24.5. Pages 525–530.
DOI: 10.1109/TIT.1978.1055927.
1978. Stephen C. POHLIG and Martin E. HELLMAN.
“An improved algorithm for computing logarithms over $GF(p)$ and its cryptographic significance”. In: *IEEE Transactions on Information Theory* 24.1. Pages 106–110.
DOI: 10.1109/TIT.1978.1055817.
1978. John M. POLLARD.
“Monte Carlo methods for index computation (mod p)”.
In: *Mathematics of Computation* 32.143. Pages 918–924.
DOI: 10.2307/2006496.
1978. Ron L. RIVEST, Adi SHAMIR, and Leonard ADLEMAN.
“A method for obtaining digital signatures and public-key cryptosystems”.
In: *Communications of the ACM* 21.2. Pages 120–126.
DOI: 10.1145/359340.359342.
1979. Leslie LAMPART.
Constructing digital signatures from a one-way function.
Technical Report № 98, Computer Science Laboratory, SRI International.
1979. Ralph C. MERKLE.
“Secrecy, authentication and public key systems”. PhD thesis. Stanford University.
URL: <http://www.merkle.com/papers/Thesis1979.pdf>.
1980. Michael O. RABIN.
“Probabilistic algorithm for testing primality”.
In: *Journal of Number Theory* 12.1. Pages 128–138.
DOI: 10.1016/0022-314X(80)90084-0.

1982. Adi SHAMIR.
“A polynomial time algorithm for breaking the basic Merkle-Hellman cryptosystem”.
In: *Foundations of Computer Science — FOCS '82*. IEEE Computer Society.
Pages 145–152. DOI: 10.1109/SFCS.1982.55.
1983. Earl CANFIELD, Paul ERDŐS, and Carl POMERANCE.
“On a problem of Oppenheim concerning ‘factorisatio numerorum’”.
In: *Journal of Number Theory* 17.1. Pages 1–28.
DOI: 10.1016/0022-314X(83)90002-1.
1985. Taher ELGAMAL.
“A public key cryptosystem and a signature scheme based on discrete logarithms”.
In: *Advances in Cryptology — CRYPTO '84*.
Edited by George Robert BLAKLEY and David CHAUM. Volume 196.
Lecture Notes in Computer Science. Springer. Pages 10–18.
DOI: 10.1007/3-540-39568-7_2.
1986. Oded GOLDREICH, Silvio MICALI, and Avi WIGDERSON.
“Proofs that yield nothing but their validity and a methodology of cryptographic protocol design”. In: *Foundations of Computer Science — FOCS '86*.
IEEE Computer Society. Pages 174–187. DOI: 10.1109/SFCS.1986.47.
1986. Victor S. MILLER.
“Use of elliptic curves in cryptography”. In: *Advances in Cryptology — CRYPTO '85*.
Edited by Hugh C. WILLIAMS. Volume 218. Lecture Notes in Computer Science.
Springer. Pages 417–426. DOI: 10.1007/3-540-39799-X_31.
1987. Neal KOBLITZ.
“Elliptic curve cryptosystems”.
In: *Mathematics of Computation* 48.177. Pages 203–209.
DOI: 10.1090/S0025-5718-1987-0866109-5.
1987. Hendrik W. LENSTRA.
“Factoring integers with elliptic curves”.
In: *Annals of Mathematics* 126.3. Pages 649–673. DOI: 10.2307/1971363.
1988. Tsutomu MATSUMOTO and Hideki IMAI.
“Public quadratic polynomial-tuples for efficient signature-verification and message-encryption”. In: *Advances in Cryptology — EUROCRYPT '88*.
Edited by Christof G. GÜNTHER. Volume 330. Lecture Notes in Computer Science.
Springer. Pages 419–453. DOI: 10.1007/3-540-45961-8_39.

1989. Neal KOBLITZ.
“Hyperelliptic cryptosystems”. In: *Journal of Cryptology* 1.3. Pages 139–150.
DOI: 10.1007/BF02252872.
1992. Hendrik W. LENSTRA and Carl POMERANCE.
“A rigorous time bound for factoring integers”.
In: *Journal of the American Mathematical Society* 5.3. Pages 483–516.
DOI: 10.1090/S0894-0347-1992-1137100-0.
1993. Don COPPERSMITH.
“Modifications to the number field sieve”.
In: *Journal of Cryptology* 6.3. Pages 169–180. DOI: 10.1007/BF00198464.
1993. Arjen K. LENSTRA and Hendrik W. LENSTRA (editors).
The Development of the Number Field Sieve. Volume 1554.
Lecture Notes in Mathematics. Springer. ISBN: 3-540-57013-4.
1994. Gilles ZÉMOR.
“Hash functions and Cayley graphs”.
In: *Designs, Codes and Cryptography* 4.3. Pages 381–394.
DOI: 10.1007/BF01388652.
1996. Miklós AJTAI.
“Generating hard instances of lattice problems”.
In: *Symposium on Theory of Computing — STOC '96*. Edited by Gary L. MILLER.
Association for Computing Machinery. Pages 99–108.
DOI: 10.1145/237814.237838.
1997. Victor SHOUP.
“Lower bounds for discrete logarithms and related problems”.
In: *Advances in Cryptology — EUROCRYPT '97*. Edited by Walter FUMY.
Volume 1233. Lecture Notes in Computer Science. Springer. Pages 256–266.
DOI: 10.1007/3-540-69053-0_18.
1999. Joan DAEMEN and Vincent RIJMEN.
The Rijndael block cipher. Advanced Encryption Standard proposal submitted to the
National Institute of Standards and Technology.
2000. Antoine JOUX.
“A one round protocol for tripartite Diffie–Hellman”.
In: *Algorithmic Number Theory — ANTS-IV*. Edited by Wieb BOSMA.
Volume 1838. Lecture Notes in Computer Science. Springer. Pages 385–393.
DOI: 10.1007/10722028_23.

2004. Manindra AGRAWAL, Neeraj KAYAL, and Nitin SAXENA.
“PRIMES is in P”. In: *Annals of Mathematics* 160.2. Pages 781–793.
DOI: 10.4007/annals.2004.160.781.
2005. Roberto M. AVANZI, Henri COHEN, Christophe DOCHE, Gerhard FREY,
Tanja LANGE, Kim NGUYEN, and Frederik VERCAUTEREN.
Handbook of Elliptic and Hyperelliptic Curve Cryptography.
Discrete Mathematics and its Applications. Chapman & Hall.
ISBN: 1-58488-518-1.
2007. Daniel J. BERNSTEIN and Tanja LANGE.
“Faster addition and doubling on elliptic curves”.
In: *Advances in Cryptology — ASIACRYPT ’07*. Edited by Kaoru KUROSAWA.
Volume 4833. Lecture Notes in Computer Science. Springer. Pages 29–50.
DOI: 10.1007/978-3-540-76900-2_3.
2007. Andrew V. SUTHERLAND.
“Order computations in generic groups”.
PhD thesis. Massachusetts Institute of Technology. URL:
<http://groups.csail.mit.edu/cis/theses/sutherland-phd.pdf>.
2009. Daniel V. BAILEY, Lejla BATINA, Daniel J. BERNSTEIN, Peter BIRKNER,
Joppe W. BOS, Hsieh-Chung CHEN, Chen-Mou CHENG, Gauthier van DAMME,
Giacomo de MEULENAER, Luis J. DOMINGUEZ PEREZ, Junfeng FAN,
Tim GÜNEYSU, Frank GURKAYNAK, Thorsten KLEINJUNG, Tanja LANGE,
Nele MENTENS, Ruben NIEDERHAGEN, Christof PAAR, Francesco REGAZZONI,
Peter SCHWABE, Leif UHSADEL, Anthony van HERREWEGE, and Bo-Yin YANG.
Breaking ECC2K-130. IACR ePrint: 2009/541.
2009. Denis X. CHARLES, Kristin E. LAUTER, and Eyal Z. GOREN.
“Cryptographic hash functions from expander graphs”.
In: *Journal of Cryptology* 22.1. Pages 93–113.
DOI: 10.1007/s00145-007-9002-x.
2009. Craig GENTRY.
“Fully homomorphic encryption using ideal lattices”.
In: *Symposium on Theory of Computing — STOC ’09*.
Edited by Michael MITZENMACHER. Association for Computing Machinery.
Pages 169–178. DOI: 10.1145/1536414.1536440.
2010. European Network of Excellence in Cryptology II.
Yearly report on algorithms and key sizes. Edited by Nigel P. SMART.
URL: <http://www.ecrypt.eu.org/documents/D.SPA.13.pdf>.

2010. Thorsten KLEINJUNG, Kazumaro AOKI, Jens FRANKE, Arjen K. LENSTRA, Emmanuel THOMÉ, Joppe W. BOS, Pierrick GAUDRY, Alexander KRUPPA, Peter L. MONTGOMERY, Dag A. OSVIK, Herman te RIELE, Andrey TIMOFEEV, and Paul ZIMMERMANN.
“Factorization of a 768-bit RSA modulus”.
In: *Advances in Cryptology — CRYPTO ’10*. Edited by Tal RABIN. Volume 6223. Lecture Notes in Computer Science. Springer. Pages 333–350.
DOI: 10.1007/978-3-642-14623-7_18.
2011. Gaetan BISSON and Andrew V. SUTHERLAND.
“A low-memory algorithm for finding short product representations in finite groups”.
In: *Designs, Codes and Cryptography*. To appear.
DOI: 10.1007/s10623-011-9527-8.

Abelian Varieties

Having established the important role of abelian varieties in modern cryptography, we turn to formally defining their properties from a mathematical standpoint.

We will present this theory concisely, in a conceptually elementary way which we believe highlights its effectiveness. For details, we refer to AVANZI, COHEN, DOCHE, FREY, LANGE, NGUYEN, and VERCAUTEREN (2005), SHAFAREVICH (1974), CORNELL and SILVERMAN (1986), SHIMURA (1998), MILNE (1998), and MUMFORD (1970), in increasing levels of abstraction.

II.1 General Theory

ALGEBRAIC VARIETIES

Fix a perfect field k , referred to as the *base field*, and a sufficiently large integer $n = \text{DIMN_MAX}$.¹ For any ideal \mathfrak{J} of the ring $k[x] = k[x_1, \dots, x_n]$ of polynomials in n variables with coefficients in k , define the *affine variety* $\mathcal{V}_{\mathfrak{J}}$ as consisting, over any extension field K/k , of the set $\mathcal{V}_{\mathfrak{J}}(K)$ of common zeroes of \mathfrak{J} in K^n called *points* of the variety. HILBERT (1893) proved the famous Nullstellensatz:

Theorem II.1.1. *When k is algebraically closed, the largest ideal of $k[x]$ vanishing on $\mathcal{V}_{\mathfrak{J}}(k)$ is the radical ideal $\sqrt{\mathfrak{J}}$ formed by polynomials of which a power lies in \mathfrak{J} .*

This puts in bijection radical ideals with affine varieties over algebraically closed fields; computationally, one might therefore use generating sets of $\sqrt{\mathfrak{J}}$ to represent $\mathcal{V}_{\mathfrak{J}}$.

¹We find it amusingly convenient to fix an integer DIMN_MAX large enough so that all varieties we consider are embedded in the projective space with that large a dimension.

Such varieties are endowed with the *Zariski topology* whose closed sets are subvarieties. Via the Nullstellensatz, the topological notion of irreducibility corresponds to its algebraic counterpart. To avoid unnecessary technical contortions, we shall exclusively consider *absolutely irreducible varieties*, that is, varieties irreducible over an algebraic closure.

Affine varieties lie in the *affine space* $\mathbb{A}(\mathbb{K}) = \mathcal{V}_0(\mathbb{K})$, also written as $\mathbb{A}^n(\mathbb{K})$ when dimension n needs to be made explicit. In many contexts, it instead proves advantageous to:

- work with projective varieties;
- use Galois action to define objects over extension fields.

Over an algebraically closed field $\bar{\mathbb{K}}$, define the *projective space* $\mathbb{P}(\bar{\mathbb{K}})$ (of dimension $n - 1$) as the set of lines passing through the origin of $\mathbb{A}(\bar{\mathbb{K}})$, and over any field \mathbb{K} as the fixed subset

$$\mathbb{P}(\mathbb{K}) = \mathbb{P}(\bar{\mathbb{K}})^{\text{Gal}(\bar{\mathbb{K}}/\mathbb{K})}$$

under its absolute Galois group. Pragmatically, the projective space $\mathbb{P}(\mathbb{K})$ can be seen as formed by equivalence classes of collinear (non-zero) vectors, which gives the projection

$$x \in \mathbb{A}(\mathbb{K}) \setminus \{0\} \mapsto \{\lambda x : \lambda \in \bar{\mathbb{K}}^\times\} \in \mathbb{P}(\mathbb{K})$$

Working *in affine coordinates* means representing projective points by distinguished elements of \mathbb{A} (typically, by enforcing $x_0 = 1$; this covers almost all of \mathbb{P} but requires inversions to compute the distinguished element); on the other hand, working *in projective coordinates* means representing projective points as non-unique n -tuples.

Similarly, *projective varieties* are projections of affine varieties invariant under coordinate-wise scalar multiplication: if \mathfrak{J} is a *homogeneous ideal* of $k[x]$, that is, generated by sums of monomials of the same degree, the projective variety $\mathcal{V}_{\mathfrak{J}} \subset \mathbb{P}$ consists of equivalence classes (under scalar multiplication) of the affine variety $\mathcal{V}_{\mathfrak{J}} \subset \mathbb{A}$ endowed with the (quotient) Zariski topology.

From now on, we will exclusively consider absolutely irreducible open subsets of projective varieties, and refer to them simply as *varieties* (they are known to part of the literature as *quasiprojective varieties*); we will always implicitly assume that they are defined over algebraically closed fields, but say that they are *defined* over smaller fields when invariant under their absolute Galois group.

MORPHISMS

Consistent with the topology, *morphisms* are algebraic maps. For the affine space, they form the ring $\text{Hom}(\mathbb{A}, \mathbb{A})$ of n -tuples of n -variate polynomials. If \mathcal{V} and \mathcal{W} are two affine varieties, $\text{Hom}(\mathcal{V}, \mathcal{W})$ consists of those morphisms of $\text{Hom}(\mathbb{A}, \mathbb{A})$ mapping \mathcal{V} to \mathcal{W} .

Morphisms of projective varieties can be seen either conceptually, *looking down from* \mathbb{A} , as equivalence classes of tuples P of polynomials of $k[x]$ of homogeneous polynomials with the same degree for the relation $P \sim P' \Leftrightarrow \{P_i P'_j - P_j P'_i\} \subset \mathcal{I}$, or visually, *looking up from specific hyperplanes of* \mathbb{A} , as compatible collections of affine morphisms.

Two cases are of particular interest:

- the *coordinate ring* $\text{Hom}(\mathcal{V}_{\mathcal{I}}, K) \simeq K[x]/\mathcal{I}$, with addition and scalar multiplication.
- the *endomorphism monoid* $\text{Hom}(\mathcal{V}, \mathcal{V}) = \text{End}(\mathcal{V})$, endowed with composition; later, when we give \mathcal{V} a group law, it will become a ring.

Rational maps are defined similarly to above from tuples of rational functions. Most important are rational maps from a variety \mathcal{V} to a field of definition K , which form its *function field*, denoted $K(\mathcal{V})$. For projective varieties $\mathcal{V} = \mathcal{V}_{\mathcal{I}}$, it can be explicitly defined as the set of fractions P/Q of homogeneous polynomials in $K[x]$ of the same degree, with $Q \notin \mathcal{I}$, up to the relation $P/Q \sim P'/Q' \Leftrightarrow PQ' - P'Q \in \mathcal{I}$.

Various properties can be read off directly from function fields, such as:

Proposition II.1.2. *The Krull dimension of an ideal is equal to the transcendence degree of the function field associated to its variety; it is called the dimension of the variety.*

Algebraic extensions have finer indicators: a morphism $\phi \in \text{Hom}(\mathcal{V}, \mathcal{W})$ induces (by composition on the right) an embedding $\phi^* : K(\mathcal{W}) \rightarrow K(\mathcal{V})$; the *degree* of ϕ is the dimension $[K(\mathcal{V}) : \phi^* K(\mathcal{W})]$ which is finite when $\phi(\mathcal{V})$ has the same dimension as \mathcal{W} .

ALGEBRAIC GROUPS

Combining algebraic varieties with group structures yields *algebraic groups*:

Definition II.1.3. *An algebraic group is an (absolutely irreducible) non-empty algebraic variety endowed with a group law (noted additively) for which the map $(x, y) \mapsto x - y$ is a morphism.*

By non-empty, we mean that it must admit one rational point over its base field, so that it contains the neutral element for the group law. An important property of algebraic groups is given by the following algebraic equivalent to the analytic notion of differentiability.

Definition II.1.4. *An irreducible algebraic variety \mathcal{V} is nonsingular if the quotient of $\{f \in \bar{k}[\mathcal{V}] : f(P) = 0\}$ by its square has the same dimension (namely $g = \dim \mathcal{V}$) for all $P \in \mathcal{V}(\bar{k})$.*

Algebraic groups are nonsingular varieties; indeed, translation maps $\tau_P : Q \mapsto P + Q$ induce isomorphisms of tangent spaces, whose dimensions are that of the quotients above.

One simply defines *morphisms* of algebraic groups as morphisms of algebraic varieties preserving the group law, and *subgroups* of algebraic groups as subgroups that are closed. From now on, we shall work with categories as a whole: when we consider algebraic groups, morphisms and subgroups will be implicitly understood to be of algebraic groups (not just of algebraic varieties).

The proposition below argues that this behaves as expected.

Proposition II.1.5. *Let \mathcal{H} be an (algebraic) normal subgroup of an algebraic group \mathcal{G} . The quotient \mathcal{G}/\mathcal{H} has a unique structure of algebraic group such that:*

- the projection map $\mathcal{G} \rightarrow \mathcal{G}/\mathcal{H}$ is a morphism;
- all morphisms from \mathcal{G} with kernel containing \mathcal{H} factor through \mathcal{G}/\mathcal{H} .

For instance, the group $\mathrm{GL}_n(\mathbb{K})$ of invertible n -by- n matrices over \mathbb{K} is a quasiprojective variety, a closed subvariety of which is $\mathrm{SL}_n(\mathbb{K})$ comprising of matrices with determinant one. In fact, all affine algebraic groups are isomorphic to subgroups of $\mathrm{GL}_n(\mathbb{K})$, and a result of CHEVALLEY (1960) states that the remaining ones are of the type we shall next discuss.

Proposition II.1.6. *Every algebraic group \mathcal{G} has a unique normal subgroup \mathcal{H} isomorphic to an affine variety such that \mathcal{G}/\mathcal{H} is projective and irreducible.*

ABELIAN VARIETIES

Definition II.1.7. *Abelian varieties are irreducible projective algebraic groups.*

Most of the rich structure of abelian varieties stems from the projectiveness condition (completeness, an algebraic equivalent to compactness, could equivalently be required).

Proposition II.1.8. *Any algebraic map from an abelian variety to another is a morphism (of algebraic groups) composed with a translation.*

In other words, morphisms of algebraic varieties are essentially morphisms of abelian varieties; this means that abelian varieties are entirely characterized by their geometry. This is a crucial fact with the notable consequence that *abelian varieties are commutative groups*; indeed, since the algebraic map $x \mapsto -x$ fixes the neutral element, it is a morphism, which implies the commutativity.

Since abelian varieties \mathcal{A} are commutative, they admit quotients by any closed subgroups \mathcal{H} . We will later be interested in the case of finite subgroups \mathcal{H} , which are evidently closed: in that case, the dimension of the quotient \mathcal{A}/\mathcal{H} is the same as that of the variety \mathcal{A} , and as we will see later, many other invariants are preserved.

As a further restriction to prevent unnecessary contortions, we henceforth assume, unless otherwise stated, that all abelian varieties we consider are *absolutely simple*, that is, do not contain any proper nontrivial abelian subvariety over an algebraic closure.

II.2 Practical Settings

Let us now focus on two types of base field: finite fields, over which abelian varieties admit efficient representations, and the complex numbers, over which their relationship to tori yields a rich theory, part of which descends to finite fields.

FINITE FIELDS

Let \mathcal{A} be an abelian variety defined over a finite field $k = \mathbb{F}_q$; its *zeta function*

$$Z_{\mathcal{A}}(t) = \exp \sum_{n=1}^{\infty} \# \mathcal{A}(\mathbb{F}_{q^n}) \frac{t^n}{n}$$

encodes its number of points, on which WEIL (1945) proved the following.

Theorem II.2.1. *The zeta function of a dimension- g abelian variety \mathcal{A} is of the form*

$$Z_{\mathcal{A}}(t) = \prod_{n=0}^{2g} P_n(t)^{(-1)^{n+1}}$$

for some polynomials $P_n \in \mathbb{Z}[t]$ whose complex zeroes have absolute value $q^{-n/2}$.

This constrains cardinalities of abelian varieties. To better see this, consider the *Frobenius endomorphism* π , which acts over any field extension K/\mathbb{F}_q by raising coordinates of points of $\mathcal{A}(K)$ to the q^{th} power; it fixes just $\mathcal{A}(\mathbb{F}_q)$, so we have $\# \mathcal{A}(\mathbb{F}_q) = \deg(1 - \pi)$.

Any endomorphism ϕ of an abelian variety of dimension g has a monic characteristic polynomial $P \in \mathbb{Z}[t]$ of degree $2g$ such that $\deg Q(\phi) = \text{Res}(P, Q)$ for all polynomials $Q \in \mathbb{Z}[t]$. For the particular Frobenius endomorphism, denoting by χ_{π} its characteristic polynomial, we obtain

$$\# \mathcal{A}(\mathbb{F}_{q^n}) = \text{Res}_u(\chi_{\pi}(u), u^n - 1)$$

which makes computing χ_{π} equivalent to counting points on \mathcal{A} over g distinct field extensions of the base field. Transcribing the theorem above to χ_{π} yields the following.

Corollary II.2.2. *The complex roots of χ_{π} all have absolute value \sqrt{q} , and the polynomial $P_{2g}(t)$ in the zeta function is $\prod (1 - \alpha t)$ where α ranges over products of $2g$ distinct such roots.*

Generalizing an algorithm of SCHOOF (1985), PILA (1990) proved that for any fixed dimension g all the above can be computed in polynomial time in the size of the base field.

Theorem 11.2.3. *The zeta function of an abelian variety defined over \mathbb{F}_q can be computed in polynomial time in $\log(q)$ where the implied exponent depends on the dimension of the projective space where it is embedded, and on the degrees of its defining equations and group law equations.*

This result is mostly of theoretical interest. Improvements on the algorithm of SCHOOF (1985) by ATKIN and ELKIES have made it possible to count points on abelian varieties of dimension $g = 1$ far beyond cryptographic range; for $g = 2$, the practicality of point counting methods on varieties of cryptographic size was only recently demonstrated by GAUDRY and SCHOST (2010) who used an extension of the algorithm of SCHOOF (1985).

From now on, we shall regard the dimension g as being fixed in complexity statements, so asymptotic analyses focus on behavior with respect to the base field; this is partly motivated by the fact that only $g = 1$ and $g = 2$ are cases of cryptographic interest.

COMPLEX NUMBERS

We have noted that abelian varieties are nonsingular. Over \mathbb{C} , abelian varieties are therefore connected compact Lie groups, which are well-understood objects; such a variety \mathcal{A} has the analytic structure of a complex torus: since the exponential map folds its tangent space onto \mathcal{A} , there is an isomorphism of Lie groups $\mathcal{A} \simeq \mathbb{C}^g / \Lambda$ where $\Lambda = \ker(\exp_{\mathcal{A}})$ is a lattice of \mathbb{C}^g , that is, a discrete subgroup of full rank.

Similarly to the algebraic case, holomorphic maps between complex tori are just group morphisms composed by translations. Holomorphic morphisms ϕ from a complex torus $\mathbb{T} = \mathbb{C}^g / \Lambda$ to another $\mathbb{T}' = \mathbb{C}^{g'} / \Lambda'$ are induced by \mathbb{C} -linear maps, denoted ϕ as well, from \mathbb{C}^g to $\mathbb{C}^{g'}$ satisfying $\phi(\Lambda) \subset \Lambda'$. Hence, as \mathbb{Z} -module, $\text{Hom}(\mathbb{T}, \mathbb{T}')$ has rank at most $4gg'$; this implies that $\text{End}(\mathcal{A})$ is a torsion-free \mathbb{Z} -algebra of dimension at most $(2g)^2$.

Even if complex abelian varieties have the analytic structure of tori, conversely, not all complex tori correspond to abelian varieties, although those that do are precisely known:

Proposition 11.2.4. *Define the Siegel upper half-space \mathbb{H}^g as the set of g -by- g symmetric matrices with positive definite imaginary part. Complex tori \mathbb{C}^g / Λ corresponding to abelian varieties are exactly those whose lattice Λ can be put under the form $\mathbb{Z}^g + \Omega\mathbb{Z}^g$ for some matrix $\Omega \in \mathbb{H}^g$.*

POLARIZATIONS

Many results on abelian varieties over finite fields exploit *reduction* from characteristic zero fields k , that is, consider varieties arising through maps $k \rightarrow k/\mathfrak{p}$ for prime ideals \mathfrak{p} of k .

For instance, the bound of HASSE (1934) which states that one-dimensional abelian varieties \mathcal{A} defined over \mathbb{F}_q satisfy

$$\left|q + 1 - \#\mathcal{A}(\mathbb{F}_q)\right| \leq 2\sqrt{q}$$

can be extended, for varieties arising as reductions from characteristic zero, into a precise description of the distribution of cardinalities: the Sato–Tate conjecture. Note that recent work of TAYLOR (2008) comes close to proving it.

Conjecture II.2.5. *Let \mathcal{A} be a non-empty abelian variety of dimension one defined over the rationals with $\text{End}(\mathcal{A}) \simeq \mathbb{Z}$. The asymptotic distribution, as the prime p goes to infinity, of*

$$\arccos\left(\frac{p + 1 - \#\mathcal{A}(\mathbb{F}_p)}{2\sqrt{p}}\right)$$

is uniform on $[0; \pi]$ where $\#\mathcal{A}(\mathbb{F}_p)$ denotes the number of points of the reduction of \mathcal{A} at p .

When $g > 1$, abelian varieties have infinite automorphism groups over algebraically closed fields. For more rigidity, we bundle them with a projective embedding or, rather, the following (simpler) analytic analog.

Definition II.2.6. *Let $\mathcal{A} \simeq \mathbb{C}^g/\Lambda$ be a complex torus. A polarization of \mathcal{A} is a positive definite Hermitian form \mathcal{P} on \mathbb{C}^g satisfying $\mathcal{P}(\Lambda, \Lambda) \subset \mathbb{Z}$. It is principal if its determinant is invertible, or equivalently if there is no $x \notin \Lambda$ satisfying $\mathcal{P}(\Lambda, x) \subset \mathbb{Z}$.*

Principally polarized abelian varieties are pairs $(\mathcal{A}, \mathcal{P})$ whose morphisms $\phi : (\mathcal{A}, \mathcal{P}) \rightarrow (\mathcal{A}', \mathcal{P}')$ are required to preserve polarizations in the sense that $\phi^* \mathcal{P}' = \lambda \mathcal{P}$ for some positive $\lambda \in \mathbb{Q}$. WEIL (1955) showed that this has the intended effect:

Proposition II.2.7. *Polarized abelian varieties have a finite automorphism group.*

For instance, on the torus $\mathbb{C}^g/(\mathbb{Z}^g + \Omega\mathbb{Z}^g)$ for $\Omega \in \mathbb{H}^g$, there is a natural polarization $\mathcal{P}(u, v) = E(iu, v) + iE(u, v)$ where the Riemann form E is expressed, on the block basis $(e_i)(\Omega e_i)$, by the block matrix

$$\begin{pmatrix} 0 & \text{Id} \\ -\text{Id} & 0 \end{pmatrix}$$

Proposition II.2.8. *Two matrices Ω and Ω' of the Siegel upper half-space \mathbb{H}^g yield isomorphic principally polarized abelian varieties if and only if they are conjugate under the action*

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} \in \text{Sp}_{2g}(\mathbb{Z}) : \Omega \mapsto (A\Omega + B)(C\Omega + D)^{-1}.$$

Polarizations are needed in actual computations, as efficient arithmetic (via theta functions or Jacobian varieties) relies on them. Worse, it is nontrivial to determine whether the varieties corresponding to two theta coordinates are isomorphic, disregarding polarizations.

Before moving on, we emphasize once more that, in dimension one, all varieties admit a unique principal polarization — so they can hopefully be forgotten altogether.

JACOBIAN VARIETIES

Theorem II.2.9. *Up to isomorphism, there is a unique abelian variety through which any morphism from a given algebraic variety \mathcal{V} to an abelian variety factors. It is the Albanese variety of \mathcal{V} .*

General Albanese varieties are hardly practical: they have no effective group law, and are not naturally endowed with a principal polarization, so there is no simple manner to identify them such as invariants (as we will see below). Cryptography is only concerned with the following subclass, on which our exposition shall now focus.

Proposition II.2.10. *Abelian varieties of dimension one or two are Jacobian varieties of hyperelliptic curves.*

Before defining hyperelliptic curves, let us briefly discuss *Jacobian varieties*: these are just Albanese varieties of *algebraic curves*, that is, one-dimensional algebraic varieties. The Jacobian variety $\text{Jac}(\mathcal{C})$ of a curve \mathcal{C} has an explicit group structure: denote by Div^0 the submodule of degree-zero divisors of the free \mathbb{Z} -module generated by points of \mathcal{C} , that is, formal sums of points whose coefficients add up to zero; it contains Princ , the set of sums of zeroes and poles (counted with multiplicities) of non-zero elements of the function field.

Proposition II.2.11. *$\text{Jac}(\mathcal{C})$ has the group structure of the quotient $\text{Div}^0 / \text{Princ}$.*

We can say much more for hyperelliptic curves; for this, we assume $\text{char } k \neq 2$.

Definition II.2.12. *Curves \mathcal{C} of the form $y^2 = f(x)$, for some squarefree polynomial f of degree $2g + 1$ or $2g + 2$, are called hyperelliptic, and g is known as the genus of \mathcal{C} .*

By the theorem of RIEMANN (1857) and ROCH (1865), g is also the dimension of $\text{Jac}(\mathcal{C})$. In the case that $g = 1$, they are known as elliptic curves, and verify $\text{Jac}(\mathcal{C}) \simeq \mathcal{C}$.

When $\deg(f)$ is odd, there is a unique projective, non-affine point (with coordinate $z = 0$); this *point at infinity* ∞ is often used as a distinguished projective point. By RIEMANN (1857) and ROCH (1865) each divisor class then has a unique *reduced* representative of the form $\sum(P_i - \infty)$ for at most g affine points $P_i \in \mathcal{C}$, none of which is conjugate to another under the *hyperelliptic involution* $(x, y) \mapsto (x, -y)$.

Assume, for simplicity, that the points $P_i = (x_i, y_i)$ are distinct. The divisor $\sum(P_i - \infty)$ can be represented by a pair of polynomials (u, v) satisfying

$$u(x) = \prod (x - x_i), \quad v(x_i) = y_i.$$

It can be checked that the P_i lie on \mathcal{C} by verifying that $u|v^2 - f$. In this representation, the group law is given by (assuming u_0 and u_1 have no common root)

$$(u_0, v_0) + (u_1, v_1) = (u_0 u_1, (u_2^{-1} \bmod v_2) u_2 v_1 + (u_1^{-1} \bmod u_2) u_1 v_2).$$

To reduce the output to a unique representative, CANTOR (1987) iterates the transformation

$$(u, v) \mapsto (u', v') \quad \text{with } u' = \frac{1}{\text{lc}(f - v^2)} \frac{f - v^2}{u} \text{ and } v' = -v \bmod u'$$

while $\deg(u) \leq g$, where $\text{lc}(\cdot)$ denotes the leading coefficient. This gives $\text{Jac}(\mathcal{C})$ an efficient group law, and an algebraic structure. Additionally, the image of the map $(P_i) \in \mathcal{C}^{g-1} \mapsto \sum(P_i - \infty)$ is a subvariety of dimension $g - 1$ that is the zero-locus of certain theta functions which naturally endow the Jacobian variety with a principal polarization \mathcal{P} .

TORELLI (1913) showed that this comprises all the information from the original curve:

Theorem II.2.13. *Up to isomorphism, the polarized abelian variety $(\text{Jac } \mathcal{C}, \mathcal{P})$ determines the curve \mathcal{C} .*

Moduli spaces are varieties whose points represent isomorphism classes of a given type of variety (we will soon discuss invariants); complementing the proposition above, we have:

<i>the moduli dimension of</i>	genus- g hyperelliptic curves	<i>is</i>	$2g - 1$
”	genus- g curves	”	$3(g - 1)$, or 1 if $g = 1$
”	abelian varieties of dimension g	”	$g(g + 1)/2$

The moduli space dimension is the same for Jacobian varieties and their underlying curves. For $g = 3$, abelian varieties are Jacobian varieties, but not all of hyperelliptic curves.

II.3 Pairings

TORSION SUBGROUPS

The center of the endomorphism ring $\text{End}(\mathcal{A})$ of an abelian variety \mathcal{A} of dimension g always contains a subring isomorphic to \mathbb{Z} formed by *scalar multiplication maps*:

$$[n] : P \in \mathcal{A} \mapsto nP = \underbrace{P + \dots + P}_{n \text{ times}}$$

for every integer n . Over an algebraic closure, the kernel of $[n]$ is the *full n -torsion subgroup* $\mathcal{A}[n]$; its structure is well understood:

Theorem II.3.1. *The degree of $[n]$ is n^{2g} . It is separable when n is coprime to $p = \text{char } k$; then $\mathcal{A}[n] \simeq (\mathbb{Z}/n)^{2g}$. When n is a power of p , then $\mathcal{A}[n] \simeq \mathbb{Z}/n^r$ where $r \leq g$ is called the p -rank of \mathcal{A} .*

The generic case is that of *ordinary* abelian varieties which have p -rank g : the moduli dimension of non-ordinary varieties is strictly smaller. Unless explicitly stated, all abelian varieties will now be assumed ordinary (this is crucial for the next chapter).

We will later compute ℓ -torsion subgroups (for primes ℓ) of abelian varieties \mathcal{A} defined over finite fields \mathbb{F}_q . The *embedding degree* $e_{\mathcal{A}}(\ell)$, which is the extension degree of the smallest field over which the points of $\mathcal{A}[\ell]$ are defined, is the primary cost factor of this process.

If χ is the characteristic polynomial of the Frobenius endomorphism π of \mathcal{A} , the morphism $\chi(\pi)$ obviously vanishes on $\mathcal{A}[\ell]$; as this only depends on the class of χ in $(\mathbb{Z}/\ell)[x]$, the embedding degree $e(\ell)$ must divide the multiplicative order of $x \in (\mathbb{Z}/\ell)[x]/(\chi)$. Consequently, it is bounded by ℓ^{2g} .

When points can be drawn uniformly at random from $\mathcal{A}(k^{\ell})$, a basis for $\mathcal{A}[\ell]$ can be found by taking random points, multiplying them by the cofactor of ℓ^∞ in $\#\mathcal{A}(k^{\ell})$, and iteratively applying $[\ell]$ until a point of ℓ -torsion is found, possibly lifting points already found along their preimage under $[\ell]$. The lifting process can either use simple baby-step giant-step computations in $\mathcal{A}[\ell]$, or faster discrete logarithm methods in k^{ℓ} via the pairing. For a fixed g , the whole method uses polynomially many operations in ℓ ; it will be described in detail in the second half of this thesis.

GENERAL PAIRINGS

Definition II.3.2. *A pairing is a non-degenerate bilinear map $\Psi : G^2 \rightarrow H$, where G and H are abelian groups.*

Strictly speaking, pairings can be defined on modules over any ring; but from a cryptographic standpoint, nothing of value is lost by restricting to \mathbb{Z} -modules. On the other hand, cryptographic use requires additional properties:

- EASY EVALUATION: Given $(x, y) \in G^2$, the pairing $\Psi(x, y)$ is easily evaluated.
- HARD INVERSION: Given $z \in H$, a preimage $(x, y) \in \Psi^{-1}(z)$ is hard to find.

These terms could be given a rigorous meaning by considering a sequence of pairings $\Psi_i : G_i^2 \rightarrow H_i$, and requesting that there exists an algorithm for evaluating Ψ_i in polynomial time in $\log(\#G_i)$ and that no algorithm finds preimages of Ψ_i in subexponential time on a

positive fraction of H_i ; however, we prefer to use the simpler and down-to-earth notion of computational infeasibility.

Similarly to the discrete logarithm problem, the pairing inversion problem has many variants, such as bilinear analogs to the computational and decisional Diffie–Hellman problems, or inversion problems where one of the parameters is fixed, not all of which are strictly equivalent to the pairing inversion problem itself. We refer to GALBRAITH, HESS, and VERCAUTEREN (2008) for a discussion of these problems.

Out of all known effective pairings, only those that arise from abelian varieties satisfy the conditions above. In fact, the *problem of pairing inversion*, that is, of inverting the map Ψ , appears to be extremely difficult for such pairings. Their cryptographic use therefore involves relying on a new hypothesis (alongside the hardness of the discrete logarithm problem) but they provide elliptic and hyperelliptic cryptography with a unique structure, which has led to the development of many novel features.

EFFECTIVE PAIRINGS

Instructional pairing examples include scalar products of vector spaces, and, if $(R, +, \times)$ is a ring, the multiplication map from $(R, +)^2$ to (R, \times) . A more interesting example is

$$(xy, x'y') \in ((\mathbb{Z}/n)^{2g})^2 \mapsto \exp\left(\frac{2\pi i}{n} (\widehat{x}y' - \widehat{y}x')\right)$$

where xy denotes the concatenation of the row vectors $x, y \in (\mathbb{Z}/n)^g$, and \widehat{x} denotes the transpose of x . This actually is the general form of the Weil pairing expressed on a symplectic basis of the n -torsion subgroup of a complex torus.

None is suitable for cryptographic use, as they are typically easy to invert; currently, the only known cryptographic pairings arise from abelian varieties:

Let \mathcal{A} be the Jacobian variety $\text{Jac}(\mathcal{C})$ of a curve \mathcal{C} of genus g , which we further assume to be a hyperelliptic curve defined over a finite field. Recall that the full n -torsion subgroup $\mathcal{A}[n]$ is isomorphic to $(\mathbb{Z}/n)^{2g}$ when n is coprime to the ambient characteristic. For cryptographic reasons we choose n to be prime, and define the map

$$\Psi_{\text{Weil}} : \begin{cases} \mathcal{A}[n] \times \widehat{\mathcal{A}}[n] & \rightarrow \mu_n \subset \bar{k}^\times \\ (P, Q) & \rightarrow f_P(Q)/f_Q(P) \end{cases}$$

where μ_n is the group of n^{th} roots of unity, and f_P and f_Q are functions of $\bar{k}(\mathcal{A})$ with disjoint support whose sum of zeroes and poles are the principal divisors nP and nQ , respectively. Its evaluation at a divisor $Q = \sum Q_i$ is explicitly $\prod f(Q_i)$.

Theorem II.3.3. Ψ_{Weil} is a Galois-invariant antisymmetric pairing called the Weil pairing.

Most of the proof relies on the reciprocity of WEIL (1940).

When \mathcal{A} is principally polarized, the polarization gives an isomorphism $\mathcal{A} \simeq \widehat{\mathcal{A}}$, and the pairing can therefore be defined on $\mathcal{A}[n] \times \mathcal{A}[n]$.

In the case of elliptic curves, points P of the variety are of the form $R - \infty$ where R is a point of the curve or the point at infinity itself. MILLER (1986) noted that the function f_i whose sum of zeroes is the principal divisor $iR - [i]R - (i-1)\infty$ can be computed iteratively by setting $f_{i+j} = f_i \cdot f_j \cdot u/v$, where u is the line containing $[i]R$ and $[j]R$ (it vanishes at $[i]R$, $[j]R$, and $-[i+j]R$, and has a pole of order 3 at ∞) and v the vertical line passing through $[i+j]R$ (it vanishes at $[i+j]R$ and $-[i+j]R$, and has a pole of order 2 at ∞).

This yields an algorithm for evaluating the Weil pairing of elliptic curves which can also be extended to Jacobian varieties of hyperelliptic curves following Cantor's algorithm for evaluating the group law. Pairings of general abelian varieties were recently shown by LUBICZ and ROBERT (2010) to be effectively computable as well.

CRYPTOGRAPHIC APPLICATIONS

Before novel cryptographic primitives exploited their structure, pairings were mainly used as a cryptanalysis tool. Indeed, if P and Q are two points in a subgroup of prime order ℓ of a variety \mathcal{A} , the bilinearity of pairings implies

$$\Psi(P, Q) = \Psi(P, P)^{\log_p Q}$$

which shows that $\log_p Q$ is also the discrete logarithm problem of $\Psi(P, Q)$ in base $\Psi(P, P)$ in an extension K/k of degree $e(\ell)$. Since discrete logarithm problems are much easier over finite fields, $e(\ell)$ must be big enough to compensate for this weakness.

The last ten years have, on the other hand, seen pairings enabling innovative cryptographic constructions, so that the extra structure they give to abelian varieties is now seen as a feature. To exploit them, the value of $e(\ell)$ is selected large enough to make attacks impracticable on the discrete logarithm of the field K , but still low enough so as to permit the efficient evaluation of the pairing.

As an example of the new features enabled by abelian varieties and their pairings, we can for instance recall the one-round tripartite Diffie–Hellman key-exchange of JOUX (2000) that we presented in the previous chapter.

II.4 Isogenies

ABSTRACT ISOGENIES

Definition II.4.1. An isogeny is a surjective morphism of abelian varieties $\phi : \mathcal{A} \rightarrow \mathcal{B}$ with finite kernel. It is separable if the corresponding function field extension $k(\mathcal{A})/\phi^*(k(\mathcal{B}))$ is.

When $\phi : \mathcal{A} \rightarrow \mathcal{B}$ is an isogeny, the abelian varieties \mathcal{A} and \mathcal{B} are said to be *isogenous*; this is an equivalence relation since there then exists a *dual isogeny* $\bar{\phi} : \mathcal{B} \rightarrow \mathcal{A}$, of the same degree n , which is simply the multiplication-by- n map of \mathcal{A} factored through ϕ .

$$\text{deg } \phi \quad \begin{array}{ccc} & \phi & \\ \curvearrowright & \mathcal{A} & \xrightarrow{\quad} \mathcal{B} \\ & \bar{\phi} & \\ & \curvearrowleft & \end{array}$$

Proposition II.4.2. If \mathcal{H} is the kernel of a separable isogeny $\phi : \mathcal{A} \rightarrow \mathcal{B}$, then ϕ is the projection map under the isomorphism $\mathcal{B} \simeq \mathcal{A}/\mathcal{H}$; in particular, we have $\text{deg}(\phi) = \#\mathcal{H}$.

The group structure of \mathcal{H} is called the type of ϕ .

From now on, the word “isogeny” should implicitly mean “separable isogeny;” this is the case for all isogenies whose degree is coprime to the characteristic of the base field.

Since composition of isogenies corresponds to inclusion of subgroups, and the latter are abelian, we deduce that all isogenies can be written as the composition of isogenies of prime degree. In dimension $g > 1$, although there is currently no known method for computing general isogenies of type \mathbb{Z}/ℓ where ℓ is a prime, there are algorithms for evaluating isogenies of type $(\mathbb{Z}/\ell)^g$ which we call ℓ -isogenies.

Recall that we assume isogenies between principally polarized abelian varieties \mathcal{A} to preserve polarizations. The induced polarization on \mathcal{A}/\mathcal{H} for a finite subgroup \mathcal{H} is principal if and only if \mathcal{H} is a maximal isotropic subgroup for the Weil pairing; when we compute isogenies from their kernel, we will first start by enumerating all such subgroups.

HONDA–TATE THEORY

Over finite fields, there is a bijection between isogeny classes of abelian varieties and their zeta functions. We have already explained the relationship between the zeta function of an abelian variety and the characteristic polynomial of its Frobenius endomorphism, and the following description of isogeny classes is due to TATE (1966).

Theorem II.4.3. Two varieties are isogenous if and only if their respective Frobenius endomorphisms have the same characteristic polynomial.

A monic polynomial with integer coefficients and $2g$ complex roots, each of absolute value \sqrt{q} , is called a *q-Weil polynomial*. Recall that this is the case of the characteristic polynomial of the Frobenius endomorphism. As a reciprocal to that statement, HONDA (1968) proved:

Theorem II.4.4. *Each q-Weil polynomial is the characteristic polynomial of the Frobenius endomorphism of a certain simple ordinary abelian variety of dimension g defined over \mathbb{F}_q .*

TATE (1968) presented these two theorems in a combined way, and this has become known as Honda–Tate theory.

The next chapter will be concerned with an *explicit* form of this theory which aims at constructing explicit abelian varieties whose Frobenius endomorphisms have prescribed characteristic polynomials. This enforces certain properties on the abelian variety, such as the cardinality.

EXPLICIT ISOGENIES

For elliptic curves \mathcal{E} , VÉLU (1971) gave explicit formulas for computing an isogeny $\phi : \mathcal{E} \rightarrow \mathcal{E}'$ defined by its kernel $\ker(\phi) \subset \mathcal{E}$: if x, y are coordinates in which an affine equation for \mathcal{E} is $y^2 = f(x)$, then there exist coordinates X, Y in which an equation for \mathcal{E}' has the form $Y^2 = g(X)$ and the isogeny can be written as

$$\phi : P \in \mathcal{E} \mapsto \begin{pmatrix} X_{\phi(P)} = \sum x_{P+Q} - x_Q \\ Y_{\phi(P)} = \sum y_{P+Q} - y_Q \end{pmatrix}$$

where the sums range over all points Q of $\ker(\phi)$, with the convention that $x_\infty = y_\infty = 0$.

This relies heavily on properties of the Weierstrass coordinates for elliptic curves, and a higher-dimensional analog was only found recently by LUBICZ and ROBERT (2009), and later made practical by COSSET and ROBERT (2011); it relies on the structure of theta functions, which we now briefly describe.

Geometric invariants identify isomorphism classes of abelian varieties. For instance, isomorphism classes of elliptic curves are identified, over an algebraic closure, by the canonical *j-invariant*. It is effective as $j(\mathcal{E})$ is a rational function in the coefficients of a Weierstrass equation for \mathcal{E} , and conversely the coefficients of such an equation are rational functions in $j(\mathcal{E})$.

In arbitrary dimension, a system of invariants for principally polarized abelian varieties is given by *theta constants*, which not only identify the isomorphism class of a variety but also part of its torsion. Theta constants are the constant terms of *theta functions* which yields a convenient *coordinate system* for points on the variety it identifies.

In the particular case of abelian varieties of dimension $g < 4$, which are all, up to isomorphism, Jacobian varieties of algebraic curves, invariants can be expressed, via Torelli's theorem, on the curves themselves, as functions of the coefficients of their equations. For $g = 2$, a popular set of invariants are the *Igusa invariants*, which consists of 10 coordinates (this bears some redundancy since the dimension of the moduli space is 3); they can be efficiently computed from the equation of a curve, but conversely, to retrieve such an equation from the invariants themselves, a specific method of MESTRE (1991) is required.

The relationship between the invariants of a curve and the theta constants of its Jacobian variety are given by formulas of THOMAE (1870).

Let $\mathcal{A} \simeq \mathbb{C}^g / (\mathbb{Z}^g + \Omega\mathbb{Z}^g)$ be a complex torus with $\Omega \in \mathbb{H}^g$. Define the *theta functions*

$$\Theta_{a,b}^{\mathcal{A}} : z \in \mathbb{C}^g \mapsto \sum_{(u+a) \in \mathbb{Z}^g} \exp i\pi \left(\frac{1}{n} \widehat{u} \Omega u + 2\widehat{u}(z+b) \right)$$

where a and b are vectors of \mathbb{Q}^g and \widehat{u} denotes the transpose of u . IGUSA (1960) proved:

Theorem II.4.5. *Fix an integer $n > 2$. The theta constants $\Theta_{a,b}^{\mathcal{A}}(0)$ for $a, b \in \{\frac{1}{n}, \dots, \frac{n}{n}\}^g$ uniquely determine the isomorphism class of \mathcal{A} as a principally polarized abelian variety.*

Details on implementing and practically computing isogenies between abelian varieties of dimension two will be found in the last chapter.

MODULAR CORRESPONDENCE

Some applications do not require to explicitly evaluate isogenies, that is, to effectively evaluate the map: it is sometimes sufficient to enumerate abelian varieties which are (rationally) ℓ -isogenous to a prescribed abelian variety \mathcal{A} , for a given prime ℓ , and there could exist a faster way than enumerating all subgroups of type $(\mathbb{Z}/\ell)^g$ and then evaluating the associated isogenies.

Ideally, this information could be encoded in polynomials via invariants $I(\mathcal{A}) \in k^n$: we would have n polynomials $\Phi_\ell^i(X_1, \dots, X_n, Y_1, \dots, Y_n)$ for $i \in \{1, \dots, n\}$ such that

$$\mathcal{A} \text{ is } \ell\text{-isogenous to } \mathcal{B} \iff \begin{cases} \Phi_\ell^1(I_1(\mathcal{A}), \dots, I_n(\mathcal{A}), I_1(\mathcal{B}), \dots, I_n(\mathcal{B})) = 0, \\ \Phi_\ell^2(I_1(\mathcal{A}), \dots, I_n(\mathcal{A}), I_1(\mathcal{B}), \dots, I_n(\mathcal{B})) = 0, \\ \vdots \\ \Phi_\ell^n(I_1(\mathcal{A}), \dots, I_n(\mathcal{A}), I_1(\mathcal{B}), \dots, I_n(\mathcal{B})) = 0, \end{cases}$$

For elliptic curves, this is achieved by the classical *modular polynomials* $\Phi_\ell(X, Y)$. ENGE (2009) computed them via the floating point method which consists in evaluating Φ over

the complex number with just enough precision so as to identify its integer coefficients. Recently, BRÖKER, LAUTER, and SUTHERLAND (2010) demonstrated the competitiveness of a method based on the Chinese remainder theorem which exploits the structure of isogeny volcanoes that we will study later.

The higher-dimensional case is not as straightforward: GAUDRY (2000) described an analog construction for $g = 2$, and the computation of explicit polynomials was later done by DUPONT (2006) and improved by BRÖKER and LAUTER (2009). However, the height of the polynomials (Φ_ℓ^i) makes their use prohibitive; currently, state-of-the-art algorithms for explicitly evaluating isogenies remain a faster alternative.

We note that this difference between elliptic curves and higher-dimensional abelian varieties is the main reason why point counting algorithms are much faster for the former than for the latter.

References

1857. Bernhard RIEMANN.
 “Theorie der Abel’schen Functionen”.
 In: *Journal für die reine und angewandte Mathematik* 1857.54. Pages 115–155.
 DOI: 10.1515/crll.1857.54.115.
1865. Gustav ROCH.
 “Über die Anzahl der willkürlichen Constanten in algebraischen Functionen”.
 In: *Journal für die reine und angewandte Mathematik* 1865.64. Pages 372–376.
 DOI: 10.1515/crll.1865.64.372.
1870. Carl J. THOMAE.
 “Beitrag zur Bestimmung von $\vartheta(0, 0, \dots, 0)$ durch die Klassenmoduln algebraischer Functionen”.
 In: *Journal für die reine und angewandte Mathematik* 1870.71. Pages 201–222.
 DOI: 10.1515/crll.1870.71.201.
1893. David HILBERT.
 “Über die vollen Invariantensysteme”.
 In: *Mathematische Annalen* 42.3. Pages 313–373. DOI: 10.1007/BF01444162.
1913. Ruggiero TORELLI.
 “Sulle varietà di Jacobi”.
 In: *Rendiconti della Reale Accademia Nazionale dei Lincei* 22.5. Pages 98–103.

1934. Helmut HASSE.
“Über die Kongruenzzetafunktionen”.
In: *Sitzungsberichte der Preussischen Akademie der Wissenschaften* 17. Pages 250–263.
1940. André WEIL.
“Sur les fonctions algébriques à corps de constantes fini”.
In: *Comptes Rendus de l'Académie des Sciences de Paris* 210. Pages 592–594.
1945. André WEIL.
Sur les courbes algébriques et les variétés qui s'en déduisent. Volume 7.
Actualités Scientifiques et Industrielles 1041.
Publications de l'Institut de Mathématique de l'Université de Strasbourg.
1955. André WEIL.
“On the theory of complex multiplication”.
In: *International Symposium on Algebraic Number Theory*. Tokyo and Nikko.
Science Council of Japan. Pages 9–22.
1960. Claude CHEVALLEY.
“Une démonstration d'un théorème sur les groupes algébriques”.
In: *Journal de Mathématiques Pures et Appliquées* 39. Pages 307–317.
1960. Jun-Ichi IGUSA.
“Arithmetic variety of moduli for genus two”.
In: *Annals of Mathematics* 72.3. Pages 612–649. DOI: 10.2307/1970233.
1966. John TATE.
“Endomorphisms of abelian varieties over finite fields”.
In: *Inventiones mathematicae* 2.2. Pages 134–144. DOI: 10.1007/BF01404549.
1968. Taira HONDA.
“Isogeny classes of abelian varieties over finite fields”.
In: *Journal of the Mathematical Society of Japan* 20.1–2. Pages 83–95.
DOI: 10.2969/jmsj/02010083.
1968. John TATE.
“Classes d'isogénie des variétés abéliennes sur un corps fini (d'après T. Honda)”.
In: *Séminaire Bourbaki*. Volume 21. 352. Pages 95–110.
1970. David MUMFORD.
Abelian Varieties. Volume 5.
Tata Institute of Fundamental Research Studies in Mathematics.
Oxford University Press.

1971. Jacques VÉLU.
“Isogénies entre courbes elliptiques”.
In: *Comptes Rendus de l'Académie des Sciences de Paris*. A 273. Pages 238–241.
1974. Igor R. SHAFAREVICH.
Basic Algebraic Geometry. Volume 2 13.
Grundlehren der mathematischen Wissenschaften. Springer.
1985. René SCHOOF.
“Elliptic curves over finite fields and the computation of square roots mod p ”.
In: *Mathematics of Computation* 44.170. Pages 483–494.
DOI: 10.2307/2007968.
1986. Gary CORNELL and Joseph H. SILVERMAN (editors).
Arithmetic Geometry. Springer. ISBN: 3-540-96311-1.
1986. Victor S. MILLER.
Short programs for functions on curves. Unpublished.
URL: <http://crypto.stanford.edu/miller/>.
1987. David G. CANTOR.
“Computing in the Jacobian of a hyperelliptic curve”.
In: *Mathematics of Computation* 48.177. Pages 95–101. DOI: 10.2307/2007876.
1990. Jonathan PILA.
“Frobenius maps of abelian varieties and finding roots of unity in finite fields”.
In: *Mathematics of Computation* 55.192. Pages 745–763.
DOI: 10.2307/2008445.
1991. Jean-François MESTRE.
“Construction de courbes de genre 2 à partir de leurs modules”.
In: *Effective methods in algebraic geometry — MEGA '90*.
Edited by Teo MORA and Carlo TRAVERSO. Volume 94. Progress in Mathematics.
Birkhäuser. Pages 313–334.
1998. James S. MILNE.
Abelian Varieties.
URL: <http://www.jmilne.org/math/CourseNotes/av.html>.
1998. Goro SHIMURA.
Abelian Varieties with Complex Multiplication and Modular Functions.
Princeton University Press. ISBN: 0-691-01656-9.

2000. Pierrick GAUDRY.
“Algorithmique des courbes hyperelliptiques et applications à la cryptologie”.
PhD thesis. École Polytechnique.
URL: <http://hal.inria.fr/tel-00514848/PDF/these.final.pdf>.
2000. Antoine JOUX.
“A one round protocol for tripartite Diffie–Hellman”.
In: *Algorithmic Number Theory — ANTS-IV*. Edited by Wieb BOSMA.
Volume 1838. Lecture Notes in Computer Science. Springer. Pages 385–393.
DOI: 10.1007/10722028_23.
2005. Roberto M. AVANZI, Henri COHEN, Christophe DOCHE, Gerhard FREY,
Tanja LANGE, Kim NGUYEN, and Frederik VERCAUTEREN.
Handbook of Elliptic and Hyperelliptic Curve Cryptography.
Discrete Mathematics and its Applications. Chapman & Hall.
ISBN: 1-58488-518-1.
2006. Régis DUPONT.
“Moyenne arithmetico-géométrique, suites de Brochardt et applications”.
PhD thesis. École Polytechnique.
URL: http://www.lix.polytechnique.fr/Labo/Regis.Dupont/these_soutenance.pdf.
2008. Steven D. GALBRAITH, Florian HESS, and Frederik VERCAUTEREN.
“Aspects of pairing inversion”.
In: *IEEE Transactions on Information Theory* 54.12. Pages 5719–5728.
DOI: 10.1109/TIT.2008.2006431.
2008. Richard TAYLOR.
“Automorphy for some l -adic lifts of automorphic mod l Galois representations, II”.
In: *Publications Mathématiques de l’Institut des Hautes Études Scientifiques* 108.1.
Pages 183–239. DOI: 10.1007/s10240-008-0015-2.
2009. Reinier BRÖKER and Kristin LAUTER.
“Modular polynomials for genus 2”. In: *London Mathematical Society Journal of
Computation and Mathematics* 12. Pages 326–339.
DOI: 10.1112/S146115700001546.
2009. Andreas ENGE.
“Computing modular polynomials in quasi-linear time”.
In: *Mathematics of Computation* 78.267. Pages 1809–1824.
DOI: 10.1090/S0025-5718-09-02199-1.

2009. David LUBICZ and Damien ROBERT.
Computing isogenies between abelian varieties. arXiv.org: 1001.2016.
2010. Reinier BRÖKER, Kristin LAUTER, and Andrew V. SUTHERLAND.
Modular polynomials via isogeny volcanoes.
To appear in *Mathematics of Computation*. arXiv.org: 1001.0402.
2010. Pierrick GAUDRY and Éric SCHOST.
Genus 2 point counting over prime fields. HAL-INRIA: 00542650.
2010. David LUBICZ and Damien ROBERT.
“Efficient pairing computation with theta functions”.
In: *Algorithmic Number Theory — ANTS-IX*.
Edited by Guillaume HANROT, François MORAIN, and Emmanuel THOMÉ.
Volume 6197. Lecture Notes in Computer Science. Springer. Pages 251–269.
DOI: 10.1007/978-3-642-14518-6_21.
2011. Romain COSSET and Damien ROBERT.
Computing (ℓ, ℓ) -isogenies in polynomial time on Jacobians of genus 2 curves.
IACR ePrint: 2011/143.

THREE

Complex Multiplication

The theory of complex multiplication describes endomorphism rings of abelian varieties; this thesis will investigate two of its applications, inverse of each other:

- constructing abelian varieties equipped with efficiently computable pairings;
- computing the endomorphism ring of prescribed abelian varieties.

There are many facets to complex multiplication theory; here, while trying to be somewhat general, we will focus on effective aspects in the case of dimension $g = 1, 2$, which are of primary interest to cryptography. For details, we refer to COX (1989) for $g = 1$, to STRENG (2010) for $g = 2$, and otherwise to SHIMURA (1998), CORNELL and SILVERMAN (1986), and MILNE (2006).

III.1 Endomorphism Rings

ABELIAN VARIETIES WITH COMPLEX MULTIPLICATION

Let us first consider the endomorphism ring structure of abelian varieties; via the following theorem of POINCARÉ and WEIL (1945), it suffices to consider simple varieties.

Theorem III.1.1. *Every abelian variety is isogenous to a product of powers of non-isogenous simple ones.*

The endomorphism ring of a perfect power \mathcal{A}^m is naturally the matrix algebra of dimension m^2 over the endomorphism ring of \mathcal{A} ; therefore, the endomorphism ring of a product $\prod \mathcal{A}_i^{m_i}$ of non-isogenous simple abelian varieties \mathcal{A}_i is $\prod \text{Mat}_{m_i}(\text{End } \mathcal{A}_i)$.

Since isogenies need not preserve endomorphism rings, the above does not completely rule out the case of non-simple varieties. Nevertheless, we will now assume that \mathcal{A} is a simple

abelian variety of dimension g . Its endomorphism ring $\text{End}(\mathcal{A})$ contains at least the scalar multiplication maps, which form a subring isomorphic to \mathbb{Z} . To better comprehend the ring $\text{End}(\mathcal{A})$, first consider the algebra $\mathbb{Q} \otimes \text{End}(\mathcal{A})$: if it contains a field K of degree $2g$, the variety \mathcal{A} is said to have *complex multiplication* by the number field K or, more precisely, by the order $K \cap \text{End}(\mathcal{A})$. Over number fields, this is a rare situation; but over finite fields, all ordinary abelian varieties have complex multiplication.

Recall that, over finite fields, the Frobenius endomorphism π of a dimension- g abelian variety \mathcal{A} admits a monic characteristic polynomial χ_π of degree $2g$, and that this polynomial uniquely identifies the isogeny class of \mathcal{A} . TATE (1966) further established the following, of which a proof can be found in WATERHOUSE and MILNE (1971).

Theorem III.1.2. *If \mathcal{A} is a simple abelian variety, the characteristic polynomial of its Frobenius endomorphism π is some power m^e of its minimal polynomial, whence $\mathbb{Q} \otimes \text{End}(\mathcal{A})$ is a division algebra of dimension $2eg$, and its center K is the field $\mathbb{Q}(\pi) \simeq \mathbb{Q}[x]/(m(x))$ of degree $2g/e$.*

The number field K is known as the *complex multiplication field* of \mathcal{A} . The structure of such fields can easily be investigated since they are quotients of $\mathbb{Q}[x]$ by q -Weil polynomials $\chi_\pi(x)$: under the embedding to $\mathbb{Q} \otimes \text{End}(\mathcal{A})$, the field K is an extension by the polynomial $X^2 - (\pi + \bar{\pi})X + q$ of the totally real field $K_+ = \mathbb{Q}(\pi + \bar{\pi})$. Therefore, complex multiplication fields are totally imaginary quadratic extensions of totally real number fields K_+ of degree g .

So far, we have not been too concerned about fields of definition; we will continue not to be, due to the following proposition.

Proposition III.1.3. *Endomorphism rings of simple ordinary abelian varieties defined over finite base fields are unaffected by base field extensions.*

COMPLEX TORI WITH COMPLEX MULTIPLICATION

Complex multiplication also concerns complex tori, and due to their simpler structure it yields a rich theory; many results concerning abelian varieties over finite fields are reductions of results on complex tori. For now, we assume that the base field is $k = \mathbb{C}$.

Let us first fix a particular embedding ι of the complex multiplication field K in $\mathbb{Q} \otimes \text{End}(\mathcal{A})$. The exponential map sends \mathcal{A} to a complex torus \mathbb{C}^g/Λ , and ι to an embedding $\iota' : K \rightarrow \text{End}(\mathbb{C}^g)$. Using representation theory, one can prove that, up to isomorphisms of \mathbb{C}^g , the map ι' is of the form

$$\iota_\Phi : \begin{cases} K & \longrightarrow & \mathbb{C}^g \\ x & \longmapsto & (\phi(x))_{\phi \in \Phi} \end{cases}$$

for a certain set Φ of g distinct embeddings of K in \mathbb{C} , no two of which are complex conjugate of each other, so that all $2g$ embeddings are in $\Phi \sqcup \overline{\Phi}$. This set Φ is called the *complex multiplication type* of the abelian variety \mathcal{A} .

Isogenies transport the embedding ι and type Φ from one variety to the next; by the following result, found for instance as Proposition 3.12 of MILNE (2006), fixing one is equivalent to fixing the other.

Proposition III.1.4. *There is a bijection between the set of isogeny classes of simple ordinary pairs (\mathcal{A}, ι) and the set of isomorphism classes of primitive types (K, Φ) .*

We will now consider abelian varieties \mathcal{A} endowed with an embedding ι or, equivalently, a complex multiplication type Φ .

Conversely, a complex torus with complex multiplication by a prescribed complex multiplication field K and type Φ can be constructed as follows. Let \mathfrak{a} be an integral ideal of K ; the g -tuple of embeddings Φ maps it to a certain lattice of \mathbb{C}^g and we may consider the complex torus $\mathbb{C}^g/\Phi(\mathfrak{a})$. To obtain a polarization as a Riemann form E on it, take an algebraic integer ξ that generates K/K_+ , whose imaginary part is totally positive, and whose square is a totally negative element of K_+ , then define E by

$$E_{\Phi}^{\xi}(\Phi(x), \Phi(y)) = \text{tr}(\xi \cdot \bar{x} \cdot y)$$

which takes integral values on $\Phi(\mathfrak{a})^2$ and thus induces a polarization on the complex torus $\mathbb{C}^g/\Phi(\mathfrak{a})$; it is obviously principal since ξ is invertible. Integral elements x of K can be seen acting as endomorphisms of the torus by

$$(z_i) \in \mathbb{C}^g \mapsto (z_i \phi_i(x))$$

where an ordering on the embeddings ϕ of Φ has been fixed by indexing them by $i \in \{1, \dots, g\}$. Since distinct orderings yield isomorphic complex tori, Φ can be simply thought of as a set.

Other transformations of the type yield isomorphic varieties as well. In the case (where we assume to be) of simple varieties, we have:

Theorem III.1.5. *All principally polarized complex tori with complex multiplication by a ring of integers \mathcal{O}_K arise, via the construction above, from a triple $(\Phi, \mathfrak{a}, \xi)$.*

Two triples $(\Phi, \mathfrak{a}, \xi)$ and $(\Phi', \mathfrak{a}', \xi')$ yield isomorphic polarized tori if and only if there exists an automorphism σ and an element γ of K such that $\Phi' = \Phi\sigma$, $\mathfrak{a}' = \gamma\mathfrak{a}$, and $\xi' = (\gamma\bar{\gamma})^{-1}\xi$.

COMPLEX MULTIPLICATION ORDERS

The complex multiplication field K embedded in $\mathbb{Q} \otimes \text{End}(\mathcal{A})$ is an important invariant; however, it fails to capture the exact isomorphism type of $\text{End}(\mathcal{A})$, which is precisely what the order $\mathcal{O} = K \cap \text{End}(\mathcal{A})$ does.

Generally-speaking, an *order* \mathcal{O} in a number field K is a lattice that is also a subring of the ring of integers \mathcal{O}_K — the latter is therefore commonly called the *maximal order*. In our context, there is also a *minimal order* due to the following result of WATERHOUSE (1969).

Proposition III.1.6. *Let K be the complex multiplication field of some ordinary abelian variety defined over a finite field k with Frobenius endomorphism π . The orders of K containing $\mathbb{Z}[\pi, \bar{\pi}]$ are exactly those that arise as endomorphism rings of abelian varieties defined over k with complex multiplication by K .*

The *Verschiebung endomorphism* $\bar{\pi}$ can also be written as $q\pi^{-1}$, since Theorem III.3.5 will show that the degree of an endomorphism is the norm of the corresponding number field element.

Now consider an abelian variety \mathcal{A} defined over a number field k . If \mathfrak{p} is a discrete place of k , its residue field k/\mathfrak{p} is finite, and we might obtain an abelian variety $\mathcal{A}_{\mathfrak{p}}$ over k/\mathfrak{p} , of the same dimension as \mathcal{A} , by pushing \mathcal{A} forward through the quotient map $k \rightarrow k/\mathfrak{p}$; when we do, we say that \mathcal{A} has *good reduction* at the prime \mathfrak{p} . Most things independent from \mathfrak{p} reduce nicely:

Proposition III.1.7. *Let \mathcal{A} and \mathcal{B} be two abelian varieties of the same dimension defined over a number field with good reduction at some discrete place \mathfrak{p} . The natural map $\text{Hom}(\mathcal{A}, \mathcal{B}) \rightarrow \text{Hom}(\mathcal{A}_{\mathfrak{p}}, \mathcal{B}_{\mathfrak{p}})$ is injective and preserves the degree of isogenies.*

Specialized to an abelian variety $\mathcal{A} = \mathcal{B}$ with complex multiplication, this states that reduction leaves the complex multiplication field unchanged and can only make the endomorphism ring larger.

When the reduction $\phi_{\mathfrak{p}}$ of an isogeny $\phi \in \text{End}(\mathcal{A})$ is separable, that is, whenever its degree is coprime to \mathfrak{p} , then the reduction map $\ker(\phi) \rightarrow \ker(\phi_{\mathfrak{p}})$ is a bijection.

NON-ORDINARY VARIETIES

For completeness, we briefly address the case of non-ordinary abelian varieties \mathcal{A} over a finite field \mathbb{F}_q ; the characteristic polynomial of the Frobenius endomorphism is then some proper power m^e with $e > 1$ of its minimal polynomial.

Contrary to the ordinary case, the endomorphism ring of non-ordinary abelian varieties might be smaller over the base field than it is over an algebraic closure.

For an elliptic curve, not being ordinary coincides with being *supersingular*, and also with the characteristic of the base field dividing the integer $\pi + \bar{\pi}$. Then, all endomorphisms are defined over \mathbb{F}_q if and only if q is a square and $\pi = \pm\sqrt{q}$.

Over fields with square cardinalities, there are thus two isogeny classes of supersingular curves with all endomorphisms defined, corresponding to the two q -Weil numbers $\pm\sqrt{q}$. Over a quadratic extension, those two become isogenous, but another isogeny class appears. Supersingular curves with not all endomorphisms defined can form up to three more isogeny classes. This has been rigorously studied by WATERHOUSE (1969), and to conclude we summarize his result concerning endomorphism rings of supersingular curves.

Proposition III.1.8. *Endomorphism rings of supersingular elliptic curves are*

- if all endomorphisms are defined: the maximal orders;
- otherwise: the p -maximal orders containing π ;

in the quaternion \mathbb{Q} -algebra ramified at infinity and p (the characteristic of the base field).

III.2 Orders and Ideals

For a moment, let us turn to topics of algebraic number theory with a computational flavor; they will later be put to use when we need to apply complex multiplication theory.

ALGEBRAIC ORDERS

Orders of a number field K are lattices (that is, discrete subgroups of full rank) with an induced ring structure; inclusion therefore yields a partial *order* on orders of K , where the italicized word is meant in the set-theoretic sense. From now on, we consider orders of a fixed complex multiplication field K , and refer to them just as “orders”; they are contained in the maximal order $\mathfrak{M} = \mathcal{O}_K$, and we are particularly interested in those containing a certain *minimal order* \mathfrak{m} of the form $\mathbb{Z}[\pi, \bar{\pi}]$. Since $K = \mathbb{Q}(\pi)$, there are finitely many such orders.

This induces a finite *lattice* structure (again, in the in the set-theoretic sense) and we will often be speaking about orders located above or below from others, meaning respectively that they contain or are contained in others. This structure extends to ideals: assuming $\mathcal{O} \subset \mathcal{O}'$ are two orders, we have natural maps

$$\begin{array}{ccc} \mathfrak{I}(\mathcal{O}') & & \mathfrak{I}(\mathcal{O}) \\ \mathfrak{a} & \mapsto & \mathfrak{a} \cap \mathcal{O} \\ \mathfrak{b} \mathcal{O}' & \longleftarrow & \mathfrak{b} \end{array}$$

and while the latter is a right inverse to the former, the converse is not true in general.

A more satisfying setting arises when we restrict to *invertible ideals* of an order \mathcal{O} , that is, fractional ideals \mathfrak{a} for which there exists another fractional ideal \mathfrak{b} satisfying $\mathfrak{a}\mathfrak{b} = \mathcal{O}$. All non-zero fractional ideals of the maximal order are invertible, but as we go down the lattice of orders, fewer and fewer are. To measure this notion of depth, we introduce the conductor, which measures how far \mathcal{O} is from its integral closure \mathfrak{M} .

Definition III.2.1. *The conductor of an order \mathcal{O} is the ideal $\mathfrak{f}_{\mathcal{O}} = \{x \in \mathfrak{M} : x\mathfrak{M} \subset \mathcal{O}\}$.*

The conductor gives a sufficient condition for invertibility: prime ideals that are coprime to $\mathfrak{f}_{\mathcal{O}}$ are invertible in \mathcal{O} . Conversely, up to principal ideals, all invertible ideals are equivalent to one coprime to the conductor. As a result, invertible ideals coprime to the conductor always have a unique decomposition into invertible prime ideals.

IDEAL CLASS GROUPS

Similarly to class groups of ring of integers, ideal class groups can be constructed from general orders. This construction resembles that of Jacobian varieties in terms of divisors, but the resulting group differs in various subtle aspects.

Definition III.2.2. *The Picard group of an order \mathcal{O} , denoted by $\text{Pic}(\mathcal{O})$, is the quotient group $\mathcal{I}(\mathcal{O})/\text{Princ}(\mathcal{O})$ of invertible ideals by principal ideals; it is finite and abelian.*

The Picard group of an order \mathcal{O} with conductor \mathfrak{f} is related to that of the maximal order $\mathfrak{M} = \mathcal{O}_{\mathbb{K}}$ via the exact sequence

$$1 \longrightarrow \mathcal{O}^{\times} \longrightarrow \mathfrak{M}^{\times} \longrightarrow (\mathfrak{M}/\mathfrak{f})^{\times}/(\mathcal{O}/\mathfrak{f})^{\times} \longrightarrow \text{Pic}(\mathcal{O}) \longrightarrow \text{Pic}(\mathfrak{M}) \longrightarrow 1$$

which shows that Picard groups grow roughly linearly in the norm of the conductor \mathfrak{f} ; more precisely, the sequence yields the following formula (which generalizes the well-known explicit formula for imaginary quadratic orders) for the *class number*:

$$\#\text{Pic}(\mathcal{O}) = \frac{\#\text{Pic}(\mathfrak{M}) \cdot \#(\mathfrak{M}/\mathfrak{f})^{\times}}{[\mathfrak{M}^{\times} : \mathcal{O}^{\times}] \cdot \#(\mathcal{O}/\mathfrak{f})^{\times}}$$

The asymptotic growth of the class number of the maximal order $h = \#\text{Pic}(\mathfrak{M})$ obeys the following conjecture of SIEGEL (1935) proved by BRAUER (1947).

Theorem III.2.3. *For any sequence of number fields \mathbb{K} whose class number, regulator, and discriminant we respectively denote by h , R , and Δ , we have:*

$$\frac{\log h + \log R}{\log \sqrt{|\Delta|}} \longrightarrow 1 \quad \text{as} \quad \frac{[\mathbb{K} : \mathbb{Q}]}{\log |\Delta|} \longrightarrow 0.$$

And we note that, for the fields K we are most interested in, namely quadratic and quartic complex multiplication fields, the regulator is respectively $R = 1$ and $R = O(\log|\Delta|)$.

Picard groups are compatible with the lattice-of-orders structure:

Proposition III.2.4. *Let $\mathcal{O} \subset \mathcal{O}'$ be two orders. The map $\mathfrak{a} \mapsto \mathfrak{a}\mathcal{O}'$, for invertible ideals \mathfrak{a} of \mathcal{O} coprime to $\mathfrak{f}_{\mathcal{O}'}$, induces a surjective morphism of Picard groups.*

Therefore, if some set \mathfrak{B} of ideals of the minimal order \mathfrak{m} generates its Picard group, it can be mapped into generating sets for each order above \mathfrak{m} . We form the free abelian group $\mathbb{Z}^{\mathfrak{B}}$, and let $\Lambda_{\mathcal{O}}$ denote the *lattice of relations of \mathcal{O}* , consisting of tuples $(\lambda)_{\mathfrak{B}}$ for which the product $\prod_{\mathfrak{B}} (\mathfrak{b}\mathcal{O})^{\lambda}$ is a principal ideal of \mathcal{O} . This gives a description of the Picard group as

$$\text{Pic}(\mathcal{O}) \simeq \mathbb{Z}^{\mathfrak{B}} / \Lambda_{\mathcal{O}}$$

and when one order is contained in another, their lattices of relations are too.

COMPUTING ORDERS

To list all possible endomorphism rings, that is, all orders containing $\mathfrak{m} = \mathbb{Z}[\pi, \bar{\pi}]$, one could simply focus on the lattice structure: subgroups of the quotient group $\mathfrak{M}/\mathfrak{m}$ can easily be enumerated, and each yields a lattice that contains \mathfrak{m} ; elementary techniques can then test whether such a lattice is closed under multiplication.

This approach is inefficient as most lattices are not orders, but also inadequate since there might be exponentially many orders above \mathfrak{m} . We can bound the conductor gap as follows:

Lemma III.2.5. *The index $[\mathfrak{M} : \mathfrak{m}]$ is bounded from above by $2^{g(g-1)} q^{g^2/2}$, where q is the norm of π and $2g$ its degree.*

Proof. Recall that $[\mathfrak{M} : \mathfrak{m}]$ is the square root of $\text{disc}(\mathfrak{m}) / \text{disc}(\mathfrak{M})$. The discriminant of the maximal order \mathfrak{M} can be small so we simply bound that of the minimal order \mathfrak{m} using

$$|\text{disc}(\mathfrak{m})| = |\text{disc}(\mathbb{Z}[\pi])| / [\mathbb{Z}[\pi, \bar{\pi}] : \mathbb{Z}[\pi]]^2.$$

The numerator can be bounded by $(2\sqrt{q})^{2g(2g-1)}$ since χ_{π} is a q -Weil polynomial of degree $2g$. For the denominator, we have $[\mathbb{Z}[\pi, \bar{\pi}] : \mathbb{Z}[\pi]] = q^{\frac{g(g-1)}{2}}$ from which the result follows. \square

Instead of enumerating all orders, we will navigate the lattice of orders and locate the endomorphism ring using complex multiplication theory. The proposition below shows that it suffices to *go up* or *down* by small powers of primes. Due to the lemma above, only polynomially many descending steps in g and $\log(q)$ are needed to reach \mathfrak{m} from \mathfrak{M} .

Proposition III.2.6. *Consider two orders $\mathcal{O}' \subset \mathcal{O}$ of relative index divisible by a prime ℓ . There exists an order \mathcal{O}'' in between whose index in \mathcal{O} is in $\{\ell, \ell^2, \dots, \ell^{2g-1}\}$ where $2g = \deg K$.*

To prove this, let \mathcal{O}'' be the order generated by $\ell\mathcal{O}$ and \mathcal{O}' : since $\ell\mathcal{O}$ has index ℓ^{2g} in \mathcal{O} and both contain \mathbb{Z} , its index in \mathcal{O} , and therefore also that of \mathcal{O}'' , must divide ℓ^{2g-1} .

Consider now the problem of *going down*, that is, enumerating all orders contained in a prescribed order \mathcal{O} with index n (to *go up* the process would be entirely equivalent).

In discussions with ENGE, we devised a simple method to enumerate all orders contained in a prescribed order \mathcal{O} with index n . The integer n should preferably be a small prime power to limit the size of the output; this amounts to considering the lattice of orders locally at this prime. When we only consider endomorphism rings of principally polarized abelian varieties, we can further restrict to those orders that are closed under complex conjugation.

Fix a \mathbb{Z} -module basis (ω_i) of \mathcal{O} so that each sublattice is uniquely identified by a basis $(\alpha_j = \sum a_{ij}\omega_i)$ in Hermite normal form, meaning that the integral matrix (a_{ij}) is upper triangular, has non-zero coefficients on the diagonal, and satisfies $a_{ij} < a_{ii}$ for $i \neq j$; see Chapter 4.7 of COHEN (1993) for details. Such a sublattice is an order if it contains all products

$$\alpha_j \alpha_{j'} = \sum_{i,i'} a_{ij} a_{i'j'} \omega_i \omega_{i'} = \sum_k \underbrace{\left(\sum_{i,i'} a_{ij} a_{i'j'} m_k^{ii'} \right)}_{b_k^{jj'}(a)} \omega_k$$

where the vector $m^{ii'}$ expresses $\omega_i \omega_{i'}$ on the basis (ω_k) ; this vector and the polynomial $b_k^{jj'}$ only depend on \mathcal{O} . Therefore, a is an order if and only if, for all j and j' , the preimage of the vector $b^{jj'}$ by the matrix a has integral coordinates; for sublattices of index $\det(a) = n$, this gives:

Proposition III.2.7. *All orders contained in \mathcal{O} with index n correspond to solutions of the polynomial system $(n \cdot a)^{-1} b^{jj'} = 0 \pmod{n^{2g} \mathbb{Z}^{2g}}$ in the coefficients of the matrix a .*

Unless there are 0 or $\Omega(n)$ such orders, this system is nonsingular and its solutions can be listed by a Gröbner basis algorithm in time polynomial in $\log n$ albeit exponential in g .

COMPUTING CLASS GROUPS

Fix an order \mathcal{O} and consider computing its Picard group; this requires a generating set of ideals for $\text{Pic}(\mathcal{O})$, an efficient ideal multiplication algorithm, and a way of finding a distinguished representative of the class of a prescribed ideal, which we call *reducing* an ideal. Under the generalized Riemann hypothesis (GRH), BACH (1990) solved the first problem:

Theorem III.2.8. *Assume the GRH and let \mathcal{O} be the ring of integers of a number field of discriminant Δ . The class group $\text{Pic}(\mathcal{O})$ is generated by prime ideals of norm at most $12 \log^2 |\Delta|$.*

Note that a less explicit, but more precise result of JAO, MILLER, and VENKATESAN (2009), which also assumes the GRH, implies that, for any $\varepsilon > 0$, the class group of any order \mathcal{O} is generated by prime ideals of norm less than $O(\log^{2+\varepsilon} |\Delta|)$, where $\Delta = \text{disc}(\mathcal{O})$.

Let \mathfrak{B} be the set of prime ideals with norm less than some bound B , and define

$$\sigma_{\mathcal{O}} : \begin{cases} \mathbb{Z}^{\mathfrak{B}} & \longrightarrow \text{Pic}(\mathcal{O}) \\ n & \longmapsto \prod_{\mathfrak{p} \in \mathfrak{B}} \mathfrak{p}^n \end{cases}$$

By the results above, when B is big enough, the map $\sigma_{\mathcal{O}}$ is surjective and therefore we have

$$\text{Pic}(\mathcal{O}) \simeq \mathbb{Z}^{\mathfrak{B}} / \Lambda_{\mathcal{O}}$$

where the lattice $\Lambda_{\mathcal{O}}$ is the kernel of $\sigma_{\mathcal{O}}$. Later, we will see how to compute the Picard group, that is, find a generating set of vectors for $\Lambda_{\mathcal{O}}$.

When the order \mathcal{O} lies in an imaginary quadratic field, its ideals can be represented as binary quadratic forms via the map

$$ax^2 + bxy + cy^2 \longmapsto a\mathbb{Z} + \frac{-b + \sqrt{b^2 - 4ac}}{2}\mathbb{Z}$$

where the right-hand side is a proper ideal of \mathcal{O} as soon as the integers a , b , and c are coprime and satisfy $b^2 - 4ac = \text{disc}(\mathcal{O})$. SCHÖNHAGE (1991) gave algorithms with *quasi-linear* runtime (that is, linear up to logarithmic factors) in $\log |\text{disc}(\mathcal{O})|$ for performing on such forms the operations which correspond to multiplying two ideals, and to reducing one into a canonical representative of its class.

When \mathcal{O} is an order of a general number field K , no such nice structure exists and a simpler approach must be used. Given a primitive element α , the field K can be represented as $\mathbb{Q}[x]/(\chi_{\alpha}(x))$, and its elements as rational vectors over the basis $(1, x, x^2, \dots, x^{\deg K - 1})$. Ideals \mathfrak{a} can then be expressed as \mathbb{Z} -modules, of which a generating set of cardinality $\deg(K)$ can be written as a matrix over a basis of the order to which they belong. As mentioned before, this matrix can be put in Hermite normal form to uniquely identify ideals.

Since there is no canonical set of ideal representatives for classes of the Picard group, it is difficult to identify ideal classes precisely. COHEN, DIAZ Y DIAZ, and OLIVIER (1997) demonstrated that this can nevertheless be done *to some extent*: the matrix that represents an ideal \mathfrak{a} can be reduced via the so-called LLL algorithm of LENSTRA, LENSTRA, and LOVÁSZ (1982), and the resulting matrix represents an ideal of the same class, but which is smaller. Such small ideals can be used as non-unique representatives of their class, and this permits one to perform most computations, notwithstanding some overhead.

III.3 Plain Complex Multiplication

We have seen that endomorphism rings of ordinary abelian varieties are isomorphic to orders in number fields, and have then considered their ideals from a computational standpoint. Let us now explain how these ideals can be seen as acting as isogenies.

This aspect of complex multiplication theory will be referred to as the *plain action*, as opposed to the *polarized action* to be discussed later. This section, does not assume that isogenies preserve any polarization structure of abelian varieties, and borrows many results of WATERHOUSE (1969).

FINITE FIELD SETTING

Let \mathcal{O} be an order isomorphic to the endomorphism ring of a simple ordinary abelian variety \mathcal{A} of dimension g defined over a finite field \mathbb{F}_q . We additionally consider an embedding $\iota : K \rightarrow \mathbb{Q} \otimes \text{End}(\mathcal{A})$ of the number field of \mathcal{O} ; its elements are then seen as endomorphisms of \mathcal{A} . An isogeny ϕ sends the variety \mathcal{A} to the variety $\mathcal{B} = \phi(\mathcal{A})$, and also maps an embedding ι for \mathcal{A} to an embedding for \mathcal{B} given as $\phi(\iota) = \frac{1}{\deg \phi} \phi \circ \iota \circ \widehat{\phi}$ where $\widehat{\phi}$ denotes the dual isogeny. In fact, we have:

Proposition III.3.1. *If ι is an embedding of K into $\mathbb{Q} \otimes \text{End}(\mathcal{A})$, all other embeddings ι' are of the form $\phi(\iota)$ for some endomorphism ϕ of \mathcal{A} .*

Let \mathcal{A} be such an abelian variety endowed with an embedding ι of \mathcal{O} into its endomorphism ring, let \mathfrak{a} be an invertible ideal of \mathcal{O} , and consider the isogeny $\phi_{\mathfrak{a}} : \mathcal{A} \rightarrow \mathcal{A} / \ker(\phi_{\mathfrak{a}})$ with kernel

$$\ker(\phi_{\mathfrak{a}}) = \bigcap_{\alpha \in \mathfrak{a}} \ker(\iota(\alpha)).$$

For instance, if \mathfrak{a} is a principal ideal (α) , then the kernel of $\phi_{\mathfrak{a}}$ is simply that of α ; therefore, $\phi_{\mathfrak{a}}$ is nothing but an endomorphism whose separable part coincides with that of α (recall that the totally inseparable part of an isogeny is not characterized by its kernel).

Now consider the composition of two such isogenies: let \mathcal{A} be an abelian variety, \mathfrak{a} be an invertible ideal of $\mathcal{O} = \iota^{-1}(\text{End}(\mathcal{A}))$, and denote the corresponding isogeny by $\phi_{\mathfrak{a}} : \mathcal{A} \rightarrow \mathcal{B}$; then, let \mathfrak{b} be an invertible element of $\phi(\iota)^{-1}(\text{End}(\mathcal{B}))$, and denote the corresponding isogeny by $\phi_{\mathfrak{b}} : \mathcal{B} \rightarrow \mathcal{C}$; in that situation, the isogeny $\phi_{\mathfrak{b}} \circ \phi_{\mathfrak{a}}$ corresponds canonically to $\phi_{\mathfrak{a}\mathfrak{b}} : \mathcal{A} \rightarrow \mathcal{C}$. In simple terms, composing isogenies corresponds to multiplying ideals.

As a consequence, there is a well-defined map

$$\mathfrak{a} \in \text{Pic}(\mathcal{O}) : \mathcal{A} \in \text{AV}_{\mathcal{O}}(k) \mapsto \phi_{\mathfrak{a}}(\mathcal{A}) \in \text{AV}(k)$$

where $AV(k)$ denotes the set of isomorphism classes of abelian varieties defined over k , and $AV_{\mathcal{O}}(k)$ the subset of such classes with endomorphism ring \mathcal{O} . Since the above is an isogeny, the complex multiplication is unchanged and we have $\mathbb{Q} \otimes \text{End}(\mathcal{A}) = \mathbb{Q} \otimes \text{End}(\phi_{\alpha}(\mathcal{A}))$; note that, for elliptic curves, $\text{End}(\phi_{\alpha}(\mathcal{A}))$ is actually always equal to $\text{End}(\mathcal{A})$ as Proposition III.3.7 will show, but in general we might only have $\text{End}(\mathcal{A}) \subset \text{End}(\phi_{\alpha}(\mathcal{A}))$.

COMPLEX ELLIPTIC TORUS

For elliptic curves, WATERHOUSE (1969) proved that the image of the map above is actually $AV_{\mathcal{O}}(k)$, and that the action of $\text{Pic}(\mathcal{O})$ on $AV_{\mathcal{O}}(k)$ this defines is transitive, which means that for any elliptic curve \mathcal{A} with endomorphism ring \mathcal{O} , the map $\alpha \mapsto \phi_{\alpha}(\mathcal{A})$ induces a bijection between $\text{Pic}(\mathcal{O})$ and $AV_{\mathcal{O}}(k)$. The specific approach that he used then enabled him to establish a similar result for (non-polarized) abelian varieties. Here, let us describe a more standard way of seeing this on elliptic curves, using complex tori.

In the elliptic case, the use of complex tori to obtain results over finite fields greatly exploits the following lifting theorem of DEURING (1941).

Theorem III.3.2. *Let α be an endomorphism of an elliptic curve \mathcal{A} defined over a finite field \mathbb{F}_p . There exists an endomorphism β of some abelian variety \mathcal{B} defined over a certain number field which, modulo some prime \mathfrak{p} above p of good reduction, reduces precisely to $\alpha \in \text{End}(\mathcal{A})$.*

In the case where $\text{End}(\mathcal{A}) = \mathbb{Z}[\alpha]$, the variety \mathcal{B} of the above theorem has $\mathbb{Z}[\beta]$ as endomorphism ring and reduction induces an isomorphism $\text{End}(\mathcal{B}) \simeq \text{End}(\mathcal{A})$, since we saw earlier that endomorphism rings of abelian varieties defined over number fields are mapped injectively into that of their good reductions at prime ideals. Endomorphism rings of ordinary elliptic curves are always of the form $\mathbb{Z}[\alpha]$, so in this case there always exist lifts with the same endomorphism ring.

Conversely, for the ordinary case, we need to reduce modulo primes totally split in \mathcal{O} :

Proposition III.3.3. *Let \mathcal{A} be an elliptic curve with endomorphism ring \mathcal{O} defined over a number field. Take an unramified prime \mathfrak{p} , and let $p = \mathfrak{p} \cap \mathbb{Z}$. Then:*

- if p splits completely in \mathcal{O} , then the reduction $\mathcal{A}_{\mathfrak{p}}$ is ordinary and defined over \mathbb{F}_p .
- if p is inert in \mathcal{O} , then the reduction $\mathcal{A}_{\mathfrak{p}}$ is supersingular and defined over \mathbb{F}_{p^2} .

Now, over the complex numbers, an elliptic curve with endomorphism ring \mathcal{O} always corresponds to a complex torus \mathbb{C}/\mathfrak{b} where \mathfrak{b} is a certain ideal of \mathcal{O} . The action of invertible ideals \mathfrak{a} of \mathcal{O} on $AV_{\mathcal{O}}(\mathbb{C})$ can then be seen as

$$\alpha : \mathbb{C}/\mathfrak{b} \in AV_{\mathcal{O}}(\mathbb{C}) \mapsto \mathbb{C}/(\alpha^{-1}\mathfrak{b}) \in AV_{\mathcal{O}}(\mathbb{C}).$$

This action is obviously transitive, and two ideals \mathfrak{a} and \mathfrak{a}' act identically if and only if they are homothetic, that is, if and only if they belong to the same class of $\text{Pic}(\mathcal{O})$. Therefore, this action factors through the Picard group into a faithful and transitive action of $\text{Pic}(\mathcal{O})$ on $\text{AV}_{\mathcal{O}}(\mathbb{C})$; modulo prime ideals \mathfrak{p} of norm p , it reduces to the action of $\text{Pic}(\mathcal{O})$ on $\text{AV}_{\mathcal{O}}(\mathbb{F}_p)$.

Theorem III.3.4. *Let \mathcal{O} be an imaginary quadratic order. For elliptic curves defined over a finite field k , the above defines a faithful and transitive action of $\text{Pic}(\mathcal{O})$ onto $\text{AV}_{\mathcal{O}}(k)$.*

We must finally mention that this action can also be seen on *invariants* of elliptic curves: if $\mathcal{B} \in \text{AV}_{\mathcal{O}}(\mathbb{C})$, its invariant $j(\mathcal{B})$ lies in the *ring class field* of \mathcal{O} , which is an abelian extension of $K = \mathbb{Q}(\mathcal{O})$ with Galois group $\text{Pic}(\mathcal{O})$. The action of $\text{Pic}(\mathcal{O})$ on $\text{AV}_{\mathcal{O}}(\mathbb{C})$ is then that of the Galois group via the Artin symbol.

GENERAL ABELIAN VARIETIES

The situation in higher dimension is far from being as nice as in the elliptic case. Certain properties nevertheless hold as they should, such as the following one of GIRAUD (1968).

Theorem III.3.5. *Let \mathcal{A} be a simple ordinary abelian variety defined over a finite field; if \mathfrak{a} is an invertible ideal of its endomorphism ring, the degree of the isogeny $\phi_{\mathfrak{a}}$ is the norm of \mathfrak{a} .*

The transitivity of the action of the Picard group, which would generalize the result on elliptic curves above, has only been shown to hold in the case that the endomorphism ring of \mathcal{A} is maximal by WATERHOUSE (1969); to prove this, he first argued that all invertible ideals are, in his terminology, *kernel ideals*, which implies the following.

Theorem III.3.6. *Let \mathcal{A} be a simple ordinary abelian variety defined over a finite field k , and assume that $\text{End}(\mathcal{A})$ is a maximal order \mathcal{O}_K ; then, for any invertible ideal \mathfrak{a} of \mathcal{O}_K :*

- the endomorphism ring of $\phi_{\mathfrak{a}}(\mathcal{A})$ is exactly that of \mathcal{A} .
- the induced action of $\text{Pic}(\mathcal{O}_K)$ on $\text{AV}_{\mathcal{O}_K}(k)$ is faithful and transitive.

The number of isomorphism classes of simple ordinary abelian varieties with endomorphism ring some maximal order \mathcal{O}_K can thus be estimated using the conjecture of SIEGEL (1935) proved by BRAUER (1947); as a direct consequence of Lemma III.2.5, we have

$$\text{disc}(\mathbb{Z}[\pi, \bar{\pi}]) < 2^{2g(g-1)} q^{g^2}$$

which gives, as g is fixed and q goes to infinity, the asymptotic behavior

$$\# \text{AV}_{\mathcal{O}_K}(\mathbb{F}_q) = \# \text{Pic}(\mathcal{O}_K) < q^{g^2/2+o(1)}.$$

EXPLICIT ACTION

In our application, we wish to use the above theory for maximal orders as well as non-maximal ones. Therefore, we rely on the following consequence of the results above, combined with the observation that, if the norm of an invertible ideal \mathfrak{a} is coprime to ℓ , since it is also the degree of the isogeny $\phi_{\mathfrak{a}}$, then the index $[\text{End}(\phi_{\mathfrak{a}} \cdot \mathcal{A}) : \text{End}(\cdot \mathcal{A})]$ cannot be divisible by ℓ . Note that we proved the contrapositive statement earlier.

Proposition III.3.7. *Let \mathcal{A} be a simple ordinary abelian variety defined over a finite field k , let π be its Frobenius endomorphism, let $K = \mathbb{Q}(\pi)$, and let $\mathcal{O} \subset K$ be its endomorphism ring.*

The invertible ideals of \mathcal{O} of norm coprime to the discriminant of $\mathbb{Z}[\pi, \bar{\pi}]$ act on $\text{AV}_{\mathcal{O}}(k)$ as isogenies of degree their norm, and this defines a faithful action of $\text{Pic}(\mathcal{O})$ on $\text{AV}_{\mathcal{O}}(k)$.

To make this proposition effective, we need to compute the isogeny $\phi_{\mathfrak{a}}$. Denote its degree by ℓ ; since $\ell = N(\mathfrak{a})$, we can start by enumerating all subgroups of cardinality ℓ of the full ℓ -torsion subgroup $\mathcal{A}[\ell]$. Recall that even when $\phi_{\mathfrak{a}}$ is rational, the points of its kernel need not be individually, but they are collectively invariant under the Galois action. Still, we need a practical way of telling $\phi_{\mathfrak{a}}$ apart from other isogenies of degree ℓ .

The improvements of ATKIN and ELKIES to the elliptic curve point counting method of SCHOOF (1985) exploit certain aspects of complex multiplication theory. In particular, they give a means to determine which specific isogeny of degree ℓ corresponds to $\phi_{\mathfrak{a}}$. It was also written as Stage 3 of the algorithm by GALBRAITH, HESS, and SMART (2002).

This result actually holds for general abelian varieties, which follows elementarily from the theory of Tate modules (from which most of the results that we stated above are derived); we therefore state it in its full generality.

Proposition III.3.8. *Let \mathcal{A} be a simple ordinary abelian variety defined over a finite field, \mathcal{O} its endomorphism ring and $\pi \in \mathcal{O}$ the element corresponding to its Frobenius endomorphism.*

Let \mathfrak{a} be an invertible prime ideal of \mathcal{O} , written as $\ell\mathcal{O} + u(\pi)\mathcal{O}$, where ℓ is its norm and u is an irreducible factor modulo ℓ of the characteristic polynomial χ_{π} of the primitive element π . Assume that ℓ is coprime to the discriminant of $\mathbb{Z}[\pi, \bar{\pi}]$.

Then, the characteristic polynomial of the Frobenius endomorphism acting on $\ker(\phi_{\mathfrak{a}})$ is u .

This proposition cannot be readily applied to non-prime ideals \mathfrak{a} , but we will explain later how this issue can be dealt with.

III.4 Polarized Complex Multiplication

In practical computations, abelian varieties are represented as Jacobian varieties of hyperelliptic curves or as theta-coordinates. Since both naturally work with principal polar-

izations, complex multiplication theory needs to be adapted to take this extra structure into account. Most of this theory originates from SHIMURA and TANIYAMA (1961).

As in the *plain* case, we start by considering complex multiplication fields before focusing on the specific endomorphism ring order and the action of its ideals.

REFLEX FIELDS AND MAPS

Recall that if \mathcal{A} is an ordinary abelian variety of dimension g , its complex multiplication field $K = \mathbb{Q} \otimes \text{End}(\mathcal{A})$ is a totally imaginary quadratic extension of a totally real number field K_+ of degree g , and that a *complex multiplication type* on K is a set of embeddings of K in \mathbb{C} satisfying $\Phi \sqcup \overline{\Phi} = \text{Hom}(K, \mathbb{C})$ where the union is disjoint.

Here, there is actually no need to involve \mathbb{C} , or even the algebraic numbers $\overline{\mathbb{Q}}$, since the image of any embedding of K is necessarily contained in its normal closure K^c . From now on, we therefore consider complex multiplication types given as sets of embeddings of K to its normal closure; this is equivalent and allows for a simpler exposition.

Definition III.4.1. *Let Φ be a type of K . The reflex field K^r is the fixed field of*

$$\{\sigma \in \text{Gal}(K^c, \mathbb{Q}) : \Phi = \Phi \circ \sigma\},$$

the automorphisms of K^c leaving Φ globally invariant. It admits a unique reflex type Φ^r which is the restriction of automorphisms of K^c whose inverses yield Φ when restricted to K , that is,

$$\{\phi \in \text{Aut}(K^c) : \phi|_{K^r} \in \Phi^r\} = \{\phi^{-1} \in \text{Aut}(K^c) : \phi|_K \in \Phi\}.$$

More generally, for any field extension K'/K , the type $\{\phi \in \text{Hom}(K', K'^c) : \phi|_K \in \Phi\}$ is called the *induced type* by Φ on K' . Types Φ which are not induced from a strictly smaller subfield are said to be *primitive*. Simple abelian varieties have primitive types, and in that case, we canonically have $K^{rr} = K$ and $\Phi^{rr} = \Phi$.

Define the *type trace* $\text{tr}_\Phi : x \in K \mapsto \sum_{\phi \in \Phi} \phi(x)$; its image actually generates the field K^r and this can be used as an equivalent definition for the reflex field; more importantly, define the *type norm*

$$N_\Phi : x \in K \mapsto \prod_{\phi \in \Phi} \phi(x) \in K^r$$

(it is elementary to verify that the images of both these maps are in K^r). There is also a *reflex type trace* tr_{Φ^r} and a *reflex type norm* $N_{\Phi^r} : K^r \rightarrow K$.

The latter is particularly important to us, as we will make great use of it via the map it induces on Picard groups: if \mathfrak{a} is an ideal of \mathcal{O}_{K^r} , there is a unique ideal of \mathcal{O}_K , which we write $N_{\Phi^r}(\mathfrak{a})$, such that

$$N_{\Phi^r}(\mathfrak{a})\mathcal{O}_{K^c} = \prod_{\phi \in \Phi} \phi(\mathfrak{a})\mathcal{O}_{K^c}$$

(see for instance Proposition 29 in Chapter II of SHIMURA (1998)). By the above, the map $N_{\Phi^r} : \mathfrak{J}(\mathcal{O}_{K^r}) \rightarrow \mathfrak{J}(\mathcal{O}_K)$ induces a morphism of Picard groups, which we also write similarly:

$$N_{\Phi^r} : \text{Pic}(\mathcal{O}_{K^r}) \rightarrow \text{Pic}(\mathcal{O}_K)$$

THE POLARIZED CLASS GROUP OF SHIMURA

Fix a primitive type Φ of a complex multiplication field K of degree $2g$, and denote the totally real subfield of K by K_+ .

Recall that a triple $(\Phi, \mathfrak{a}, \xi)$ yields the principally polarized complex torus $\mathbb{C}^g/\Phi(\mathfrak{a})$ with the polarization E_{Φ}^{ξ} ; Theorem III.1.5 explained that all tori arise in this way and gave necessary and sufficient conditions for two triples to yield isomorphic polarized varieties.

Following Section 14 of SHIMURA (1998), a group $\mathfrak{C}(\mathcal{O})$ can be constructed so as to naturally act on this set of triples representing isomorphism classes of principally polarized abelian varieties:

1. Let P be the group of pairs (\mathfrak{a}, ρ) where $\rho \in K_+$ is totally positive and \mathfrak{a} is a fractional ideal of \mathcal{O} satisfying $\mathfrak{a}\bar{\mathfrak{a}} = \rho\mathcal{O}$, endowed with component-wise multiplication.
2. Let I be the subgroup formed by the $(\mu\mathcal{O}, \mu\bar{\mu})$ for $\mu \in K^{\times}$.
3. Let $\mathfrak{C}(\mathcal{O})$ be the quotient group P/I .

As a consequence to Theorem III.1.5, we therefore have:

Corollary III.4.2. *For $\mathcal{O} = \mathcal{O}_K$, the group $\mathfrak{C}(\mathcal{O})$ acts faithfully and transitively on the set of isomorphism classes of principally polarized abelian varieties having complex multiplication by \mathcal{O} with type Φ . In particular, they have the same cardinality.*

It might be easier to understand the group $\mathfrak{C}(\mathcal{O})$ as part of the exact sequence

$$\mathbb{U}(K) \longrightarrow \mathbb{U}^+(K_+) \longrightarrow \mathfrak{C}(\mathcal{O}) \longrightarrow \text{Pic}(\mathcal{O}) \longrightarrow \text{Pic}^+(\mathcal{O}_+)$$

where the implied maps are, respectively, the norm of K/K_+ , the embedding $\rho \mapsto (\mathcal{O}, \rho)$, the projection $(\mathfrak{a}, \rho) \mapsto \mathfrak{a}$, and the map $\mathfrak{a} \mapsto \mathfrak{a}\bar{\mathfrak{a}} \cap K_+$; also, $\mathbb{U}^+(K_+)$ denotes the totally positive units of the totally real subfield K_+ , and $\text{Pic}^+(\mathcal{O}_+)$ denotes the quotient of the group of fractional ideals of $\mathcal{O} \cap K_+$ by those that admit a totally positive generator.

Intuitively, the class group $\text{Pic}(\mathcal{O})$ acts on the set of abelian varieties up to isomorphism, as proven by WATERHOUSE (1969) for $\mathcal{O} = \mathcal{O}_K$; the subgroup $\text{Pic}^+(\mathcal{O}_+)$ encodes the different ways an isogeny can alter polarizations, and the group $\mathbb{U}^+(K_+)/N_{K/K_+}(\mathbb{U}(K))$ corresponds to isomorphism classes of principal polarization.

For instance, in the case of dimension $g = 2$, when the totally-real subfield K_+ contains a unit of norm -1 , which exactly means that its fundamental unit is not totally positive, the quotient $\mathbb{U}^+(K_+)/N_{K/K_+}(\mathbb{U}(K))$ is trivial so we have a one-to-one map:

$$\mathfrak{C}(\mathcal{O}) \longrightarrow \ker(\text{Pic}(\mathcal{O}) \rightarrow \text{Pic}^+(\mathcal{O}_+))$$

Although the computation of the polarized class group $\mathfrak{C}(\mathcal{O})$ of Shimura is a much less classical topic than that of Picard groups, it is not more difficult; for instance, we note that similar groups have been studied from an algorithmic viewpoint by COHEN, DIAZ Y DIAZ, and OLIVIER (1998).

POLARIZED ACTION

There is a particular subgroup of the polarized class group of Shimura formed by elements arising as Galois actions. Here, we give a simplified exposition of this general theory and refer to Section 15 of SHIMURA (1998) for a more robust construction.

Let \mathcal{A} be a principally polarized abelian variety defined over \mathbb{C} with complex multiplication by the maximal order \mathcal{O}_K of a field K with type Φ . In fact, the abelian variety \mathcal{A} can be defined over the Hilbert class field \mathcal{H}_{K^r} which is the maximal abelian unramified extension of the reflex field, and in particular its *invariants* lie in that field; the action that we now describe can be seen as that of the Galois group of \mathcal{H}_{K^r} via the Artin symbol.

Theorem III.4.3. *Invertible ideals of K^r act on polarized tori with complex multiplication by \mathcal{O}_K with type Φ via*

$$\mathfrak{r} \in \mathfrak{I}(K^r) : \mathbb{C}^g / \Phi(\mathfrak{a}), E_\Phi^\xi \longmapsto \mathbb{C}^g / \Phi(N_{\Phi^r}(\mathfrak{r})^{-1}\mathfrak{a}), E_{\Phi}^{N_{K^r/\mathbb{Q}}(\mathfrak{r})^\xi};$$

an ideal \mathfrak{r} acts trivially when its reflex type norm ideal $N_{\Phi^r}(\mathfrak{r})$ is a principal ideal of \mathcal{O}_K generated by an invertible element $\mu \in K^\times$ which satisfies $\mu\bar{\mu} = N_{K^r/\mathbb{Q}}(\mathfrak{r})$.

Recall that the set of principally polarized abelian varieties with endomorphism ring \mathcal{O}_K is acted upon faithfully and transitively by the polarized class group $\mathfrak{C}(\mathcal{O}_K)$ of Shimura. The isogenies that arise via the reflex type norm (by theorem above) therefore act as the subgroup of $\mathfrak{C}(\mathcal{O}_K)$ formed by the elements

$$(N_\Phi(\mathfrak{r}), N_{K^r/\mathbb{Q}}(\mathfrak{r}))$$

where \mathfrak{r} ranges over ideals of \mathcal{O}_{K^r} . We emphasize that other elements of $\mathfrak{C}(\mathcal{O}_K)$ also act as isogenies, but that they might not be rational.

For instance, in dimension two, if $(\mathfrak{a}, \ell) \in \mathfrak{C}(\mathcal{O}_K)$, and ℓ totally splits as $\mathfrak{p}\bar{\mathfrak{p}}\mathfrak{q}\bar{\mathfrak{q}}$ in K , then the possible values for \mathfrak{a} are $\mathfrak{p}\mathfrak{q}$, $\mathfrak{p}\bar{\mathfrak{q}}$, and their respective conjugates; in that case, ℓ also splits

completely in K^r and the reflex type norm maps the prime factors of $\ell \mathcal{O}_{K^r}$ onto those four elements of \mathfrak{C} with norm ℓ^2 . In other cases, elements of $\mathfrak{C}(\mathcal{O}_K)$ of norm ℓ^2 might not be in the image of the reflex type norm.

REDUCTION TO FINITE FIELDS

We briefly review how the action that we have just defined transports to finite fields, in the case of simple ordinary abelian varieties of dimension two. For details, we refer to the work of GOREN (1997) and GOREN and LAUTER (2007).

We first consider a principally polarized abelian variety \mathcal{A}_p defined over a finite field of characteristic p ; given any embedding ι_p of \mathcal{O}_K into $\text{End}(\mathcal{A}_p)$, implying that \mathcal{A}_p has complex multiplication by \mathcal{O}_K , there exists an abelian variety \mathcal{A} defined over a number field and an embedding $\iota : \mathcal{O}_K \rightarrow \text{End}(\mathcal{A})$ which, at a certain prime, reduce to \mathcal{A}_p and ι_p respectively.

Conversely, if \mathcal{A} is a simple polarized abelian variety with complex multiplication by the maximal order of some field K , its invariants lie in the Hilbert class field \mathcal{H}_{K^r} which is the maximal abelian unramified extension of the reflex field. For almost all primes \mathfrak{p} of its field of definition, the abelian variety \mathcal{A} has good reduction modulo \mathfrak{p} .

Now, let p denote the rational prime below \mathfrak{p} , that is $p\mathbb{Z} = \mathfrak{p} \cap \mathbb{Z}$; when p splits completely in the complex multiplication field K this reduction is a simple ordinary abelian variety. Denote by \mathcal{A}_p the reduction of \mathcal{A} modulo \mathfrak{p} ; due to the injective map $\text{End}(\mathcal{A}) \rightarrow \text{End}(\mathcal{A}_p)$, we know that \mathcal{A}_p also has complex multiplication by \mathcal{O}_K . In that case, all elements of norm ℓ^s of the polarized class group of Shimura arise from the reflex type norm, and they give all isogenies of type $(\mathbb{Z}/\ell)^s$.

There is another splitting case for \mathfrak{p} which can result in the reduction \mathcal{A}_p being a simple ordinary abelian varieties: that where p is inert in K_+ but splits as $q\bar{q}$ in K and as $\tau\bar{\tau}'$ in K^r , where τ' has norm p^2 . In that case, reduction modulo a prime above τ or $\bar{\tau}$ also yield a simple ordinary abelian variety. However, the reduction of \mathcal{A} modulo τ' is a *superspecial variety*, that is, far from being ordinary.

If \mathcal{A} is a simple ordinary abelian variety of dimension $g = 2$ defined over a finite field k of sufficiently large characteristic, we will later exploit complex multiplication theory to predict the structure of its isogeny graph from that of its polarized class group of Shimura, or rather do the converse: predict the structure of the group $\mathfrak{C}(\mathcal{O})$ from that of the isogeny graph. For this, we have seen that we can always use isogenies of type $(\mathbb{Z}/\ell)^2$ for primes ℓ which split completely in the reflex field K^r .

However, we observe that elements of $\mathfrak{C}(\mathcal{O})$ of the form (\mathfrak{a}, ℓ) , where ℓ is a prime, which are not in the image of the reflex type norm, often also act as rational isogenies of type $(\mathbb{Z}/\ell)^2$, and we make use of these as well. In certain cases, this approach can be fully rigorous by

solely exploiting the action of $\mathfrak{C}(\mathcal{O})$ under the type norm, or that of certain elements (q, ℓ) for primes ℓ splitting in K as $q\bar{q}$. In other cases, this requires additional hypotheses, which we will then specify.

References

1935. Carl L. SIEGEL.
 “Über die Classenzahl quadratischer Zahlkörper”.
 In: *Acta Arithmetica* 1. Pages 83–86.
1941. Max DEURING.
 “Die Typen der Multiplikatorenringe elliptischer Funktionenkörper”.
 In: *Abhandlungen aus dem mathematischen Seminar der hamburgischen Universität* 14. Pages 197–272.
1945. André WEIL.
Sur les courbes algébriques et les variétés qui s'en déduisent. Volume 7.
 Actualités Scientifiques et Industrielles 1041.
 Publications de l'Institut de Mathématique de l'Université de Strasbourg.
1947. Richard BRAUER.
 “On the zeta-functions of algebraic number fields”.
 In: *American Journal of Mathematics* 69.2. Pages 243–250.
 DOI: 10.2307/2371849.
1961. Goro SHIMURA and Yutaka TANIYAMA.
Complex multiplication of abelian varieties and its applications to number theory.
 Volume 6. Publications of the Mathematical Society of Japan.
 The Mathematical Society of Japan.
1966. John TATE.
 “Endomorphisms of abelian varieties over finite fields”.
 In: *Inventiones mathematicae* 2.2. Pages 134–144. DOI: 10.1007/BF01404549.
1968. Jean GIRAUD.
 “Remarque sur une formule de Shimura-Taniyama”.
 In: *Inventiones Mathematicae* 5.3. Pages 231–236. DOI: 10.1007/BF01425552.
1969. William C. WATERHOUSE.
 “Abelian varieties over finite fields”.
 In: *Annales Scientifiques de l'École Normale Supérieure* 2.4. Pages 521–560.

1971. William C. WATERHOUSE and James S. MILNE.
“Abelian varieties over finite fields”. In: *1969 Number Theory Institute*.
Edited by Donald J. LEWIS. Volume 20.
Proceedings of Symposia in Pure Mathematics. American Mathematical Society.
Pages 53–64.
1982. Arjen K. LENSTRA, Hendrik W. LENSTRA, and László LOVÁSZ.
“Factoring polynomials with rational coefficients”.
In: *Mathematische Annalen* 261.4. Pages 515–534. DOI: 10.1007/BF01457454.
1985. René SCHOOF.
“Elliptic curves over finite fields and the computation of square roots mod p ”.
In: *Mathematics of Computation* 44.170. Pages 483–494.
DOI: 10.2307/2007968.
1986. Gary CORNELL and Joseph H. SILVERMAN (editors).
Arithmetic Geometry. Springer. ISBN: 3-540-96311-1.
1989. David A. COX.
Primes of the form $x^2 + ny^2$. John Wiley & Sons. ISBN: 0-471-19079-9.
1990. Eric BACH.
“Explicit bounds for primality testing and related problems”.
In: *Mathematics of Computation* 55.191. Pages 355–380.
DOI: 10.1090/S0025-5718-1990-1023756-8.
1991. Arnold SCHÖNHAGE.
“Fast reduction and composition of binary quadratic forms”.
In: *Symbolic and Algebraic Computation — ISSAC '91*. Edited by Stephen M. WATT.
Association for Computing Machinery. Pages 128–133.
DOI: 10.1145/120694.120711.
1993. Henri COHEN.
A course in computational algebraic number theory. Volume 138.
Graduate Texts in Mathematics. Springer. ISBN: 3-540-55640-0.
1997. Henri COHEN, Francisco DIAZ Y DIAZ, and Michel OLIVIER.
“Subexponential algorithms for class group and unit computations”.
In: *Journal of Symbolic Computation* 24.3–4. Special issue on computational algebra
and number theory: proceedings of the first MAGMA conference. Pages 433–441.
DOI: 10.1006/jscs.1996.0143.

1997. Eyal Z. GOREN.
“On certain reduction problems concerning abelian surfaces”.
In: *Manuscripta Mathematica* 94.1. Pages 33–43. DOI: 10.1007/BF02677837.
1998. Henri COHEN, Francisco DIAZ Y DIAZ, and Michel OLIVIER.
“Computation of relative quadratic class groups”.
In: *Algorithmic Number Theory — ANTS-III*. Edited by Joe P. BUHLER.
Volume 1423. Lecture Notes in Computer Science. Springer. Pages 433–440.
DOI: 10.1007/BFb0054882.
1998. Goro SHIMURA.
Abelian Varieties with Complex Multiplication and Modular Functions.
Princeton University Press. ISBN: 0-691-01656-9.
2002. Steven D. GALBRAITH, Florian HESS, and Nigel P. SMART.
“Extending the GHS Weil descent attack”.
In: *Advances in Cryptology — EUROCRYPT ’02*. Edited by Lars R. KNUDSEN.
Volume 2332. Lecture Notes in Computer Science. Springer. Pages 29–44.
DOI: 10.1007/3-540-46035-7_3.
2006. James S. MILNE.
Complex Multiplication.
URL: <http://www.jmilne.org/math/CourseNotes/cm.html>.
2007. Eyal Z. GOREN and Kristin E. LAUTER.
“Class invariants for quartic CM fields”.
In: *Annales de l’Institut Fourier* 57.2. Pages 457–480.
2009. David JAO, Stephen D. MILLER, and Ramarathnam VENKATESAN.
“Expander graphs based on GRH with an application to elliptic curve cryptography”.
In: *Journal of Number Theory* 129.6. Pages 1491–1504.
DOI: 10.1016/j.jnt.2008.11.006.
2010. Marco STRENG.
“Complex multiplication of abelian surfaces”. PhD thesis. Universiteit Leiden.
ISBN: 90-5335-291-0.
URL: <http://www.math.leidenuniv.nl/~streng/thesis.pdf>.

FOUR

Pairing-Friendly Varieties

IV.1 Cryptographic Requirements

The use of pairings enables many cryptographic protocols; as we have mentioned before, cryptography-grade pairings, that is, pairings which can be evaluated efficiently and are hard to invert, are only known to be defined on abelian varieties.

Here, we first review cryptographic requirements for pairing-based constructions, and then consider how abelian varieties satisfying these conditions can be generated.

GENERAL CONSIDERATIONS

Let \mathcal{A} be an abelian variety defined over a finite field \mathbb{F}_q and containing a cyclic subgroup of order r . The *embedding degree* $e(r)$, also written e when there is no ambiguity on the subgroup, is defined as the smallest integer such that the Weil pairing

$$\Upsilon_{\text{Weil}} : \mathcal{A}[r](\mathbb{F}_{q^e}) \times \mathcal{A}[r](\mathbb{F}_{q^e}) \longrightarrow \mu_r \subset \mathbb{F}_{q^e}^\times$$

is non-degenerate; extending a result of BALASUBRAMANIAN and KOBLITZ (1998), RUBIN and SILVERBERG (2009) proved that, if r does not divide $q - 1$ and the degree of the polarization of \mathcal{A} is coprime to r , then e divides the order of q modulo r .

Using this pairing for cryptographic purposes imposes the following:

1. It must be computationally infeasible to solve discrete logarithm problems in $\mathcal{A}[r]$.
2. It must be computationally infeasible to solve discrete logarithm problems in $\mu_r \subset \mathbb{F}_{q^e}^\times$.
3. It must be practical to compute over the field \mathbb{F}_{q^e} .

The last condition ensures that the algorithm of MILLER (1986) evaluates the Weil pairing efficiently. Note that many constructions do not directly use the Weil pairing but rather variants of it that enable evaluation speedups by small factors; however, from a variety generation point of view, this makes little difference: so long as field operations in \mathbb{F}_q can efficiently be computed, pairings with embedding degree e can be evaluated with more or less effort.

Later, it will be convenient to allow r to be a prime times a small cofactor; this does not invalidate the above: the security simply rests on the largest prime factor of r .

There are two big decisions to be made:

Binary or prime fields? Fields of characteristic two (also known as binary fields) are suited to efficient hardware implementations; on the other hand, software implementations work equally well with prime fields.

Supersingular or ordinary varieties? Supersingular varieties are easy to generate and readily have small embedding degrees; however, they are quite special and have an easy decisional Diffie–Hellman problem.

We choose to work with ordinary varieties defined over prime fields. Some authors argue that prime powers with exponent greater than one have density zero amongst prime powers, but here we justify this choice by its convenience and the fact that it avoids Weil-descent attacks altogether. Although attractive for the design of cryptographic protocol, the properties of supersingular curves can be seen unnecessarily special; they are mostly interesting over fields of small characteristic, and it is not so challenging to generate them.

To avoid wasting bits, we wish to balance the expected hardness the discrete logarithm problem in the abelian variety $\mathcal{A}(\mathbb{F}_q)$ and in the group $\mu_r \subset \mathbb{F}_q^\times$ as they are rendered equivalent by the pairing. When q is a prime power, HITT (2007) warned that μ_r might reside in a strict subfield of \mathbb{F}_q^\times , leading to faster attacks on its discrete logarithm problem. However, this problem does not arise when q is prime.

ASYMPTOTICS

Suppose \mathcal{A} is an ordinary abelian variety of dimension g defined over a prime field \mathbb{F}_q of which the discrete logarithm problem and pairing are considered for cryptographic use. By the Pohlig–Hellman reduction, it is sufficient to consider its largest prime subgroup \mathcal{H} ; we denote its order by r and its embedding degree by e . In order avoid attacks on high-genus varieties, we furthermore assume that $g = 1, 2$; this conveniently enables us to use the fast arithmetic of Jacobian varieties of hyperelliptic curves.

To measure the cryptographic efficiency, fix g and let q go to infinity: the complexity of additions in $\mathcal{A}(\mathbb{F}_q)$ is polynomial in $\log q$; disregarding the pairing, the discrete logarithm

problem in $\mathcal{A}(\mathbb{F}_q)$ achieves an expected security of $\frac{1}{2} \log_2 r$ bits. Hence, we introduce the quantity

$$\rho = \frac{g \log_2 q}{\log_2 r}$$

which, since $\#\mathcal{A}(\mathbb{F}_q) \sim q^g$, also indicates the proportion of bits used to represent points of $\mathcal{A}(\mathbb{F}_q)$ that actually contribute to the security of scheme: if $\rho \approx 1$ then nearly all of the variety is put to use; if $\rho \approx 2$ then only half of the bits are needed to identify points of \mathcal{A} .

Recall the best-known bounds on the complexity of solving discrete logarithm problems:

1. Discrete logarithm problems in $\mathcal{A}(\mathbb{F}_q)$ can be solved in $O(r^{1/2+o(1)} \log q)$.
2. Discrete logarithm problems in $\mathbb{F}_{q^e}^\times$ can be solved heuristically in $L_{1/3}^c(q^e)$.

To solve the first problem, in general, no better algorithm than generic ones is known, for which a lower bound of \sqrt{r} is proven; the other term in the complexity denotes the cost of operations in $\mathcal{A}(\mathbb{F}_q)$. Many variants of the number field sieve can be used to solve the second problem: the method of MATYUKHIN (2003) applies to prime fields, and that of JOUX and LERCIER (2006) is particularly adapted to extension fields such as here.

In the most effective case that $\rho \approx 1$, balancing the two complexities above requires

$$\frac{1}{2} g \log q \log \log q \approx c (e \log q)^{1/3} (\log e + \log \log q)^{2/3}$$

which implies $e \sim \left(\frac{g}{2c}\right)^3 \left(\frac{1}{3} \log q\right)^2 \log \log q$ and shows that the embedding degree should grow quadratically in the size of the base field; this is another reason to avoid supersingular varieties: since their embedding degrees are uniformly bounded as g is fixed (see below), they do not scale well to higher levels of security.

PRACTICE

To select the parameters g and e according to the level of security chosen (or equivalently the desired date until when the cryptosystem should withstand attacks), the cost of attacks on the discrete logarithm problems in both finite fields and abelian varieties must be carefully considered. Various agencies and organizations regularly publish updated tables listing parameter tuples for various security levels, such as ECRYPT II (2010) whose table was featured in the first chapter. Most agree that pairing-based cryptosystems aimed at being secure beyond 2030 should have a 256-bit r and a 3248-bit q^e ; as usual, more is better.

The practical cost of an attack can be estimated by using timings of previous attacks to calibrate the big-O (and possibly other) constants in the asymptotic complexity; this usually gives a fair estimation for larger instances. Here, we need to control both the hardness of




FIGURE 1. The abscissa bounds the security level of the discrete logarithm problem in \mathbb{F}_q^\times while the ordinate does the same in \mathcal{E}/\mathbb{F}_q . The diagonal represents the optimal case that these are balanced. The curves plot what elliptic curves achieve for selected values of e/ρ .

the discrete logarithm problem in the curve and the embedding field. Figure 1 does such a rough analysis for the parameters (ρ, e, q) of pairing-friendly curves. It shows, for instance, that 128 bits of security are best achieved by elliptic curves for which $e/\rho \approx 12$, with the most preferable choice of $\rho \approx 1$ implying that $e = 12$ and $q \approx 2^{256}$.

Before explaining how to generate elliptic curves and abelian varieties with the above properties, let us first say a bit more on supersingular varieties.

SUPERSINGULAR VARIETIES

While ordinary varieties are the generic case, supersingular varieties are the other extreme: recall that *supersingular abelian varieties* are defined as being isogenous to powers of supersingular elliptic curves (elliptic curves with zero p -rank) or, equivalently, as having Frobenius endomorphisms that satisfy $\pi^n = \pm q^{n/2}$ for some integer n .

Their cryptographic interest stems from the following result of GALBRAITH (2001).

Proposition IV.1.1. *The embedding degree of any subgroup of any g -dimensional supersingular abelian variety defined over a finite field is uniformly bounded by some quantity e_g .*

We have for instance $e_1 = 6$, $e_2 = 12$, $e_3 = 30$, $e_4 = 60$.

For certain types of base fields, the bound e_g can be lowered: the optimal bound for e_1 is 4 in characteristic two, 6 in characteristic three, 3 in higher characteristic, and 2 over prime fields with more than three elements.

An interesting feature of supersingular varieties is the existence of *distortion maps*, that is, non-rational endomorphisms. For ordinary varieties, we have seen that all endomorphisms defined over an algebraic closure are also defined over the base field, so their field of definition makes no difference. However, for supersingular varieties, there exist endomorphisms which do not commute with the Frobenius endomorphism.

Such *distortion maps* ψ are useful in cryptography because they send points of the rational r -torsion subgroup to points of $\mathcal{A}[r](\mathbb{F}_{q^r})$ which might not be rational. Then, the application

$$(P, Q) \in \mathcal{A}[r](\mathbb{F}_q)^2 \mapsto \Psi_{\text{Weil}}(\psi(P), Q) \in \mu_r$$

is a “self” pairing which is a very attractive object to build cryptographic primitives on, as its domain is the Cartesian product of two copies of the same cyclic group of order r , rather than the Cartesian product of two different ones.

On the other hand, this makes the decisional Diffie–Hellman problem easy, since for any triple of integers (a, b, c) and point P on \mathcal{A} , one can verify whether $c = ab$ given P, aP, bP, cP by checking whether

$$\Psi_{\text{Weil}}(\psi(aP), bP) = \Psi_{\text{Weil}}(\psi(P), cP);$$

from a security viewpoint, this can be seen as an undesirable property. Naturally, many protocols take advantage of that situation as well.

Since embedding degrees of supersingular curves are bounded, the base field size must grow more than linearly in the desired security level in order to avoid discrete logarithm attacks in \mathbb{F}_q^\times via the pairing; this lack of scalability is unpractical in the long term, and we now shift our focus to the ordinary case.

IV.2 Complex Multiplication Method

The problem of constructing ordinary abelian varieties defined over a finite field on which pairings are efficiently computable (meaning that the embedding degree is small) is an active topic of research.

This section describes the so-called *complex multiplication method* for generating ordinary abelian varieties with prescribed endomorphism rings; as a consequence, it also generates varieties whose Frobenius endomorphism have prescribed polynomials. Since the existence of a subgroup of order r with embedding degree e only depends on this polynomial, the next section will exploit this method to generate pairing-friendly varieties.

SCARCITY OF PAIRING-FRIENDLY VARIETIES

As we have argued before, abelian varieties of dimension $g = 1$ and 2 are the most suitable for cryptosystems which rely on the discrete logarithm problem. When no additional structure (such as a pairing) is required, abelian varieties need just have a near-prime group order, and are best generated by random search, which additionally reduces their likelihood of having undesirable special properties. For elliptic curves, such computations are classical, and for $g = 2$ it was recently demonstrated practical by GAUDRY and SCHOST (2010).

When, on top of a near-prime group order, one seeks a small embedding degree, this approach is not feasible anymore due to the scarcity of abelian varieties with the desired condition. More precisely, BALASUBRAMANIAN and KOBLITZ (1998) proved the following.

Theorem IV.2.1. *There are at most $M^{1/2+o(1)}$ isogeny classes of elliptic curves \mathcal{E}/\mathbb{F}_p with prime order and embedding degree less than $\log^2 p$, where p is a prime in $\{M/2, \dots, M\}$.*

Since there are roughly $M^{3/2}$ isogeny classes of elliptic curves defined over \mathbb{F}_p with $p \in \{M/2, \dots, M\}$, this is a pretty slim fraction of the total. LAUTER and SHANG (2010) recently gave a similar result for dimension-two abelian varieties:

Theorem IV.2.2. *Let H and K be positive integers, the number of pairs (p, N) where N is the order of a dimension-two abelian variety defined over \mathbb{F}_p with $p \in \{M/2, \dots, M\}$, such that $N = hr$ where $h \leq H$, r is prime and has embedding degree less than K is at most $M^{3/2+o(1)}HK^2$ for M large enough.*

Since there are roughly $M^{5/2}$ pairs (p, N) arising as orders of two-dimensional abelian varieties, this gives, similarly to the one-dimensional case, a probability of $p^{-1+o(1)}$ of finding a pairing-friendly abelian variety by random search over \mathbb{F}_p .

The theory of complex multiplication provides a method for generating such varieties efficiently. This involves two steps: we will first describe how varieties with prescribed endomorphism rings and prescribed fields of definition can be constructed using the so-called complex multiplication method, and we will then consider characterizing pairing-friendly varieties in terms of their endomorphism ring and base field.

CLASS POLYNOMIALS

Since abelian varieties of dimension three or more are not interesting for cryptography, we restrict to Jacobian varieties of hyperelliptic curves \mathcal{C} since all principally polarized abelian variety of dimension one or two are of this type. This allows to use invariants which uniquely identify the isomorphism class of such a variety and are expressed as rational functions of the coefficients of an equation for \mathcal{C} .

Fix a genus g and a family of invariants (I_i) that uniquely identify birationally equivalent classes of hyperelliptic curves. For instance, in dimension one, the *j-invariant*

$$\mathcal{C} : y^2 = x^3 + ax + b \mapsto j(\mathcal{C}) = \frac{2^8 3^3 a^3}{2^2 a^3 + 3^3 b^2}$$

(where we have assumed the characteristic to be different from 2 and 3) alone suffices. In higher dimension, as we have mentioned before, more invariants are necessary.

Let \mathcal{O} be the order of a complex multiplication field K of degree $2g$, that is, a totally imaginary quadratic extension of a totally real number field. SPALLEK (1994) first proposed to encode the information about all abelian varieties \mathcal{A} of dimension g defined over the complex numbers into the following polynomial

$$\mathcal{H}_i^{\mathcal{O}}(x) = \prod_{\{\mathcal{A} : \text{End } \mathcal{A} = \mathcal{O}\}} (x - I_i(\mathcal{A})),$$

where \mathcal{A} ranges over isomorphism classes of abelian varieties. In dimension one, they are usually called *Hilbert class polynomials* when \mathcal{O} is the maximal order of K , as their roots, the invariants of abelian varieties with endomorphism ring \mathcal{O} , generate the Hilbert class field of \mathcal{O} ; for non-maximal orders and in higher dimension, these lie in the ring class field of \mathcal{O} and the polynomials are simply known as *class polynomials*.

WENG (2001) later developed this theory and explained how these polynomials could be used to generate abelian varieties over finite fields with prescribed endomorphism ring, as we will soon explain. When there are two invariants or more (that is, for $g > 1$), these polynomials do not encode which root of $\mathcal{H}_1^{\mathcal{O}}$ corresponds to which root of $\mathcal{H}_i^{\mathcal{O}}$ for $i > 1$; in other words, the invariant tuples we are interested in are lost amongst tuples of unrelated invariants.

To address this issue, GAUDRY, HOUTMANN, KOHEL, RITZENTHALER, and WENG (2006) interpolated the values $I_i(\mathcal{A})$ at the $I_1(\mathcal{A})$: they defined

$$\mathcal{H}_i^{\prime\mathcal{O}}(x) = \sum_{\text{End } \mathcal{A} = \mathcal{O}} I_i(\mathcal{A}) \prod_{\substack{\text{End } \mathcal{B} = \mathcal{O} \\ \mathcal{B} \neq \mathcal{A}}} (x - I_1(\mathcal{B}))$$

for $i > 1$. This encodes exactly the information wanted.

REDUCTION TO PRIME FIELDS

Let \mathcal{A} be an ordinary abelian variety with complex multiplication by \mathcal{O} defined over some number field, and let \mathfrak{p} be a prime of degree one at which the reduction $\mathcal{A}_{\mathfrak{p}}$ of \mathcal{A} is itself an ordinary abelian variety defined over \mathbb{F}_p , where p is the rational prime below \mathfrak{p} . Since invariants are compatible with reduction, we have $I_i(\mathcal{A}_{\mathfrak{p}}) = I_i(\mathcal{A})_{\mathfrak{p}}$.

As the endomorphism ring of \mathcal{A} is mapped injectively into that of $\mathcal{A}_{\mathfrak{p}}$, we have $\mathcal{O} \subset (\text{End } \mathcal{A}_{\mathfrak{p}})$; when \mathcal{O} is the maximal order, equality must hold, and this is also the case for any order when $\mathcal{A}_{\mathfrak{p}}$ is an elliptic curve, due to the Deuring lifting theorem.

Consequently, an abelian variety with complex multiplication by \mathcal{O} defined over a finite field can be found using the following algorithm.

Algorithm IV.2.3.

INPUT: A prime p , and an order \mathcal{O} , either imaginary quadratic or maximal in a quartic complex multiplication field.

OUTPUT: An abelian variety $\mathcal{A} / \mathbb{F}_p$ with $\text{End } \mathcal{A} \simeq \mathcal{O}$.

1. Compute the class polynomials $\mathcal{H}_i^{\mathcal{O}}(x)$.
2. For each root I_1 of $\mathcal{H}_1^{\mathcal{O}}(x) \bmod p$:
3. For all $i > 1$, let $I_i = \mathcal{H}_i^{\mathcal{O}}(I_1) / \mathcal{H}_1^{\mathcal{O}}(I_1)$.
4. Use the method of MESTRE (1991) to compute a hyperelliptic curve whose Jacobian variety has invariants (I_i) .

Note that the output of this algorithm might be empty; for instance, when there are no abelian varieties with endomorphism ring \mathcal{O} defined over the field with p elements. In other cases, the number of curves returned might not be constant as \mathcal{O} is fixed and p varies. The conceptually simplest case is that where p completely splits in the ring class field of \mathcal{O} : then, the $\mathcal{H}_i^{\mathcal{O}}$ split into linear factors modulo p .

COMPUTATION OF CLASS POLYNOMIALS

Before making use of the method above, let us briefly describe the current methods available for computing class polynomials in dimension one and two.

Since the class polynomials $\mathcal{H}_i^{\mathcal{O}}$ are defined over the complex numbers and have good reduction to finite fields, there are, as with modular polynomials, two methods to compute them: a complex analytic method and one based on the Chinese remainder theorem.

The complex analytic version evaluates the invariants $I_i(\mathcal{A})$ for complex tori verifying $\text{End } \mathcal{A} \simeq \mathcal{O}$ to sufficient precision to identify the coefficients of the class polynomial; it requires tight bounds on the height of these coefficients. COUVEIGNES and HENOCQ

(2002) also proposed a p -adic version which proceeds similarly but uses the canonical lift of an abelian variety defined over a small extension of \mathbb{F}_p to transport the computation to \mathbb{Q}_p .

The Chinese remainder theorem version reconstructs the polynomials $\mathcal{H}_i^\mathcal{O} \in \mathbb{Q}[x]$ from their reduction to many small prime fields \mathbb{F}_p by enumerating the abelian varieties with endomorphism ring \mathcal{O} in each such field; typically, a first variety with complex multiplication $\mathbb{Q} \otimes \mathcal{O}$ is found by sheer luck (this requires computing the endomorphism ring of many random curves), and isogenies are then used to find a curve with endomorphism ring exactly \mathcal{O} and to enumerate all other such varieties.

When the dimension of \mathcal{O} is fixed, the complexity of all methods mainly depends on the order of the Picard group of \mathcal{O} , which dictates the number of roots of the class polynomials.

For elliptic curves, all methods have a quasi-linear runtime in the size of the output; see the careful analyses of ENGE (2009), BRÖKER (2008), and SUTHERLAND (2011). A practical advantage of the Chinese remainder theorem version is that it need not keep the full polynomials $\mathcal{H}_i^\mathcal{O} \in \mathbb{Q}(x)$ in memory: only their reductions modulo many primes are required; from these, $\mathcal{H}_i^\mathcal{O}$ can be directly reconstructed in the prime field where we seek an abelian variety with endomorphism ring \mathcal{O} . This is particularly useful as memory requirements are the current bottleneck of the other two methods.

In dimension two, WENG (2001) introduced the complex analytic method, CHAO, MATSUO, KAWASHIRO, and TSUJII (2000) the Chinese remainder theorem one, and GAUDRY, HOUTMANN, KOHEL, RITZENTHALER, and WENG (2006) a 2-adic method. All have since been improved by many researchers. Their respective speeds do not support a range of orders \mathcal{O} as wide as for elliptic curves, but quite a fair number of class polynomials have been computed and made available, for instance in the ECHIDNA (2008) package.

IV.3 Elliptic Curve Generation

Let us now explain how to apply the material of the previous section to generate pairing-friendly elliptic curves; very satisfying results can be obtained in this case. This is however not the case for higher-dimensional varieties, as the next section will discuss.

THE COCKS–PINCH METHOD

We have explained how an ordinary elliptic curve with prescribed order \mathcal{O} can be generated over a prescribed finite field \mathbb{F}_p when \mathcal{O} has small class number or, equivalently, small discriminant. We now consider which parameters p and \mathcal{O} should be chosen in order for the resulting curve to be pairing-friendly.

Let \mathcal{E} be an ordinary elliptic curve over the prime field with p elements; the characteristic polynomial $\chi_\pi(x)$ of its Frobenius polynomial is of the form $x^2 - tx + p$ where the integer t satisfies $|t| < 2\sqrt{p}$. Conversely, for each such nonzero integer, there exists an ordinary curve \mathcal{E}/\mathbb{F}_p with cardinality $p + 1 - t$ (we assume $p \neq 2, 3$). If r is the largest prime factor of $\#\mathcal{E}$, we require that its embedding degree be small, that is, $r \mid p^e - 1$ for some small integer e .

Additionally, for the complex multiplication method to be practical, there must exist orders of small discriminants in $\mathbb{Q}(\pi)$, that is, the squarefree part of $4p - t^2$ must be small.

Therefore, we require that:

1. p be a prime number.
2. t be a nonzero integer less than $2\sqrt{p}$ in absolute value.
3. r be a prime factor of $p + 1 - t$ such that $r \mid p^e - 1$ for a small e .
4. the squarefree part Δ of $t^2 - 4p$ be small in absolute value.

Since Δ and e need to be small, we first fix them: if an integer p can be derived as a function of Δ and e and it is not prime, we can always rerun the algorithm on a different input and hope that it takes a prime value after roughly $\log p$ trials; however, fixing p and deriving Δ or e would have little chances of producing small numbers.

Once Δ and e have been fixed, the method of COCKS and PINCH (2001) consists in rewriting the above set of conditions to the equivalent one:

$$\begin{cases} t^2 - 4p = v^2 \Delta \\ r \mid \Phi_e(t - 1) \\ r \mid v^2 \Delta - (t - 2)^2 \end{cases}$$

where Φ_e denotes the e^{th} cyclotomic polynomial; the second condition asserts that e is the smallest integer such that $r \mid p^e - 1$ but this stronger condition is not as important as the construction that it enables: since Φ_e is irreducible it yields a number field where to work. This gives the following algorithm.

Algorithm IV.3.1.

INPUT: A negative and a positive integer, Δ and e .

OUTPUT: A prime p and an order \mathcal{O} such that there exists a pairing-friendly elliptic curve with endomorphism ring \mathcal{O} over \mathbb{F}_p .

1. Choose a prime field \mathbb{F}_r containing $\sqrt{\Delta}$ and an e^{th} root of unity ζ_e .
2. Put $t = 1 + \zeta_e$ and $v = (t - 2)/\sqrt{\Delta}$ in \mathbb{F}_r .
3. Lift t and v to \mathbb{Z} and put $p = \frac{1}{4}(t^2 - v^2 \Delta)$.
4. Unless p is prime, go back to Step 1.
5. Output p and the order $\mathcal{O} = \mathbb{Z} + u^2 \mathcal{O}_{\mathbb{Q}(\sqrt{\Delta})}$ where u is any divisor of v .

Due to p being a sum of squares lifted from \mathbb{F}_p , the resulting elliptic has $\rho \approx 2$ on average.

FAMILIES OF PAIRING-FRIENDLY CURVES

Better ρ values are achieved by *families of curves* with a constant embedding degree e and discriminant Δ over fields \mathbb{F}_p for increasing primes p . Families of elliptic curves are given by tuples $(\Delta, e, p(x), t(x), r(x), v(x))$ where the last four parameters are polynomials in a formal variable x ; additionally to the conditions above, since p and r are expected to take prime values, they are required to be irreducible. The density of primes they produce can be estimated using Conjecture 1.4.3.

Before explaining how to adapt the method above to this context, let us give two explicit families; for a broader coverage, we refer to FREEMAN, SCOTT, and TESKE (2009).

MNT curves. Shortly before the constructive use of pairings in cryptography was uncovered, MIYAJI, NAKABAYASHI, and TAKANO (2001) warned that certain explicit families of curves had a small embedding degree and therefore were probably unsuitable for cryptographic use: they exhaustively studied the case that $\rho \approx 1$ and the cyclotomic polynomial Φ_e is quadratic, that is, $e \in \{3, 4, 6\}$; they gave an explicit description of all such ordinary elliptic curves; it was later noticed that they provide interesting pairing-friendly curves. For example, they proved that a curve features $\rho \approx 1$ and $e = 6$ if and only if $p(x) = 4x^2 + 1$ is prime and $t(x) = 1 \pm 2x$ for some integer x .

The Barreto–Naehrig family. BARRETO and NAEHRIG (2006) exhibited a family of ordinary elliptic curves with $\rho \approx 1$ and $e = 12$; as we have seen before, this is optimal to achieve the 128-bit security level using primes p of 256 bits. Their family has $\Delta = -3$ and

$$p(x) = 6^2 x^4 + 6^2 x^3 + 4 \cdot 6x^2 + 6x + 1 \quad t(x) = 1 + 6x^2$$

with $r(x) = p(x) + 1 - t(x)$.

The advantage of such families is that they fix the discriminant Δ and the asymptotic value of ρ , as we indeed have the limit $\rho \rightarrow \deg p / \deg r$ as $x \rightarrow \infty$. This enables the generation of good pairing-friendly curves with ρ bounded below 2 over large prime fields.

Deriving curves from such an explicit family is easy: for an expected p of n bits, take a random integer x having $n / \deg(p)$ bits, evaluate $p(x)$ and $r(x)$ and repeat the process until both $p(x)$ and $r(x)$ are primes; this requires an expected $O(n^2)$ trials.

THE BREZING–WENG METHOD

BREZING and WENG (2005) adapted the method of COCKS and PINCH (2001) to generate families of polynomials as defined above. Their construction follows the above except that the arithmetic is done over polynomial rings rather than over the integers.

Algorithm IV.3.2.

INPUT: A negative and a positive integer, Δ and e .

OUTPUT: A pairing-friendly family of curves given by $p(x)$, $t(x)$, and $r(x)$.

1. Choose an irreducible polynomial $r(x)$ with positive leading coefficient such that the field $\mathbb{Q}(x)/r$ contains $\sqrt{\Delta}$ and an e^{th} root of unity ζ_e .
2. Put $t = 1 + \zeta_e$ and $v = (t - 2)/\sqrt{\Delta}$, as elements of $\mathbb{Q}(x)/r$.
3. Lift t and v to $\mathbb{Z}[x]$ and put $p = \frac{1}{4}(t^2 - v^2\Delta)$.
4. Unless p is irreducible, go back to Step 1.
5. Output $p(x)$, $t(x)$, and $r(x)$.

Since the polynomial $p(x)$ is constructed as a sum of squares of lifts from $\mathbb{Q}(x)/r$, its degree is roughly twice that of r . However, when $\deg(r)$ is small, the degree of $p(x)$ can be much smaller and yield ρ values below 2; note that $\deg(p)$ being smaller is not a problem: curves defined over large prime fields can still be obtained by evaluating $p(x)$ at large integers x ; in fact, this is preferable since the slower increase of polynomials gives more flexibility.

LARGER CONDUCTORS

To conclude this section, we discuss the results of B. and SATOH (2008).

In this paper, we noted that the two methods described above only fix the complex multiplication field or, equivalently, the isogeny class, but not a specific endomorphism ring order \mathcal{O} which the complex multiplication method takes as input. Actually, our presentation of the Cocks–Pinch method above already showed that fact, since it stated that the order to be output could be of the form $\mathbb{Z} + u\mathcal{O}_{\mathbb{Q}(\pi)}$ for any divisor u of v , where $t^2 - 4p = v^2\Delta$ is the discriminant of the minimal order $\mathbb{Z}[\pi]$.

This means that, once parameters for a pairing-friendly curve or family have been computed, before applying the complex multiplication method and obtaining an actual elliptic curve, there is still some choice to be made on the specific endomorphism ring desired. In the Brezing–Weng method, since $v(x)$ is constructed as $(t - 2)/\sqrt{\Delta}$, its degree as polynomial is likely to be roughly that of r ; this typically gives a large (and predictable in size) pool of factors to choose from as the conductor of the endomorphism ring.

Therefore, pairing-friendly curves with non-maximal endomorphism rings \mathcal{O} can be generated as easily as maximal ones as long as \mathcal{O} is in the range of the complex multiplication method.

Denote by \mathcal{E}_1 and \mathcal{E}_u the elliptic curves with trace t and endomorphism rings respectively $\mathcal{O}_{\mathbb{Q}(\pi)}$ and $\mathcal{O} = \mathbb{Z} + u\mathcal{O}_{\mathbb{Q}(\pi)}$; there is an isogeny of degree u going from \mathcal{E}_1 to \mathcal{E}_u . Computing this isogeny takes essentially quadratic time in the largest prime factor of u , as we will see in subsequent chapters. Therefore, as it takes $u^{2+o(1)}\Delta$ time to generate the curve \mathcal{E}_u via class polynomials, using different values for u does not yield fundamentally new cryptosystems; it simply shows that a small range of conductors is readily available from pairing-friendly curve generation methods.

IV.4 Variety Generation

As a natural generalization of the problem of pairing-friendly elliptic curves generation, we now consider generating higher-dimensional pairing-friendly abelian varieties. We will first give general statements before mentioning state-of-the-art results.

MOTIVATION AND SETTING

From a mathematical viewpoint, it is only natural to switch our focus to abelian varieties when we feel the pool of interesting elliptic curves has been depleted, since abelian varieties with an efficient arithmetic (such as Jacobian varieties of genus-2 hyperelliptic curves) have equally effective and secure pairings; they can even be evaluated faster than that of elliptic curves as FREY and LANGE (2006) demonstrated.

Originally, abelian varieties were proposed for cryptographic use not only as alternatives to elliptic curves but also as a potential improvement: since the size of the group is g times the size of the base field, where g is the dimension, the parameters of a cryptosystem based on dimension-two abelian varieties need only be of half the size of an equivalently secure elliptic cryptosystem; in addition, the smaller base field can possibly be exploited to yield a faster (or at least competitive) arithmetic to that of elliptic curves.

Although abelian varieties readily provide a good framework for cryptosystems based on the discrete logarithm problem only, other factors need to be taken into account for pairing-based cryptography. Before explaining how the situation degrades for ordinary varieties, let us recall that two-dimensional supersingular abelian varieties have an embedding degree of at most 12 and ρ values which can be close to 1; they are currently the only kind of two-dimensional abelian varieties suitable for cryptographic use.

All known constructions of ordinary pairing-friendly varieties of dimension two have large ρ values: we will see that none has $\rho \leq 2$, and that ρ values close to 2 are only achieved by special constructions; generic constructions feature $\rho \geq 4$, at the time of this writing.

It therefore appears as if genus-two constructions had a lot of room for improvement.

COMPLEX MULTIPLICATION METHOD

We have seen that the computation of class polynomials, although harder for abelian varieties of dimension two than for elliptic curves, can be done (and has been done) for a limited number of orders \mathcal{O} , all of which are ring of integers of quartic complex multiplication fields with relatively small discriminant.

Therefore, it is even more important to fix \mathcal{O} as a first step of any construction than it was with elliptic curves. We distinguish two types of constructions:

1. Generic constructions, which take an arbitrary maximal quartic complex multiplication order as input, and output generic pairing-friendly abelian varieties.
2. Specific constructions, which focus on varieties of a particular form (usually implying that \mathcal{O} is fixed too) and exploit explicit results due to this form.

Here, by “generic” we mean that the former methods output varieties with no particular properties other than those required; in particular, the varieties are usually absolutely simple and ordinary. This is to be compared to the varieties obtained by the latter method which are typically simple but not absolutely simple.

GENERIC CONSTRUCTIONS

The first construction of ordinary pairing-friendly abelian varieties of dimension $g > 1$ with cryptographic size are due to FREEMAN (2007). It can be considered a genus-two analog to the Cocks–Pinch method, and proceeds by solving explicit equations which arise by writing the characteristic polynomial of the Frobenius endomorphism in terms of parameters for the desired complex multiplication field. The abelian varieties it generates have a typical ρ value of 8.

Later, FREEMAN, STEVENHAGEN, and STRENG (2008) provided a cleaner framework for constructing pairing-friendly ordinary abelian varieties of dimension two by using more of the theory of complex multiplication.

Let π be the Frobenius endomorphism of a simple ordinary abelian variety \mathcal{A} over a finite field. Their idea was to write the condition that \mathcal{A} has a subgroup of order r with embedding degree e as

$$\begin{cases} r \mid N_{\mathbb{Q}(\pi)/\mathbb{Q}}(\pi - 1) \\ r \mid \Phi_e(\pi\bar{\pi}) \end{cases} .$$

Now let Φ be a type on the complex multiplication field K , and denote by Φ^r and K^r their respective reflexes. The key observation is that, if r is a prime congruent to one modulo

e that splits completely in K , and if

$$\prod_{\phi \in \Phi^r} (\xi \bmod \mathfrak{r}_\phi) = 1 \quad \text{and} \quad \prod_{\phi \in \Phi^r} (\xi \bmod \overline{\mathfrak{r}_\phi}) = \zeta_e$$

where ζ_e is an e^{th} root of unity and $\prod_{\phi \in \Phi^r} \mathfrak{r}_\phi \overline{\mathfrak{r}_\phi}$ denotes the factorization of r in K^r , then the type norm $\pi = N_{\Phi^r}(\xi)$ of ξ is a q -Weil number (that is, a root of a q -Weil polynomial) satisfying the conditions above asserting that it represents an ordinary pairing-friendly abelian variety.

Computationally, numbers ξ can be constructed from their reductions modulo the prime factors of r so as to satisfy the above requirement; after sufficiently many trials, the integer $q = N_{K^r/\mathbb{Q}}(\xi)$ is expected to be prime, and when it is additionally unramified in K and π generates K , this yields, by Honda–Tate theory, an isogeny class of ordinary pairing-friendly abelian varieties with complex multiplication by K .

The method above still produces varieties whose embedding degree is 8 or more, but FREEMAN (2008) soon adapted it to generate families of pairing-friendly varieties similarly to the Brezing–Weng method for elliptic curves. He applies it to find many families with ρ less than 8, and a particular one with an asymptotic ρ value of 4 for $e = 5$.

SPECIFIC CONSTRUCTIONS

To improve on the ρ values obtained by constructions applicable to arbitrary complex multiplication fields, one way is to consider abelian varieties \mathcal{A} of a particular form and exploit explicit results regarding this form as much as possible. Usually, \mathcal{A} is taken as the Jacobian variety $\text{Jac}(\mathcal{C})$ of a hyperelliptic curve \mathcal{C} of genus two with a particular shape of Weierstrass polynomial.

For instance, consider curves \mathcal{C} of the form $y^2 = x^5 + ax$ for some number $a \in \mathbb{F}_p$, where p is a prime congruent to one modulo eight; in that situation, the associated Jacobian variety $\text{Jac}(\mathcal{C})$ is ordinary and simple, and KAWAZOE and TAKAHASHI (2008) exploited explicit formulas for the characteristic polynomial of the Frobenius endomorphism in terms of a and p to obtain an analog of the Cocks–Pinch method for that specific type of curves. They obtained a ρ value of 3 with the embedding degree $e = 24$.

The varieties they constructed are not absolutely simple: over an extension containing fourth roots of e , they split as products of two elliptic curves. FREEMAN and SATOH (2011) studied such varieties from a much more general perspective: from an elliptic curve \mathcal{E} which is pairing-friendly over some extension of its base field, they explain how to derive a simple ordinary pairing-friendly abelian variety which becomes isomorphic to a power of \mathcal{E} over some extension of the same base field. As an application, they construct families of such

abelian varieties with $\rho \approx 2.22$ and $e = 27$, which are to date the best known ordinary pairing-friendly varieties of dimension two.

References

1986. VICTOR S. MILLER.
Short programs for functions on curves. Unpublished.
URL: <http://crypto.stanford.edu/miller/>.
1991. JEAN-FRANÇOIS MESTRE.
“Construction de courbes de genre 2 à partir de leurs modules”.
In: *Effective methods in algebraic geometry — MEGA '90*.
Edited by Teo MORA and Carlo TRAVERSO. Volume 94. Progress in Mathematics.
Birkhäuser. Pages 313–334.
1994. ANNE-MONIKA SPALLEK.
“Kurven vom Geschlecht 2 und ihre Anwendung in Public-Key-Kryptosystemen”.
PhD thesis. Universität Duisburg-Essen.
URL: <http://www.iem.uni-due.de/zahlentheorie/AES-KG2.pdf>.
1998. RAMACHANDRAN BALASUBRAMANIAN and NEAL KOBLITZ.
“The improbability that an elliptic curve has subexponential discrete log problem under the Menezes–Okamoto–Vanstone algorithm”.
In: *Journal of Cryptology* 11.2. Pages 141–145. DOI: 10.1007/s001459900040.
2000. JINHUI CHAO, KAZUTO MATSUO, HIROTO KAWASHIRO, and SHIGEO TSUJII.
“Construction of hyperelliptic curves with CM and its application to cryptosystems”.
In: *Advances in Cryptology — ASIACRYPT '00*. Edited by Tatsuaki OKAMOTO.
Volume 1976. Lecture Notes in Computer Science. Springer. Pages 259–273.
DOI: 10.1007/3-540-44448-3_20.
2001. CLIFFORD COCKS and RICHARD G. E. PINCH.
Identity-based cryptosystems based on the Weil pairing. Unpublished manuscript.
2001. STEVEN D. GALBRAITH.
“Supersingular curves in cryptography”.
In: *Advances in Cryptology — ASIACRYPT '01*. Edited by Colin BOYD.
Volume 2248. Lecture Notes in Computer Science. Springer. Pages 495–513.
DOI: 10.1007/3-540-45682-1_29.

2001. Atsuko MIYAJI, Masaki NAKABAYASHI, and Shunzou TAKANO.
“New explicit conditions of elliptic curve traces for FR-reductions”.
In: *IEICE Transactions on Fundamentals of Electronics* E84-A.5. Pages 1234–1243.
2001. Annegret WENG.
“Konstruktion kryptographisch geeigneter Kurven mit komplexer Multiplikation”.
PhD thesis. Universität Duisburg-Essen. URL: <http://www.iem.uni-due.de/zahlentheorie/preprints/wengthesis.pdf>.
2002. Jean-Marc COUVEIGNES and Thierry HENOCQ.
“Action of modular correspondences around CM-points”.
In: *Algorithmic Number Theory — ANTS-V*.
Edited by Claus FIEKER and David R. KOHEL. Volume 2369.
Lecture Notes in Computer Science. Springer. Pages 234–243.
DOI: 10.1007/3-540-45455-1_19.
2003. Dmitry V. MATYUKHIN.
“On the asymptotic complexity of computing discrete logarithms in the field $\text{GF}(p)$ ”.
In: *Diskretnaya Matematika* 15.1. Pages 28–49.
2005. Friederike BREZING and Annegret WENG.
“Elliptic curves suitable for pairing based cryptography”.
In: *Design, Codes and Cryptography* 37.1. Pages 133–141.
DOI: 10.1007/s10623-004-3808-4.
2006. Paulo S. L. M. BARRETO and Michael NAEHRIG.
“Pairing-friendly elliptic curves of prime order”.
In: *Selected Areas in Cryptography — SAC ’05*.
Edited by Bart PRENEEL and Stafford TAVARES. Volume 3897.
Lecture Notes in Computer Science. Springer. Pages 319–331.
DOI: 10.1007/11693383_22.
2006. Gerhard FREY and Tanja LANGE.
“Fast bilinear maps from the Tate-Lichtenbaum pairing on hyperelliptic curves”.
In: *Algorithmic Number Theory — ANTS-VII*.
Edited by Florian HESS, Sebastian PAULI, and Michael POHST. Volume 4076.
Lecture Notes in Computer Science. Springer. Pages 466–479.
DOI: 10.1007/11792086_33.
2006. Pierrick GAUDRY, Thomas HOUTMANN, David R. KOHEL,
Christophe RITZENTHALER, and Annegret WENG.
“The 2-adic CM method for genus 2 curves with application to cryptography”.
In: *Advances in Cryptology — ASIACRYPT ’06*.

- Edited by Xuejia LAI and Kefei CHEN. Volume 4284.
Lecture Notes in Computer Science. Springer. Pages 114–129.
DOI: 10.1007/11935230_8.
2006. Antoine JOUX and Reynald LERCIER.
“The function field sieve in the medium prime case”.
In: *Advances in Cryptology — EUROCRYPT ’06*. Edited by Serge VAUDENAY.
Volume 4004. Lecture Notes in Computer Science. Springer. Pages 254–270.
DOI: 10.1007/11761679_16.
2007. David M. FREEMAN.
“Constructing pairing-friendly genus 2 curves with ordinary Jacobians”.
In: *Pairing-Based Cryptography — PAIRING ’07*. Edited by Tsuyoshi TAKAGI,
Tatsuaki OKAMOTO, Eiji OKAMOTO, and Takeshi OKAMOTO. Volume 4575.
Lecture Notes in Computer Science. Springer. Pages 152–176.
DOI: 10.1007/978-3-540-73489-5_9.
2007. Laura HITT.
“On the minimal embedding field”. In: *Pairing-Based Cryptography — PAIRING ’07*.
Edited by Tsuyoshi TAKAGI, Tatsuaki OKAMOTO, Eiji OKAMOTO, and
Takeshi OKAMOTO. Volume 4575. Lecture Notes in Computer Science. Springer.
Pages 294–301. DOI: 10.1007/978-3-540-73489-5_16.
2008. Gaetan BISSON and Takakazu SATOH.
“More discriminants with the Brezing-Weng method”.
In: *Progress in Cryptology — INDOCRYPT ’08*.
Edited by Dipanwita R. CHOWDHURY, Vincent RIJMEN, and Abhijit DAS.
Volume 5365. Lecture Notes in Computer Science. Springer. Pages 389–399.
DOI: 10.1007/978-3-540-89754-5_30.
2008. Reinier BRÖKER.
“A p -adic algorithm to compute the Hilbert class polynomial”.
In: *Mathematics of Computation* 77.264. Pages 2417–2435.
DOI: 10.1090/S0025-5718-08-02091-7.
2008. David M. FREEMAN.
“A generalized Brezing-Weng method for constructing pairing-friendly ordinary
abelian varieties”. In: *Pairing-Based Cryptography — PAIRING ’08*.
Edited by Steven D. GALBRAITH and Kenny G. PATERSON. Volume 5209.
Lecture Notes in Computer Science. Springer. Pages 146–163.
DOI: 10.1007/978-3-540-85538-5_11.

2008. David M. FREEMAN, Peter STEVENHAGEN, and Marco STRENG.
“Abelian varieties with prescribed embedding degree”.
In: *Algorithmic Number Theory — ANTS-VIII*.
Edited by Alfred J. van der POORTEN and Andreas STEIN. Volume 5011.
Lecture Notes in Computer Science. Springer. Pages 60–73.
DOI: 10.1007/978-3-540-79456-1_3.
2008. Mitsuru KAWAZOE and Tetsuya TAKAHASHI.
“Pairing-friendly hyperelliptic curves with ordinary Jacobians of type $y^2 = x^5 + ax$ ”.
In: *Pairing-Based Cryptography — PAIRING ’08*.
Edited by Steven D. GALBRAITH and Kenny G. PATERSON. Volume 5209.
Lecture Notes in Computer Science. Springer. Pages 164–177.
DOI: 10.1007/978-3-540-85538-5_12.
2008. David R. KOHEL.
ECHIDNA: Databases for elliptic curves and higher dimensional analogues.
URL: <http://echidna.maths.usyd.edu.au/>.
2009. Andreas ENGE.
“The complexity of class polynomial computation via floating point approximations”.
In: *Mathematics of Computation* 78.266. Pages 1089–1107.
DOI: 10.1090/S0025-5718-08-02200-X.
2009. David FREEMAN, Michael SCOTT, and Edlyn TESKE.
“A taxonomy of pairing-friendly elliptic curves”.
In: *Journal of Cryptology* 23.2. Pages 224–280.
DOI: 10.1007/s00145-009-9048-z.
2009. Karl RUBIN and Alice SILVERBERG.
“Using abelian varieties to improve pairing-based cryptography”.
In: *Journal of Cryptology* 22.3. Pages 330–364.
DOI: 10.1007/s00145-008-9022-1.
2010. European Network of Excellence in Cryptology II.
Yearly report on algorithms and key sizes. Edited by Nigel P. SMART.
URL: <http://www.ecrypt.eu.org/documents/D.SPA.13.pdf>.
2010. Pierrick GAUDRY and Éric SCHOST.
Genus 2 point counting over prime fields. HAL-INRIA: 00542650.
2010. Kristin LAUTER and Ning SHANG.
Generating pairing-friendly parameters for the CM construction of genus 2 curves over prime fields. IACR ePrint: 2010/529.

2011. David M. FREEMAN and Takakazu SATOH.
“Constructing pairing-friendly hyperelliptic curves using Weil restriction”.
In: *Journal of Number Theory* 131.5. Pages 959–983.
DOI: 10.1016/j.jnt.2010.06.003.
2011. Andrew V. SUTHERLAND.
“Computing Hilbert class polynomials with the Chinese remainder theorem”.
In: *Mathematics of Computation* 80. Pages 501–538.
DOI: 10.1090/S0025-5718-2010-02373-7.

COMPUTATION OF
ENDOMORPHISM RINGS

Exponential Methods

The last chapter was concerned with constructing abelian varieties with prescribed endomorphism rings and we now turn to the inverse problem: that of computing the endomorphism ring of a prescribed variety. Our contribution is covered by the next three chapters; here, we review prior state-of-the-art algorithms, all of which have a worst case running time exponential in the size of the base field.

All sections but the last solely consider ordinary varieties, and our complexity analyses concern a fixed dimension g and a cardinality q of the base field going to infinity.

If \mathcal{A} is an ordinary abelian variety with complex multiplication field K , an isomorphism $\mathbb{Q}(\pi) \simeq K$ between the field of fractions of $\text{End}(\mathcal{A})$ and K will be understood throughout this chapter; this identifies endomorphism rings uniquely as orders of K .

v.1 Isogeny Volcanoes

Let us first describe the structure of the connected component of the isogeny graph containing a prescribed simple ordinary abelian variety over a finite field; we will emphasize *vertical* isogenies and their role in the algorithm of KOHEL (1996) for computing endomorphism rings in the dimension-one case.

VERTICAL ISOGENIES

Following FOUQUET and MORAIN (2002), we say that an isogeny is *horizontal* when its domain and codomain have isomorphic endomorphism rings, and that it is *vertical* otherwise; we first focus on the latter kind, in the context of computing endomorphism rings. Later, we will use horizontal isogenies, via complex multiplication theory, as the key to our subexponential-time algorithm for computing endomorphism rings.

To put to light the relationship between endomorphism rings and vertical isogenies, we use an observation of KOHEL:

Lemma v.1.1. *Let $\phi : \mathcal{A} \rightarrow \mathcal{B}$ be an isogeny of type $(\mathbb{Z}/\ell)^g$ between ordinary abelian varieties defined over a finite field. The order $\text{End}(\mathcal{B})$ is bounded below by $\mathbb{Z} + \ell \text{End}(\mathcal{A})$.*

Indeed, since ϕ splits multiplication by ℓ , we have $\ell \text{End}(\mathcal{A}) \subset \text{End}(\mathcal{B})$, and since the latter is an order it must also contain \mathbb{Z} . Note that applying this lemma to the dual isogeny $\widehat{\phi}$ gives a bound on $\text{End}(\mathcal{B})$ from above. To encompass both bounds, we generalize the inclusion index to the following distance on the lattice of orders.

Definition v.1.2. *For any two orders \mathcal{O} and \mathcal{O}' of the same field, define the order distance $\text{dist}(\mathcal{O}, \mathcal{O}')$ as $[\mathcal{O} : \mathcal{O} \cap \mathcal{O}'] + [\mathcal{O}' : \mathcal{O} \cap \mathcal{O}']$.*

Corollary v.1.3. *Let $\phi : \mathcal{A} \rightarrow \mathcal{B}$ be an isogeny of type $(\mathbb{Z}/\ell)^g$ between ordinary abelian varieties defined over a finite field. The distance $\text{dist}(\text{End } \mathcal{A}, \text{End } \mathcal{B})$ is divisible by ℓ^{4g-2} .*

This follows from the lemma, since $\mathbb{Z} + \ell \mathcal{O}$ has index ℓ^{2g-1} in \mathcal{O} , for any order \mathcal{O} . By exploiting the symmetry of the lattice of orders, the distance could even be proven to divide ℓ^{2g-1} . However, this simple result is sufficient for us; as a consequence, there can only be finitely many vertical isogenies of a given type leaving from any given variety \mathcal{A} since:

- only finitely many orders of K are endomorphism rings, that is, contain $\mathbb{Z}[\pi, \bar{\pi}]$;
- therefore there are only finitely many possible degrees for vertical isogenies;
- since $\mathcal{A}[\ell] = (\mathbb{Z}/\ell)^{2g}$ there are finitely many suitable subgroups.

Recall the results of TATE (1966) and WATERHOUSE (1969):

Theorem v.1.4. *Isogeny classes of abelian varieties defined over a finite field are identified by the characteristic polynomial of their Frobenius endomorphism. Endomorphism rings of ordinary varieties \mathcal{A} are exactly those orders of the complex multiplication field K that contain $\mathbb{Z}[\pi, \bar{\pi}]$.*

This shows that the structure of vertical isogenies is quite rigid: the possible degrees are fixed per isogeny class by the index of the minimal order $\mathbb{Z}[\pi, \bar{\pi}]$ in the maximal one of K . Worse, they can be as large as $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ which Lemma III.2.5 showed can only be bounded by $q^{g^2/2+o(1)}$ where q is the cardinality of the base field and g the dimension of the variety. This does not give much flexibility for working with vertical isogenies, and can make it quite costly to evaluate them.

On the other hand, we will later argue that horizontal isogenies are convenient to work with, as there are infinitely many with domain any given variety.




FIGURE 2. Structure of the graph of vertical isogenies and of the lattice of orders.

GLOBAL STRUCTURE

As a consequence to the above, the structure of the vertical-isogeny graph can be described as resembling that of the lattice of orders which contain the minimal order $\mathbb{Z}[\pi, \bar{\pi}]$.

Corollary v.1.5. *Let G be the graph whose vertices are classes of varieties with Frobenius endomorphism π , up to horizontal isogenies, with edges the vertical isogenies of type $(\mathbb{Z}/\ell)^g$. Similarly, let H be the graph whose vertices are the orders containing $\mathbb{Z}[\pi, \bar{\pi}]$, with edges between two orders $\mathcal{O} \subsetneq \mathcal{O}'$ when there is no \mathcal{O}'' satisfying $\mathcal{O} \subset \mathcal{O}'' \subset \mathcal{O}'$.*

The map $(\mathcal{A} \rightarrow \mathcal{A}') \in G \mapsto (\text{End } \mathcal{A} \rightarrow \text{End } \mathcal{A}') \in H$ is bijective on the vertices, and splits edges into sequences of at most $2g - 1$ edges.

Figure 2 is probably worth all the above words: it depicts the graph of vertical isogenies (the big circles denote horizontal isogenies classes) to the left, and the corresponding lattice of orders to the right. In fact, this is a simple case, similar to the situation in dimension one: each order above $\mathbb{Z}[\pi, \bar{\pi}]$ is uniquely identified by its index in \mathcal{O}_K , and vertical isogenies are in bijection with edges of the lattice of orders, that is, they do not jump orders.

Computing the endomorphism ring of a variety is therefore equivalent to determining its location up to horizontal isogenies in the isogeny graph.

To see how big this structure can be, consider the typical case of ordinary varieties of dimension $g = 2$ defined over the prime field with p elements. From the conditions on p -Weil polynomials, we deduce that there must be $p^{3/2+o(1)}$ isogeny classes. Since there are $p^{3+o(1)}$

isomorphism classes of curves, each isogeny class contains, on average, $p^{3/2+o(1)}$ isomorphism classes.

From now on, we will assume that the discriminant of $\mathbb{Z}[\pi, \bar{\pi}]$ (and therefore its index in the maximal order) has been factored, so that we can make use of the various algorithms for lattices of orders developed earlier.

From a cryptanalysis viewpoint, if \mathcal{A} is an abelian variety of which the discrete logarithm problem is to be used in a cryptographic scheme, and \mathcal{A}' is a variety in the same isogeny class for which this problem is known to be weak, it should be ensured that it is infeasible to compute any isogeny $\mathcal{A} \rightarrow \mathcal{A}'$.

By the theory of complex multiplication, there are many horizontal isogenies of small degree going from any abelian variety \mathcal{A} to others with the same endomorphism ring; therefore, horizontal isogeny classes can be “walked around” quite easily. Note, however, that finding an explicit path from a prescribed variety to another might be a difficult task when the horizontal isogeny class is big, since only generic methods are available.

However, when \mathcal{A} and \mathcal{A}' have different endomorphism rings, denoting by ℓ the largest prime factor of $\text{dist}(\text{End } \mathcal{A}, \text{End } \mathcal{A}')$, any isogeny chain going from \mathcal{A} to \mathcal{A}' must contain an isogeny of degree ℓ . Since current isogeny-computing algorithms require exponential time in $\log(\ell)$, this bounds below the time needed to transport the discrete logarithm problem.

LOCAL STRUCTURE IN DIMENSION ONE

FOUQUET and MORAIN (2002) gave a metaphorical interpretation of the work of KOHEL (1996) on the structure of the graph of isogenies of type \mathbb{Z}/ℓ , for a fixed prime ℓ , between ordinary elliptic curves defined over a finite field. In dimension one, a number of properties which we sum up in the proposition below indeed give graphs of degree- ℓ isogenies a distinctive *volcano* look.

Recall that the complex multiplication fields of ordinary elliptic curves are exactly the imaginary quadratic number fields; orders of such fields are of the form $\mathbb{Z} + f\mathcal{O}_K$ where f is the index in the maximal order \mathcal{O}_K .

The following rephrases Proposition 23 of KOHEL (1996) and, for short, refers to isomorphism classes of elliptic curves as *curves* and to the valuation at a fixed prime ℓ of the conductor of their endomorphism ring as their *depth*.

Proposition v.1.6. *Consider the graph of isogenies of prime degree ℓ between isomorphism classes of elliptic curves defined over a finite field with complex multiplication by the imaginary quadratic field $K = \mathbb{Q}(\sqrt{D})$ of discriminant D , and denote by v the valuation of $[\mathcal{O}_K : \mathbb{Z}[\pi]]$ at ℓ . The following exhaustively describes all edges of this graph.*




FIGURE 3. Typical volcano structure in dimension one when the discriminant is a square modulo ℓ (the prime degree of isogenies); here, in the case that $\ell = 3$.

1. From a curve at depth $u > 0$, there is one isogeny going up to a curve at depth $u - 1$.
2. From a curve at depth $u < v$, there are ℓ isogenies going down to ℓ curves at depth $u + 1$, unless $u = 0$ in which case there are $\ell - 1$, ℓ , or $\ell + 1$ when D is respectively a square, zero, or a non-square modulo ℓ .
3. From a curve at depth 0, there are two isogenies going to curves at depth 0 when D is a square modulo ℓ , and one when D is divisible by ℓ .

Again, Figure 3 is likely worth the above words: it displays one connected component of the graph that we discussed; note that by the proposition and results of complex multiplication theory, all connected components of this graph are isomorphic.

The algorithm of KOHEL (1996) computes the endomorphism ring of an ordinary curve \mathcal{E} by determining the valuation of its conductor at certain primes ℓ , for which it probes the location of \mathcal{E} in the graph structure that we have just described.

This relies on the vertical structure of this graph being that of trees rooted on the (possibly degenerated) cycle of curves with locally maximal endomorphism rings. Note that this structure is lost in higher dimension, as we will later see.

KOHEL'S ALGORITHM

KOHEL (1996) introduced many ideas and results related to the computation of endomorphism rings of elliptic curves over finite fields. Let us just describe two of them which lead to his deterministic algorithm for computing the endomorphism ring $\text{End}(\mathcal{E})$ of an ordinary elliptic curve \mathcal{E} over \mathbb{F}_q in time $q^{1/3+\epsilon}$.

The first idea directly exploits the structure of the volcano discussed above: the valuation of the conductor of $\text{End}(\mathcal{E})$ at some prime ℓ can be found by determining on which level of the graph of degree- ℓ isogenies \mathcal{E} lies. To this extent, compute three chains of degree- ℓ isogenies starting from \mathcal{E} ; one chain necessarily descends to levels of higher depth, and eventually hits a leaf, that is, a curve with depth v from which no isogeny leaves but the dual of that with which we arrived. The set of leaves is called the *floor of rationality*; its curves only have one rational subgroup of order ℓ (whence the expression), and the ℓ remaining subgroups define isogenies over an extension of the base field. This gives the following algorithm.

Algorithm v.1.7.

INPUT: An ordinary elliptic curve $\mathcal{E} / \mathbb{F}_q$.

OUTPUT: The conductor of its endomorphism ring.

1. Count the points of \mathcal{E} and deduce its complex multiplication field K .
2. For each prime ℓ dividing $[\mathcal{O}_K : \mathbb{Z}[\pi]]$:
3. Compute three curves ℓ -isogenous to \mathcal{E} .
4. Keep walking a non-backward chain of ℓ -isogenies from each.
5. Denote by u_ℓ the length of the chain that ends first.
6. Return $[\mathcal{O}_K : \mathbb{Z}[\pi]] / \prod \ell^{u_\ell}$.

By *non-backward*, we mean that we avoid duals of isogenies already computed. The first step uses polynomially many operations in $\log(q)$. Each isogeny can then be computed in time $\ell^{2+o(1)}$ using the independent improvement of DEWAGHE (1995) and KOHEL (1996), Section 2.4, on the formulas of VÉLU (1971); this process will be detailed in the next chapter. Since ℓ can be as large as \sqrt{q} , the overall complexity is only bounded by $q^{1+o(1)}$.

The second idea then comes to the rescue by trading off vertical isogenies for horizontal ones; the concise presentation below is largely inspired by a talk of KOHEL (2010).

Recall from complex multiplication theory that there are exactly $\#\text{Pic}(\mathcal{O})$ curves with endomorphism ring \mathcal{O} , and that they form a connected component of the horizontal isogeny graph. Therefore, when ℓ is large, the value of u_ℓ can be tested by comparing the class number of the order \mathcal{O} with valuation u_ℓ to the number of curves in the horizontal isogeny component. Formally, this gives the algorithm below.

Algorithm v.1.8.

INPUT: An ordinary elliptic curve $\mathcal{E} / \mathbb{F}_q$.

OUTPUT: The conductor of its endomorphism ring.

1. Count the points of \mathcal{E} and deduce its complex multiplication field K .
2. For each prime-power factor $\ell^v \leq q^{1/6}$ of $[\mathcal{O}_K : \mathbb{Z}[\pi]]$:
3. Apply the former algorithm.
4. For each prime-power factor $\ell^v > q^{1/6}$ of $[\mathcal{O}_K : \mathbb{Z}[\pi]]$:
5. Count the number of curves having horizontal isogenies to \mathcal{E} .
6. Determine the order whose class group matches.

The horizontal isogenies of Step 5 can be constructed as chains of isogenies of degree up to $12 \log^2 \Delta$, where $\Delta = \text{disc}(K)$, by Theorem III.2.8. In addition, not the whole horizontal isogeny class need be enumerated: it is sufficient to compute enough of it so as to rule out other orders with smaller class number.

KOHEL (1996) concludes that:

Theorem v.1.9 (GRH). *For any real number $\varepsilon > 0$, endomorphism rings of ordinary elliptic curves can be computed in deterministic time $q^{1/3+\varepsilon}$.*

v.2 Higher Dimension

Before presenting methods for computing endomorphism rings in arbitrary dimension, let us describe more of the structure of isogeny graphs. We start by formalizing the localization of the lattice of orders at a prime; this isolates a subgraph of the corresponding isogeny graph structure. Then, we move on to describing those specific aspects of the isogeny graph which differ from dimension one to dimension two and more.

LOCAL ORDER STRUCTURE

Fix a number field K and consider the lattice L of orders \mathcal{O} that contain a prescribed *minimal order* \mathfrak{m} , which will be $\mathbb{Z}[\pi, \bar{\pi}]$ in our applications. The index of any such order in the *maximal order* $\mathfrak{M} = \mathcal{O}_K$ then obviously divides $w = [\mathfrak{M} : \mathfrak{m}]$.

Now if ℓ is a prime factor of w , we can *localize* the lattice of orders via the map

$$\begin{aligned} L &\longrightarrow L_\ell = \{\mathcal{O} \in L : [\mathfrak{M} : \mathcal{O}] \mid \ell^\infty\} \\ \mathcal{O} &\longmapsto \mathcal{O}_\ell = \mathcal{O} + \mathfrak{m}_\ell \end{aligned}$$

where \mathfrak{m}_ℓ is the smallest order of the codomain, that is, the smallest order with index in \mathfrak{M} a power of ℓ . This projects \mathcal{O} onto the maximal order \mathfrak{M} locally at all primes but ℓ , thus

isolating the local information at ℓ . This information can be recombined by the isomorphism

$$\begin{array}{ccc} L & \simeq & \prod_{\ell} L_{\ell} \\ \mathcal{O} & \mapsto & \mathcal{O} + \mathfrak{m}_{\ell} \\ \bigcap_{\ell} \mathcal{O}_{\ell} & \longleftarrow & (\mathcal{O}_{\ell}) \end{array}$$

which can be evaluated in time polynomial in $\log |\Delta|$, where $\Delta = \text{disc}(\mathfrak{m})$, using the classical algorithms from Chapter III.

For us, K is the complex multiplication field of an ordinary abelian variety \mathcal{A} over a finite field, and $\mathfrak{m} = \mathbb{Z}[\pi, \bar{\pi}]$. We will often say that we consider the endomorphism ring of \mathcal{A} *locally at* ℓ to mean that we consider the localization $\text{End}(\mathcal{A})_{\ell}$; by the above, knowing $\text{End}(\mathcal{A})_{\ell}$ for each prime factor ℓ of w is sufficient to identify $\text{End}(\mathcal{A})$ exactly.

Since isogenies of degree ℓ^n can only move endomorphism rings by distances that are powers of ℓ , the endomorphism rings of abelian varieties in a connected degree- ℓ vertical isogeny class are injectively projected to L_{ℓ} . Therefore, for the purpose of identifying the endomorphism ring using vertical isogenies, those of degree ℓ can be considered one prime ℓ at a time.

In dimension one, K is an imaginary quadratic field in which orders are uniquely identified by their index in \mathcal{O}_K . The local lattice L_{ℓ} is then the chain

$$\mathcal{O}_K \supset \mathbb{Z} + \ell \mathcal{O}_K \supset \mathbb{Z} + \ell^2 \mathcal{O}_K \supset \cdots \supset \mathbb{Z} + \ell^{\text{val}_{\ell} w} \mathcal{O}_K.$$

Consequently, it is really worthwhile for many algorithms dealing with imaginary quadratic orders to work locally, so as to benefit from this simple structure: this usually yields conceptually simpler algorithms. However, from dimension two on, the local lattice is not a tree but a general lattice itself, so it makes no conceptual difference whether one works locally or not, although it is advantageous for performance reasons.

LOCAL ISOGENY STRUCTURE

Let us now briefly present the major differences between the degree- ℓ isogeny graph structure for elliptic curves and for higher-dimensional abelian varieties. Part of the last chapter will be devoted to giving details and results of computations on these aspects.

Let \mathcal{O} be the endomorphism ring of an ordinary elliptic curve defined over a finite field. The distinctive look of its isogeny volcanoes stems from two properties:

- Rational primes ℓ split in at most two ideals of \mathcal{O} .
- Ideals of prime norm dividing the index $[\mathcal{O}_K : \mathcal{O}]$ are not invertible in \mathcal{O} .




FIGURE 4. Graph of isogenies of type $(\mathbb{Z}/3)^2$ containing the Jacobian variety of the curve $y^2 = 8x^6 + 3x^5 + 7x^4 + 5x^3 + 12x^2 + 5x + 5$ over the field with 23 elements. Red circle varieties have maximal endomorphism ring, and blue triangle ones have index 9 in the maximal order.

By the theory of complex multiplication, the first property implies that elliptic curves with locally maximal endomorphism ring lie on (possibly degenerated) circles: the *crater* of the volcano. When the prime ℓ is inert, these circles degenerate into single vertices; when it splits as $\mathfrak{p}\bar{\mathfrak{p}}$, then each circle has length the order of \mathfrak{p} in $\text{Pic}(\mathcal{O})$. The second property implies that there are no horizontal isogenies of prime degree between elliptic curves with locally non-maximal endomorphism rings, that is, other than at the crater of the volcano.

Both properties are lost in higher dimension; indeed, if \mathcal{O} is an order in a complex multiplication field of degree $2g$ for $g > 1$, then:

- Rational primes ℓ can split in up to $2g$ ideals of \mathcal{O} .
- Ideals of prime norm not coprime to the index $[\mathcal{O}_K : \mathcal{O}]$ may be invertible in \mathcal{O} .

This implies that horizontal degree- ℓ isogenies between varieties with locally maximal endomorphism rings now have a slightly more involved structure than a cycle, and that they might also exit other than at the top of the volcano. Both features are displayed on Figure 4.

We shall say more on this topic in the last chapter. In the meantime, the reader should not be misled into thinking that all higher dimensional local isogeny graphs portray the same structure as this specific one; however, this gives an idea why generalizing the algorithm of KOHEL (1996) for computing endomorphism rings cannot be done straightforwardly.

PAIRINGS VIA TORSION STRUCTURE

Although endomorphism rings of higher-dimensional abelian varieties cannot be determined by their vertical isogeny graph structure alone, other structures can be involved in a hope to adapt the method of KOHEL (1996) to this generalized setting.

IONICA and JOUX (2010) recently gave a method for finding subgroups of order ℓ in ordinary elliptic curves over finite fields that are kernels of ascending or horizontal isogenies, meaning that they lead to curves with larger (or equal) endomorphism rings. Essentially, they exploit the relationship between the rational ℓ^∞ -torsion subgroup structure of an elliptic curve and the valuation at ℓ of its endomorphism ring. To obtain the subgroup structure, they rely on pairing computations and on the algorithm of COUVEIGNES (2009) for computing the torsion, which will be discussed in the next section.

This permits one to navigate in the volcano not just blindly relying on the tree structure of vertical isogenies, but with “some sense of orientation.” Since we believe their method should be, to some extent, applicable to higher dimension varieties, we briefly present it.

A theorem of LENSTRA (1996) states the following.

Theorem v.2.1. *Let π be the Frobenius endomorphism of an ordinary elliptic curve \mathcal{E} defined over \mathbb{F}_q and put $\mathcal{O} = \text{End}(\mathcal{E})$. The \mathcal{O} -modules $\mathcal{E}(\mathbb{F}_{q^n})$ and $\mathcal{O}/(\pi^n - 1)$ are isomorphic.*

Since \mathcal{O} is a quadratic order, the group structure of the elliptic curve $\mathcal{E}(\mathbb{F}_q)$ is therefore of the form $\mathbb{Z}/N_0 \times \mathbb{Z}/N_1$ where $N_0 \mid N_1$. In particular, its ℓ^∞ -torsion subgroup structure is of the form $\mathbb{Z}/\ell^{\alpha_0} \times \mathbb{Z}/\ell^{\alpha_1}$ and IONICA and JOUX (2010) derive explicit formulas for the integers α_0 and α_1 which show that they only depend on the valuation at ℓ of the conductor of $\text{End}(\mathcal{E})$.

To give an example of the specific way in which α_0 and α_1 are affected by vertical isogenies, let us reproduce Proposition 4 of IONICA and JOUX (2010).

Proposition v.2.2. *Let \mathcal{E} be an elliptic curve of rational ℓ^∞ -torsion subgroup $\mathbb{Z}/\ell^{\alpha_0} \times \mathbb{Z}/\ell^{\alpha_1}$ with $\alpha_1 > \alpha_0$. If P is a point of order ℓ^{α_0} , then the isogeny with kernel generated by $\ell^{\alpha_0-1}P$ is descending.*

The computational ingredients are simple: we will present a torsion-finding method in the next chapter, as it is needed in our own algorithms, and pairing evaluations are used to test relations between the order of ℓ^∞ -torsion points. Therefore, we believe this method has a good potential of being generalized to higher dimension, at least partially.

Since it is based on vertical isogenies, this approach is probably not best suited to computing endomorphism rings, as we argue below. Nevertheless, it has other interesting applications which can be found in the original article.

LIMITATIONS OF VERTICAL ISOGENIES

Isogeny computation is currently a topic in active development for abelian varieties of dimension $g > 1$. The state-of-the-art algorithm of COSSET and ROBERT (2011) can only compute isogenies of type $(\mathbb{Z}/\ell)^g$ and requires the prime ℓ to be reasonably small: although the asymptotic complexity is polynomial in ℓ and exponential in g , the constant factors and exponents are such that only a much more restricted range of isogenies can be computed than in dimension one.

We have argued before that vertical isogenies have constrained degrees; if certain isogenies are not within reach of known isogeny-computing methods, then their local vertical isogeny volcano is simply not computable. After our review of previous methods, the next chapter will present an algorithm which addresses this issue by relying on horizontal isogenies, whose degrees can be chosen with much more flexibility.

Another obstruction arises from the type of the isogenies that can be evaluated: consider a chain of orders

$$\mathcal{O}_K = \mathcal{O}_1 \supset \cdots \supset \mathcal{O}_v = \mathbb{Z}[\pi, \bar{\pi}]$$

where each order is contained in the following one with prime order ℓ ; this is a simple case, as we have mentioned that there are others for $g > 1$, but it suffices to make our point.

WATERHOUSE (1969) proved the existence of abelian varieties \mathcal{A}_i with endomorphism ring \mathcal{O}_i and TATE (1966) proved that there exist isogenies between all of the \mathcal{A}_i ; the degrees of these isogenies are necessarily powers of ℓ .

However, the kernels of these isogenies need not be of type $(\mathbb{Z}/\ell)^g$ or a combination of such subgroups. In other words, in dimension g , we might “skip” up to $g - 1$ orders when computing vertical isogenies. In the case that $g = 2$, for instance, starting from an abelian variety with endomorphism ring \mathcal{O}_0 and following isogenies of type $(\mathbb{Z}/\ell)^2$ we might only reach abelian varieties with endomorphism ring \mathcal{O}_i for i even, and fail to reach those with i odd. The last chapter will give several examples illustrating this.

v.3 General Methods

Two methods were previously known for computing endomorphism rings of general abelian varieties \mathcal{A} defined over finite fields. Both test whether elements α of the complex multiplication field $K = \mathbb{Q}(\pi)$ correspond to endomorphisms of \mathcal{A} ; doing so for generating sets of orders permits one to eventually recover the full endomorphism ring.

To find whether $\alpha \in \text{End}(\mathcal{A})$, the method of EISENTRÄGER and LAUTER (2009) tests if some easy-to-evaluate multiple $n\alpha$ kills the full n torsion subgroup of \mathcal{A} .

Recently, WAGNER (2009) designed a new method which can loosely be understood as a Chinese remainder theorem variant of the latter: to determine whether $\alpha \in \text{End}(\mathcal{A})$, it tries to interpolate the potential corresponding endomorphism over small torsion subgroups.

EVALUATING ENDOMORPHISMS

Let \mathcal{A} be a simple ordinary principally polarized abelian variety defined over the field with q elements. Since the endomorphism ring of \mathcal{A} always contains the order $\mathbb{Z}[\pi, \bar{\pi}]$, let us explain how the action on \mathcal{A} of an endomorphism α of this subring can be evaluated.

Evaluating the Frobenius endomorphism π is straightforward: it suffices to put the coordinates of a point to the q^{th} power, which, using a double-and-add approach, only requires a number of base field multiplications that is polynomial in $\log(q)$. On the other hand, evaluating the Verschiebung endomorphism $\bar{\pi} = q/\pi$ is more involved but can be avoided, unless p divides the conductor of $\mathbb{Z}[\pi, \bar{\pi}]$ where p is the prime of which q is a power.

Since $K = \mathbb{Q}(\pi)$, any element $\alpha \in K$ can be written as a rational polynomial in the Frobenius endomorphism π : if $2g$ is the degree of the field, there exist an integer n and integers α_i , for $i \in \{0, \dots, 2g - 1\}$ such that

$$\alpha = \frac{1}{n} \sum_i \alpha_i \pi^i.$$

Computing α therefore amounts to evaluating the Frobenius endomorphism, scalar multiplications, endomorphism compositions, and one division. Note that division by n is easily computed on torsion subgroups of \mathcal{A} of order coprime to n : simply multiply by the inverse of n modulo the order. Subgroups of order not coprime to n will soon be addressed.

In the following, α will always be an algebraic integer of K , and we assume this from now on. Put $w' = [\mathcal{O}_K : \mathbb{Z}[\pi]]$; as a group, $\frac{1}{w'} \mathbb{Z}[\pi]$ then contains \mathcal{O}_K . Therefore, α can be written in the form above for some integer n dividing w' . And this is in fact always the case when the above expression is reduced, meaning that $\gcd(\alpha, n) = 1$.

Recall from Lemma III.2.5 that $w'/w = [\mathbb{Z}[\pi, \bar{\pi}] : \mathbb{Z}[\pi]] = q^{g(g-1)/2}$ where $w = [\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ as before. As a consequence, the prime factors of the denominator n are those of w (that is, the degrees of vertical isogenies) plus, possibly, q .

THE EISENTRÄGER–LAUTER METHOD

We now present the method of EISENTRÄGER and LAUTER (2009); it was first targeted at testing whether endomorphism rings of abelian varieties over finite fields are maximal, but it applies to other orders as well. It relies on Corollary 9 which reads as follows.

Proposition v.3.1. *Let \mathcal{A} be an abelian variety defined over an algebraically closed field. If α is an endomorphism of \mathcal{A} and n is coprime to the ambient characteristic, then $\mathcal{A}[n] \subset \ker(\alpha)$ if and only if $\alpha/n \in \text{End}(\mathcal{A})$, that is, if there exists an endomorphism β such that $\alpha = n \circ \beta = \beta \circ n$.*

In other words, the endomorphism corresponding to the algebraic integer α kills the full n -torsion subgroup if and only if α/n belongs to the endomorphism ring.

As we have mentioned before, when \mathcal{A} is ordinary, assuming the base field to be algebraically closed does not affect the endomorphism ring; it only demands that we compute the full n -torsion of \mathcal{A} , possibly over an extension of the actual (finite) base field.

Consequently, an order \mathcal{O} of the complex multiplication field K of \mathcal{A} can be tested to be contained in $\text{End}(\mathcal{A})$ by computing a generating set for \mathcal{O} , writing its elements α in the form $\frac{1}{n} \sum_i \alpha_i \pi^i$, and testing whether $\sum_i \alpha_i \pi^i$ kills the full n -torsion of \mathcal{A} for all such α . A module basis for \mathcal{O} has cardinality $2g$, but since \mathbb{Z} is contained in both \mathcal{O} and $\mathbb{Z}[\pi]$, only $2g - 1$ tests are really required; furthermore, as only an algebra basis is required, much fewer elements actually need to be tested.

The proposition requires denominators n to be coprime to the order q of the base field. When the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ is coprime to q , this can always be made the case: since the index of $\mathbb{Z}[\pi, \bar{\pi}, q\alpha]$ in $\mathbb{Z}[\pi, \bar{\pi}, \alpha]$ divides q and both orders contain $\mathbb{Z}[\pi, \bar{\pi}]$, this index must be one, which means that $q\alpha$ and α belong exactly to the same orders above $\mathbb{Z}[\pi, \bar{\pi}]$; therefore, the factor of n divisible by a power of q can simply be dropped.

This method is suited to local computations: similarly to what we did above, if ℓ is a prime, one can show that $\text{End}(\mathcal{A})_\ell = \mathcal{O}_\ell$ can be determined only using elements whose denominators are powers of ℓ . We will later rely on this local version to determine the endomorphism ring locally at small primes ℓ where our own algorithm fails to compute it.

When g is fixed and we work over base fields of increasing prime cardinality q , it becomes increasingly rare for q to divide the index $\mathbb{Z}[\pi, \bar{\pi}]$, although this can be seen to happen. In those cases where we want to determine the endomorphism ring locally at a large prime, the present method is probably not the best suited in the first place.

Two building blocks remain to be explained: computing the full ℓ -torsion, and efficiently finding the endomorphism ring by testing whether $\mathcal{O} \subseteq \text{End}(\mathcal{A})$ for chosen orders \mathcal{O} ; algorithms for both will be described and analyzed in the next chapter. When g is fixed and q goes to infinity, we deduce that the worst-case overall complexity of this method is

$$\ell^{2g+o(1)} \log^{2+o(1)} q \quad \text{where} \quad \ell = q^{g^2/2+o(1)}.$$

Note that in the case that we only wish to test whether $\text{End}(\mathcal{A})$ is maximal, FREEMAN and LAUTER (2007) subsequently improved this method using specific probabilistic tests.

CORRESPONDENCES AND ENDOMORPHISMS

Let us now briefly introduce elements of the theory of correspondences as background material for the work of WAGNER (2009), which will be discussed below.

First define a *function field* K over k (which we write K/k) as a finitely generated extension of transcendence degree one. In Chapter II, we saw that function fields arise from algebraic varieties, but here we will work with them abstractly. For details on the following, we refer to Chapter IV of the collection of lectures by DEURING (1973).

Definition v.3.2. *Let K/k be a function field, and K'/k an extension field. There exists a function field L/l such that L contains K , $l \cap K = k$, and L is the composite extension of K and a subfield of l that is k -isomorphic to K' .*

The function field L/l is called the constant field extension of K/k by K'/k .

DEURING (1937) introduced *correspondences* as ideals of maximal orders of function fields L/l , up to both principal ideals and *constant ideals*, that is, ideals with nontrivial intersection with l . When L is the constant field extension of a function field $k(\mathcal{C})/k$ by another $k(\mathcal{C}')/k$ where \mathcal{C} and \mathcal{C}' are two algebraic curves defined over a finite field k , he showed that correspondence classes represent isogenies from the Jacobian variety of \mathcal{C} to that of \mathcal{C}' .

In the particular case that $\mathcal{C} = \mathcal{C}'$, this gives a bijection

$$C : \text{End}(\text{Jac } \mathcal{C}) \xrightarrow{\sim} \{\text{correspondence classes}\} = \mathfrak{I}(\mathcal{O}_L) / \sim$$

which is compatible with the ring structure in the sense that for all endomorphisms α and β we have $C(\alpha + \beta) = C(\alpha) \cdot C(\beta)$, and similarly there exists a computational way of deriving the composition $C(\alpha \circ \beta)$ from $C(\alpha)$ and $C(\beta)$.

For instance, correspondences representing the Frobenius endomorphism π , the Verschiebung endomorphism $\bar{\pi}$, and the identity I are easily obtained; multiplication-by- n is then represented by $C(I)^n$, and so on.

Finally, and this is maybe the most crucial point for what follows, the action of a correspondence on a point, that is, that of the endomorphism it represents can be evaluated simply in terms of elementary function field operations.

WAGNER'S ALGORITHM

To determine whether some prescribed algebraic number of $\mathbb{Q} \otimes \text{End}(\text{Jac } \mathcal{C})$ represents an endomorphism, start as before by writing it as an element $\alpha \in \mathbb{Z}[\pi]$ divided by some integer n ; the correspondence class $C(\alpha)$ is easily computed from $C(\pi)$, so it remains to determine whether it can be divided by n .

The main idea of WAGNER (2009) is to interpolate the hypothetical correspondence class $C(\alpha/n)$ over a set of small-torsion points: let P_i be a point of $\text{Jac}(\mathcal{C})$ of order m_i ; if it exists, $C(\alpha/n)$ should act as

$$P_i \mapsto (n^{-1} \bmod m_i)C(\alpha)(P_i);$$

and we can write equations asserting that a formal correspondence class D acts this way. WAGNER (2009) gives an upper bound on the number of points P_i required to completely characterize the action of α/n , that is, ensuring that if the system admits a solution D , then we must have $D = C(\alpha/n)$, and as a consequence $\alpha/n \in \text{End}(\text{Jac } \mathcal{C})$.

He exhibits correspondence class representatives which are compatible with the above operations and therefore allow efficient correspondence class computations. These representatives are written in Hermite normal form and are almost entirely determined by their norms due to the restrictive conditions required for being a representative.

Therefore, WAGNER (2009) focuses on *interpolating the norm*, which is of the form

$$N_{L/k(\mathcal{C})}(C(\alpha/n)) = x^l + \sum_{i=0}^{l-1} \frac{f_i}{g_i} x^i$$

for some degree $l \leq g$, where the indeterminates f_i and g_i are polynomials of bounded degree with coefficients in $k(\mathcal{C})$; see “Abschätzung der Grade der Polynome in x_2 ” in Section 4.5 on page 99.

The whole procedure is summarized in “Algorithmus 5: Approximation” of the same section on page 103. That algorithm takes as input a \mathbb{Z} -basis β of an order \mathcal{O} of which $C(\beta)$ is known, an element α of some order \mathcal{O} , and an integer n ; if α/n is an endomorphism, it returns a correspondence representing it, or returns `false` otherwise.

As we will describe in the next chapter, being able to test whether prescribed orders \mathcal{O} are contained in the endomorphism ring suffices to determine it in a polynomial number of steps in the size of the base field.

A short analysis of the method can be found in Section 4.9; in brief, the degree of the norm of α/n is polynomial in n and it thus requires interpolating a number of points which is polynomial in n . In the worst case, the overall algorithm therefore uses exponential time in the size of the base field.

Nevertheless, it has the interesting feature that, as n grows, testing whether α/n is an endomorphism becomes easier; indeed, the size of the hypothetical correspondence representing it then gets smaller, so a shorter system of equations can be used. Note that all methods we have previously seen showed the reverse phenomenon.

V.4 Supersingular Methods

For the sake of completeness, let us address the case of supersingular elliptic curves in this section (and this section only). Known methods for computing endomorphism rings of such curves all have an exponential asymptotic running time in the size of the base field; however, contrary to the ordinary case, we are quite pessimistic about the possibilities of improvement.

In addition to the methods presented here, we note that KOHEL (1996) has an algorithm that gives some information on the endomorphism ring of supersingular curves which suffices to determine it only in specific cases; however, we are unaware of further developments of this technique.

ISOGENOUS SUPERSINGULAR CURVES

We first present background results on supersingular elliptic curves, their isogeny classes, and their endomorphism rings. Most results originate from DEURING (1941).

Recall that an elliptic curve \mathcal{E} defined over a finite field of characteristic p is *supersingular* when it has no p -torsion. As a meager compensation for the troubles ahead, we have:

Proposition v.4.1. *Up to isomorphism, every supersingular elliptic curve defined over a finite field of characteristic p is defined over \mathbb{F}_{p^2} .*

As a consequence, it is simple to enumerate all such isomorphism classes. Endomorphism rings of supersingular curves can similarly be enumerated simply.

Proposition v.4.2. *Endomorphism rings of supersingular curves correspond bijectively to maximal orders of $\mathbb{Q}_{p,\infty}$, the quaternion algebra ramified only at p and ∞ . Two such curves defined over \mathbb{F}_{p^2} have the same endomorphism ring if and only if they are conjugate under $\text{Gal}(\mathbb{F}_{p^2}/\mathbb{F}_p)$.*

This is why we are sceptical as to the possibilities of substantial improvements on the computation of endomorphism rings in this case: since all orders are maximal, and there are exponentially many of them, there seems to be no way around considering each, one at a time. Although we have not yet presented our method which exploits the structure of the lattice of orders in the ordinary case, the localization that we have described earlier (and which suffices in dimension one) should convince the reader of the benefit of having such a structure.

As for ordinary curves, there is a theory of complex multiplication; however, care must be taken due to its non-commutativity.

Proposition v.4.3. *Fix a supersingular curve \mathcal{E} . For any left ideal \mathfrak{a} of $\text{End}(\mathcal{E})$ coprime to p , the degree of the isogeny $\phi_{\mathfrak{a}}$ with kernel $\ker(\phi_{\mathfrak{a}}) = \bigcap_{\alpha \in \mathfrak{a}} \ker(\alpha)$ is the norm of \mathfrak{a} ; all isogenies between supersingular curves arise in this way.*

If $\mathcal{E}' = \phi_{\mathfrak{a}}(\mathcal{E})$, then $\text{End}(\mathcal{E}')$ is the right order of \mathfrak{a} , that is, $\{x \in \mathbb{Q}_{p,\infty} : \mathfrak{a}x \subset \mathfrak{a}\}$. If additionally $\mathcal{E}'' = \phi_{\mathfrak{b}}(\mathcal{E})$, the curves \mathcal{E}' and \mathcal{E}'' are isomorphic if and only if \mathfrak{a} and \mathfrak{b} are in the same left ideal class.

Much more can be said on the structure of this isogeny graph: for instance, when $p = 1 \pmod{12}$, it is a Ramanujan graph, a particular case of expander graph with desirable properties, such as mixing properties for random walks, which makes it notably a suitable building block for a hash function, as was proposed by CHARLES, LAUTER, and GOREN (2009).

QUATERNION ALGEBRAS

To give the above an effective flavor, let us briefly recall various results related to the structure of quaternion algebras.

The structure of the quaternion algebra $\mathbb{Q}_{p,\infty}$ is readily given by a result of PIZER (1980) which states that

$$\mathbb{Q}_{p,\infty} \simeq \mathbb{Q}[i, j, k] / (i^2 - a, j^2 - b, ij + ji, ij - k),$$

for $(a, b) = \begin{cases} (-1, -1) & \text{if } p = 2, \\ (-1, -p) & \text{if } p = 3 \pmod{4}, \\ (-2, -p) & \text{if } p = 5 \pmod{8}, \\ (-p, -q) & \text{if } p = 1 \pmod{8}, \end{cases}$

where q can be any prime congruent to three modulo four, modulo which p is not a square.

To enumerate maximal orders of this algebra, we can exploit the proposition above which states that the isogeny graph is connected. Therefore, if \mathcal{O} is any maximal order of $\mathbb{Q}_{p,\infty}$ and \mathfrak{a} ranges through representatives of each left ideal class of \mathcal{O} , then the right order of \mathfrak{a} ranges through all maximal orders of the quaternion algebra $\mathbb{Q}_{p,\infty}$.

THE MCMURDY-LAUTER METHOD

To find out which specific maximal order of $\mathbb{Q}_{p,\infty}$ is isomorphic to $\text{End}(\mathcal{E})$, MCMURDY and LAUTER (2004) proposed to

- count the number of endomorphisms of \mathcal{E} of degree ℓ ;
- compare it to the number of elements of \mathcal{O} of norm ℓ .

By the proposition we saw earlier, isogenies correspond to ideals, and endomorphisms correspond to principal ideals. Therefore, when \mathcal{O} is the particular order isomorphic to $\text{End}(\mathcal{E})$, the two numbers must be equal.

Repeating the above for various primes ℓ different from the characteristic rules out orders \mathcal{O} from the candidate list, so that eventually the endomorphism ring alone remains. This formally proceeds as the following procedure.

Algorithm v.4.4.

INPUT: A supersingular elliptic curve $\mathcal{E}/\mathbb{F}_{p^2}$.

OUTPUT: An order isomorphic to its endomorphism ring.

1. Let L be the list of maximal orders of $\mathbb{Q}_{p,\infty}$.
2. Until L is a singleton:
3. Pick a prime ℓ , and count the degree- ℓ endomorphisms of \mathcal{E} .
4. Rule out orders of L with a different count of elements of norm ℓ .
5. Return the only element in L .

For Step 4 MCMURDY and LAUTER (2004) derive an explicit method in Section 3.2; it boils down to finding integer solutions of a quadratic equation.

This procedure behaves quite well in practice: its bottleneck is the enumeration of isogenies of degree ℓ from \mathcal{E} to \mathcal{E} ; MCMURDY and LAUTER (2004) give explicit formulas for $\ell = 2$ and $\ell = 3$, and the isogeny-computing machinery for elliptic curves is nowadays at a stage of development where such operations can be performed quickly for a large range of ℓ .

However, we stress that its termination is not guaranteed, as two distinct maximal orders of $\mathbb{Q}_{p,\infty}$ might have the same number of ideals of norm ℓ for infinitely many primes ℓ .

CERVIÑO'S ALGORITHM

Although testing the norm of ideals alone is not sufficient to guarantee the termination of the endomorphism-ring identifying process, CERVIÑO (2004) observed in his Proposition 4.5 that considering both the norm *and* the trace yields a sufficient amount of information after finitely many tests. More precisely, he proved the following.

Proposition v.4.5. *No two maximal orders of the quaternion algebra $\mathbb{Q}_{p,\infty}$ have the same set*

$$\{(\mathrm{tr}(\alpha), N(\alpha)) : \alpha \in \mathcal{O}, N(\alpha) \leq b\}$$

where b is a certain bound which is $\mathcal{O}(p)$.

The norm and trace of such numbers map to the norm and trace of the characteristic polynomial of the corresponding endomorphism: we have

$$\phi_{\alpha}^{(2)} - \mathrm{tr}(\phi_{\alpha})\phi_{\alpha} + N(\phi_{\alpha}) = 0$$

since the degree (or norm) of an isogeny is always known (as we construct them from their kernels), the trace of ϕ can be found by testing the possible values in turn over a sufficiently large extension of the base field.

This gives the following algorithm.

Algorithm v.4.6.

INPUT: A supersingular elliptic curve $\mathcal{E}/\mathbb{F}_{p^2}$.

OUTPUT: An order isomorphic to its endomorphism ring.

1. Let L be the list of maximal orders of $\mathbb{Q}_{p,\infty}$.
2. For successive primes ℓ , starting from $\ell = 2$:
3. Compute the multiset $I = \{\text{tr}(\phi)\}$,
where ϕ ranges over degree- ℓ endomorphisms of \mathcal{E} .
4. Rule out from L those orders \mathcal{O} for which $I \neq \{\text{tr}(\beta)\}$
where β ranges over the elements of norm ℓ in \mathcal{O} .
5. Return the only element in L .

By the proposition above, this algorithm terminates after $O(p)$ operations. Nevertheless, since computing the trace of the endomorphisms is extremely costly, the former procedure is more suited to a large range of practical problems, although it is not guaranteed to terminate.

References

1937. Max DEURING.
“Arithmetische Theorie der Korrespondenzen algebraischer Funktionenkörper”.
In: *Journal für die reine und angewandte Mathematik* 1937.177. Pages 161–192.
DOI: 10.1515/crll.1937.177.161.
1941. Max DEURING.
“Die Typen der Multiplikatorenringe elliptischer Funktionenkörper”.
In: *Abhandlungen aus dem mathematischen Seminar der hamburgischen Universität*
14. Pages 197–272.
1966. John TATE.
“Endomorphisms of abelian varieties over finite fields”.
In: *Inventiones mathematicae* 2.2. Pages 134–144. DOI: 10.1007/BF01404549.
1969. William C. WATERHOUSE.
“Abelian varieties over finite fields”.
In: *Annales Scientifiques de l'École Normale Supérieure* 2.4. Pages 521–560.

1971. Jacques VÉLU.
 “Isogénies entre courbes elliptiques”.
 In: *Comptes Rendus de l'Académie des Sciences de Paris*. A 273. Pages 238–241.
1973. Max DEURING.
Lectures on the Theory of Algebraic Functions of One Variable. Volume 314.
 Lecture Notes in Mathematics. Springer. ISBN: 3-540-06152-5.
1980. Arnold PIZER.
 “An algorithm for computing modular forms on $\Gamma_0(N)$ ”.
 In: *Journal of Algebra* 64.2. Pages 340–390.
 DOI: 10.1016/0021-8693(80)90151-9.
1995. Laurent DEWAGHE.
Un corollaire aux formules de Vélu. Preprint.
1996. David R. KOHEL.
 “Endomorphism rings of elliptic curves over finite fields”.
 PhD thesis. University of California at Berkeley.
 URL: <http://echidna.maths.usyd.edu.au/kohel/pub/thesis.pdf>.
1996. Hendrik W. LENSTRA.
 “Complex multiplication structure of elliptic curves”.
 In: *Journal of Number Theory* 56.2. Pages 227–241.
 DOI: 10.1006/jnth.1996.0015.
2002. Mireille FOUQUET and François MORAIN.
 “Isogeny volcanoes and the SEA algorithm”.
 In: *Algorithmic Number Theory — ANTS-V*.
 Edited by Claus FIEKER and David R. KOHEL. Volume 2369.
 Lecture Notes in Computer Science. Springer. Pages 47–62.
 DOI: 10.1007/3-540-45455-1_23.
2004. Juan M. CERVIÑO.
On the correspondence between supersingular elliptic curves and maximal quaternionic orders. arXiv.org: math/0404538.
2004. Ken McMURDY and Kristin E. LAUTER.
 “Explicit generators for endomorphism rings of supersingular elliptic curves”.
 In: *Number Theory Conference in Honor of Harold Stark*. University of Minnesota.
2007. David M. FREEMAN and Kristin E. LAUTER.
 “Computing endomorphism rings of Jacobians of genus 2 curves over finite fields”.
 In: *Algebraic Geometry and its Applications — SAGA '07*.

- Edited by Jean CHAUMINE, James HIRSCHFELD, and Robert ROLLAND. Volume 5. Number Theory and Its Applications. World Scientific. Pages 29–66.
DOI: 10.1142/9789812793430_0002.
2009. Denis X. CHARLES, Kristin E. LAUTER, and Eyal Z. GOREN.
“Cryptographic hash functions from expander graphs”.
In: *Journal of Cryptology* 22.1. Pages 93–113.
DOI: 10.1007/s00145-007-9002-x.
2009. Jean-Marc COUVEIGNES.
“Linearizing torsion classes in the Picard group of algebraic curves over finite fields”.
In: *Journal of Algebra* 321.8. Pages 2085–2118.
DOI: 10.1016/j.jalgebra.2008.09.032.
2009. Kirsten EISENTRÄGER and Kristin E. LAUTER.
“A CRT algorithm for constructing genus 2 curves over finite fields”.
In: *Arithmetic, Geometry and Coding Theory — AGCT ’10*.
Edited by François RODIER and Serge VLADUT. Volume 21. Séminaires et Congrès. Société Mathématique de France. Pages 161–176.
2009. Markus WAGNER.
“Über Korrespondenzen zwischen algebraischen Funktionenkörpern”.
PhD thesis. Technische Universität Berlin.
URL: <http://www.math.tu-berlin.de/~wagner/Diss.pdf>.
2010. Sorina IONICA and Antoine JOUX.
“Pairing the volcano”. In: *Algorithmic Number Theory — ANTS-IX*.
Edited by Guillaume HANROT, François MORAIN, and Emmanuel THOMÉ.
Volume 6197. Lecture Notes in Computer Science. Springer. Pages 201–218.
DOI: 10.1007/978-3-642-14518-6_18.
2010. David R. KOHEL.
“Endomorphisms, isogeny graphs, and moduli”.
In: *Workshop on Elliptic Curves and Computation — ECC ’10*. URL:
<http://research.microsoft.com/apps/video/dl.aspx?id=140496>.
2011. Romain COSSET and Damien ROBERT.
Computing (ℓ, ℓ) -isogenies in polynomial time on Jacobians of genus 2 curves.
IACR ePrint: 2011/143.

Subexponential Method

We have so far discussed endomorphism-ring computation methods with an exponential worst-case runtime, and will now present one of subexponential complexity.

This method was first introduced in B. and SUTHERLAND (2009) under a form quite specific to elliptic curves, and relying on several unproven assumptions. All assumptions but the GRH were later removed in B. (2011) by modifying parts of the algorithm. Here, we present a variant of this algorithm which applies to general abelian varieties.

We stress that this chapter considers abelian varieties without taking polarizations into account, which is not an effective approach in dimension $g > 1$, but allows for a conceptually simpler presentation. For $g = 1$, where polarizations are unneeded, it is highly effective, and the next chapter will be devoted to rigorously proving its probabilistic runtime under the generalized Riemann hypothesis, and its unconditional correctness.

Modifications that make our method practical for $g = 2$ will be presented in the last chapter; they are expectedly slower and rely on more unproven hypotheses.

VI.1 Algorithm Overview

Let \mathcal{A} be a simple ordinary abelian variety defined over a finite field; denote by K its complex multiplication field and fix an isomorphism $\iota : K \rightarrow \mathbb{Q} \otimes \text{End}(\mathcal{A})$, which will be implicitly understood from now on.

To locate $\text{End}(\mathcal{A})$ amongst candidate orders of K , the main idea to our subexponential method is to compute certain properties describing the Picard groups of candidate orders, and to test them via complex multiplication in the horizontal isogeny graph. Since there exist subexponential algorithms for computing Picard groups we are done... Almost so.

We now give the main ingredients enabling this approach. Computational details are given in subsequent sections, while proofs and rigorous analysis are in the next chapter.

LATTICE OF ORDERS

Let us first briefly recall results that express where the endomorphism ring is to be sought.

Let \mathcal{A} be a simple ordinary abelian variety of dimension g defined over a finite field with q elements. The Frobenius endomorphism π acts on geometric points of \mathcal{A} by raising their coordinates to the q^{th} power; its characteristic polynomial $\chi_\pi(x)$ is a q -Weil polynomial, which means that it is monic, has integer coefficients, and has $2g$ complex roots, each of absolute value \sqrt{q} .

Computing this polynomial is equivalent to counting the number of points on the variety over \mathbb{F}_{q^n} for $n \in \{1, 2, \dots, g\}$, as we have

$$\#\mathcal{A}(\mathbb{F}_{q^n}) = \text{Res}_u(\chi_\pi(u), u^n - 1).$$

SCHOOF (1985) proved that that this can be done in deterministic polynomial time in $\log(q)$ for elliptic curves; his algorithm was later generalized to abelian varieties by PILA (1990).

Many endomorphisms stem from the Frobenius endomorphism, since $\mathbb{Q} \otimes \text{End}(\mathcal{A}) \simeq \mathbb{Q}(\pi)$. Since the complex multiplication field $K = \mathbb{Q}(\pi)$ is isomorphic to $\mathbb{Q}[x]/(\chi_\pi(x))$, by computing the Weil polynomial of \mathcal{A} we have already determined the endomorphism ring *up to fractions*. Fixing $\iota : K \rightarrow \mathbb{Q} \otimes \text{End}(\mathcal{A})$ means fixing this isomorphism; here, we simply put $x = \pi$ and make this implicit from now on.

This isomorphism maps $\text{End}(\mathcal{A})$ to an order in K so we have

$$\mathbb{Z}[\pi, \bar{\pi}] \subset \text{End}(\mathcal{A}) \subset \mathcal{O}_K;$$

the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]$] is the square part of the quotient $\text{disc}(\mathbb{Z}[\pi, \bar{\pi}]) / \text{disc}(\mathcal{O}_K)$, and it measures how broad the search-range is. As a simple upper bound, we use $\Delta = \text{disc}(\mathbb{Z}[\pi, \bar{\pi}])$ which Lemma III.2.5 proved can be as big as $q^{g^2/2+o(1)}$ in the worst case.

The orders of K containing $\mathbb{Z}[\pi, \bar{\pi}]$ form the *lattice of orders*. Since it might contain exponentially many orders, we need to devise a better way of finding $\text{End}(\mathcal{A})$ than testing each order in turn. Computing $\text{End}(\mathcal{A})$ locally at many primes ℓ helps, but is not sufficient since (apart from the case that $g = 1$) the local lattices themselves might not have any nicer structure than the general one.

Instead of localizing, we use a lattice-ascending algorithm designed to only test polynomially many orders. For those orders \mathcal{O} , it *tests* whether $\mathcal{O} \subset \text{End}(\mathcal{A})$ using tools derived from complex multiplication theory.

PRINCIPAL IDEALS AND CERTIFICATES

We exclusively consider ideals of norm coprime to Δ , so that they are unramified and invertible in $\mathbb{Z}[\pi, \bar{\pi}]$. Recall that such ideals of \mathcal{O} act on the set $\text{AV}_{\mathcal{O}}(k)$ of abelian varieties

defined over the finite field k with endomorphism ring \mathcal{O} by $\alpha : \mathcal{A} \mapsto \phi_\alpha(\mathcal{A})$ where ϕ_α denotes the isogeny with kernel $\bigcap_{\alpha \in \mathfrak{a}} \ker(\alpha)$. We assume that this induces a faithful and transitive action of $\text{Pic}(\mathcal{O})$ on $AV_{\mathcal{O}}(k)$; by complex multiplication theory, this is always the case when \mathcal{O} is an imaginary quadratic order, or a ring of integers.

Intuitively, the structure of the Picard group of $\text{End}(\mathcal{A})$ therefore dictates that of the horizontal isogeny graph component containing \mathcal{A} . Our approach is essentially to look at the latter and deduce information on the former, which might eventually lead to the identification of $\text{End}(\mathcal{A})$. We formalize the notion of *structure* by the following concept.

Definition VI.1.1. *An ideal \mathfrak{a} of $\mathbb{Z}[\pi, \bar{\pi}]$ is said to be principal in \mathcal{O} if the ideal $\mathfrak{a}\mathcal{O}$ is principal; it is said to be principal in the isogeny graph when the isogeny ϕ_α is an endomorphism of \mathcal{A} .*

In fact, we meant $\phi_{\mathfrak{a}\text{End}(\mathcal{A})}$ rather than ϕ_α since we want it to act on \mathcal{A} even though \mathfrak{a} is an ideal of $\mathbb{Z}[\pi, \bar{\pi}]$. Obviously, since we are looking for $\text{End}(\mathcal{A})$ we cannot really compute $\mathfrak{a}\text{End}(\mathcal{A})$, but we will see later that $\phi_{\mathfrak{a}\text{End}(\mathcal{A})}$ can be computed regardless.

Therefore, an ideal is principal in $\text{End}(\mathcal{A})$ if and only if it is principal in the isogeny graph, which gives a way to tell the endomorphism ring apart from other orders of the lattice. To avoid testing all orders, we rely on this simple result.

Lemma VI.1.2. *If an ideal is principal in some order, it is principal in all orders containing it.*

Indeed, if $\mathcal{O} \subset \mathcal{O}'$ are two orders containing $\mathbb{Z}[\pi, \bar{\pi}]$, the map $\mathfrak{a} \in \mathcal{I}(\mathcal{O}) \mapsto \mathfrak{a}\mathcal{O}' \in \mathcal{I}(\mathcal{O}')$ induces, as we have mentioned before, a surjective morphism of Picard groups. Intuitively, this means that more and more ideals become principal as we ascend the lattice of orders, or equivalently that Picard groups get smaller. This is why we chose $\mathbb{Z}[\pi, \bar{\pi}]$ to be the ring of our ideals: via the morphism $\mathfrak{a} \mapsto \mathfrak{a}\mathcal{O}$ we can map ideals of $\mathbb{Z}[\pi, \bar{\pi}]$ to any order of the lattice.

Computationally, the lemma above implies that by verifying whether principal ideals of \mathcal{O} are also principal in the isogeny graph, we can convince ourselves that \mathcal{O} is contained in $\text{End}(\mathcal{A})$. However, this approach does not prove anything (in fact, it fails in certain rare cases that we will cover later); to rigorously assert the location of the endomorphism ring, we use the following concept.

Definition VI.1.3. *A certificate for the order \mathcal{O} consists of:*

- a family of orders \mathcal{O}_i and ideals \mathfrak{a}_i principal in \mathcal{O}_i but not in \mathcal{O} ,
- a family of orders \mathcal{O}_j and ideals \mathfrak{a}_j principal in \mathcal{O} but not in \mathcal{O}_j ,

such that \mathcal{O} is the only order above $\mathbb{Z}[\pi, \bar{\pi}]$ satisfying $\mathcal{O}_i \not\subset \mathcal{O}$ and $\mathcal{O}_j \not\supset \mathcal{O}$ for all indices.

It is said to be verified on the abelian variety \mathcal{A} if the ideals \mathfrak{a}_j are principal in its isogeny graph whereas the \mathfrak{a}_i are not.

If a certificate for the order \mathcal{O} is verified on the abelian variety \mathcal{A} , by the contrapositive of the lemma above, then we have $\text{End}(\mathcal{A}) = \mathcal{O}$. In fact, the family $(\mathcal{O}_i, \mathfrak{a}_i)$ is effectively constructed when one executes the lattice-ascending walk that we are about to describe; the family \mathcal{O}_i is then typically chosen to consist of all orders immediately below \mathcal{O} , that is, just one level below \mathcal{O} in the lattice of orders.

The next section will address the search for ideals and, as a consequence, show that it takes $L(q^{g^2})^{1/4\gamma+o(1)}$ time to generate a certificate that can subsequently be verified within $L(q^{g^2})^{3g\gamma+o(1)}$ operations, as q goes to infinity and γ is any positive constant real number. This eliminates the need to carefully ensure the correctness of our algorithm: we can simply run an algorithm that is only proven to return a correct result with probability $\varepsilon > 0$ and, when it does return a result, verify it using our certificate method; if it proves to be incorrect, we start over. The expected overhead on the complexity is $1/\varepsilon$.

COMPUTING FROM BELOW

To search for the endomorphism ring $\text{End}(\mathcal{A})$ in the lattice of orders, we *test* whether orders \mathcal{O} lie below it by selecting principal ideals of them and checking whether they are principal in the isogeny graph.

It remains to design a general strategy to select the orders to be tested.

We shall say that an order \mathcal{O} lies directly above another \mathcal{O}' if we have $\mathcal{O} \supset \mathcal{O}'$ but there exists no order \mathcal{O}'' different from \mathcal{O} and \mathcal{O}' satisfying $\mathcal{O} \supset \mathcal{O}'' \supset \mathcal{O}'$; we also define the corresponding notion of “directly below” where inclusions are reversed. As an example, when an order contains another with prime index, then it must lie directly above it.

To ascend the lattice of orders, we proceed one step at a time: each step consists in enumerating all orders lying directly above a prescribed order \mathcal{O}' . We have seen that the index of \mathcal{O}' in any order directly above it is a divisor of ℓ^{2g-1} where ℓ is a prime factor of $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$. By factoring Δ we therefore obtain the possible values of ℓ , and we can then use the algorithm described earlier that lists those orders containing \mathcal{O}' with a prescribed index.

Our strategy to locate the endomorphism ring in this lattice by testing orders and ascending in corresponding directions works as follows: given some order \mathcal{O}' contained in $\text{End}(\mathcal{A})$ (we start with $\mathcal{O}' = \mathbb{Z}[\pi, \bar{\pi}]$), find some order \mathcal{O} directly above \mathcal{O}' which lies below $\text{End}(\mathcal{A})$; then replace \mathcal{O}' by \mathcal{O} and iterate the process. The ascension ends when no \mathcal{O} is found to be contained in $\text{End}(\mathcal{A})$; then, we must have $\text{End}(\mathcal{A}) \simeq \mathcal{O}'$. See Figure 5 where we start from the bottom and ascend towards orders \mathcal{O} for which the statement $\mathcal{O} \subset \text{End}(\mathcal{A})$ holds.

Formally, we obtain the following algorithm.




FIGURE 5. Locating $\text{End}(\mathcal{A})$ by ascending a test-sequence of orders.

Algorithm VI.1.4.

INPUT: A simple ordinary abelian variety \mathcal{A} over a finite field \mathbb{F}_q .

OUTPUT: An order isomorphic to its endomorphism ring.

1. Compute the Frobenius polynomial $\chi_\pi(x)$ of \mathcal{A} .
2. Factor the discriminant Δ and construct the order $\mathcal{O}' = \mathbb{Z}[\pi, \bar{\pi}]$.
3. For orders \mathcal{O} directly above \mathcal{O}' :
4. If $\mathcal{O} \subset \text{End}(\mathcal{A})$ set $\mathcal{O}' \leftarrow \mathcal{O}$ and go to Step 3.
5. Return \mathcal{O}' .

To test whether an order lies above \mathcal{O} we compute sufficiently many principal ideals of it and test whether they are principal in the isogeny graph. Before detailing this process, let us present an alternative approach to locating the endomorphism ring in the lattice of orders.

The next sections will show that it requires $L(|\Delta|)^{1/4\gamma+o(1)}$ time to find random principal ideals \mathcal{O} whose associated isogenies can be computed within $L(|\Delta|)^{3g\gamma+o(1)}$ operations; to balance these costs, we set $\gamma = 1/\sqrt{12g}$ and since $|\Delta| < q^{g^2+o(1)}$ we find an overall runtime of

$$L(q)^g \sqrt{3g/2+o(1)}.$$

Note that for $g = 1$ we can do better by using a faster isogeny computing method whose exponent is just 2γ instead of $3g\gamma$ for the arbitrary-dimension method.

COMPUTING FROM ABOVE

Rather than start at the bottom of the lattice and ascend towards the endomorphism ring, we can generate certificates for each order starting from the top and attempt to verify them; to ensure this only uses subexponentially many operations, we *trim* the lattice of orders as we go. The runtime is then bounded in the size of the output, rather than the input. The method of WAGNER (2009) had a similar feature; however, our bound is subexponential.

In most cases, there are only polynomially many orders in $\log|\Delta|$, but to give a subexponential bound on the complexity of our algorithm when there are exponentially many, we *eliminate* small branches of orders as we go; these branches correspond to small prime power factors ℓ of the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$; by “eliminating them,” we mean computing the endomorphism locally at ℓ using the method of EISENTRÄGER and LAUTER (2009). Formally, we proceed as follows.

Notation. Let $b_x(f(x))$ denote any function satisfying $f(x) < b_x(f(x)) < f(x)^{1+o(1)}$ that can be evaluated in essentially linear time in $f(x)$.

Algorithm VI.1.5.

INPUT: A simple ordinary abelian variety \mathcal{A} over a finite field \mathbb{F}_q .

OUTPUT: An order isomorphic to its endomorphism ring.

1. Compute the Frobenius polynomial $\chi_\pi(x)$, and factor $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ as $\prod \ell^{v_\ell}$.
2. Set $S \leftarrow \emptyset$ and $r \leftarrow 2$.
3. For all primes ℓ with $\ell^{2g v_\ell} < b_r(\exp \sqrt{\log(r)})$:
4. If $\ell \notin S$, compute $\text{End}(\mathcal{A})_\ell$ and add ℓ to S .
5. For all orders \mathcal{O} with $\forall \ell \in S, \mathcal{O}_\ell = \text{End}(\mathcal{A})_\ell$ and $|\text{disc}(\mathcal{O})| < r$:
6. Test whether $\text{End}(\mathcal{A}) = \mathcal{O}$; if yes, then return \mathcal{O} .
7. Set $r \leftarrow r^{1+1/b_q(\log q)}$ and go back to Step 2.

Step 4 applies the method of Eisenträger and Lauter locally at ℓ ; its complexity is therefore $\ell^{2g v_\ell + o(1)}$, omitting polynomial factors in $\log(q)$. The inequality of Step 3 thus ensures that no more than $L^{o(1)}(r)$ operations are spent there.

The cost of generating a certificate for \mathcal{O} is bounded by $L(\text{disc}(\mathcal{O}))^{1/4\gamma + o(1)}$ when the verification time is bounded by $L(\text{disc}(\mathcal{O}))^{3\gamma + o(1)}$; to balance these, Step 6 uses $\gamma = 1/\sqrt{12g}$ which gives it a complexity bound of $L(\text{disc}(\mathcal{O}))^{\sqrt{3g/2} + o(1)}$. Step 5 ensures that:

- only orders that match the local information obtained in Step 3 are tested;
- testing them all uses at most $L^{O(1)}(r)$ computing time.

Step 7 increments r little by little so that, on the one hand, it never goes much beyond the discriminant of $\text{End}(\mathcal{A})$, and, on the other hand, it takes only $O(g \log q)^2$ iterations for r to reach $|\text{disc}(\mathbb{Z}[\pi, \bar{\pi}])| = O(q^{g^2+o(1)})$ and thus for our algorithm to have considered all orders.

To bound the number of orders to be tested in Step 6, assume that there are at most $n^{1+o(1)}$ orders contained in \mathcal{O} with index n ; this is a classical fact for $g = 1$ (since orders are identified by their index in \mathcal{O}_K) and it has been proven by NAKAGAWA (1996) for $g = 2$. We thus find that for $r = n^2$ the number of orders satisfying the condition of Step 5 is bounded, up to exponent $1 + o(1)$, by the number of divisors of

$$[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]] \prod_{\ell^{g\ell} < \exp \sqrt{\log r}} \ell^{v_\ell}$$

that are less than n , where the denominator removes prime powers from S ; a crude calculation shows that this number is bounded polynomially in $\log(q)$.

Ignoring the cost of factoring the discriminant Δ , and omitting polynomial factors in $\log(q)$, we obtain an overall complexity of

$$L(\text{disc}(\text{End } \mathcal{A})) \sqrt{3g/2+o(1)}.$$

VI.2 Finding Principal Ideals

To test whether some prescribed order \mathcal{O} lies below the endomorphism ring of a simple ordinary abelian variety \mathcal{A} , we first compute principal ideals \mathfrak{a} that discriminate the structure of $\text{Pic}(\mathcal{O})$ from that of other orders containing $\mathbb{Z}[\pi, \bar{\pi}]$. Then, we evaluate the corresponding isogenies; for this reason, we compute the factorization $\mathfrak{a} = \prod \mathfrak{p}^{z_{\mathfrak{p}}}$ and then evaluate $\phi_{\mathfrak{a}}$ as the composition of $z_{\mathfrak{p}}$ times the isogeny $\phi_{\mathfrak{p}}$, for all \mathfrak{p} .

We therefore consider smooth ideals with small exponents, which we call *short ideals*.

GENERIC METHODS

Let \mathfrak{B} be a generating set of ideals for the Picard group of an order \mathcal{O} in a number field K ; for instance, under the generalized Riemann hypothesis, we can take for \mathfrak{B} the set of prime ideals of norm less than $12 \log^2 |\text{disc } \mathcal{O}|$. By *computing relations* of \mathfrak{B} , we mean finding products of ideals of \mathfrak{B} that are principal.

For convenience of the exposition and of the implementation, let \mathfrak{B} actually generate the Picard group of the minimal order \mathfrak{m} ; this way, the set $\{\mathfrak{b}\mathcal{O} : \mathfrak{b} \in \mathfrak{B}\}$ generates the Picard

group of any order \mathcal{O} containing \mathfrak{m} , and its relations are vectors under the product map

$$\sigma: \begin{cases} \mathbb{Z}^{\mathfrak{B}} & \rightarrow \mathcal{I}(\mathfrak{m}) \\ x & \mapsto \prod_{\mathfrak{p} \in \mathfrak{B}} \mathfrak{p}^{x_{\mathfrak{p}}} \end{cases}.$$

If we let $\sigma_{\mathcal{O}}(x)$ denote the ideal class of $\text{Pic}(\mathcal{O})$ containing the ideal $\sigma(x)\mathcal{O}$, then the set of relation vectors $x \in \mathbb{Z}^{\mathfrak{B}}$ for \mathcal{O} is exactly the lattice $\Lambda_{\mathcal{O}} = \ker(\sigma_{\mathcal{O}})$. Note that since \mathfrak{B} generates the Picard group, the map $\sigma_{\mathcal{O}}$ is surjective and we have

$$\text{Pic } \mathcal{O} \simeq \mathbb{Z}^{\mathfrak{B}} / \Lambda_{\mathcal{O}}$$

which means that computing relations is essentially equivalent to computing the group structure of $\text{Pic}(\mathcal{O})$. The principal ideals of \mathcal{O} we search for will be obtained in the form $\sigma_{\mathcal{O}}(z)$, where $z \in \Lambda_{\mathcal{O}}$ is a relation vector to be found.

To find kernel vectors of $\sigma_{\mathcal{O}}$, we first need to identify a finite subset of $\mathbb{Z}^{\mathfrak{B}}$ which is big enough to contain a generating set for $\Lambda_{\mathcal{O}}$. Let n denote the class number of \mathcal{O} ; since $\text{Pic}(\mathcal{O})$ is generated by \mathfrak{B} and its elements have order n at most, the box $\{0, \dots, n-1\}^{\mathfrak{B}}$ maps surjectively onto the Picard group via $\sigma_{\mathcal{O}}$. As a consequence, there exists a generating set for $\Lambda_{\mathcal{O}}$ contained in the box $B = \{0, \dots, n\}^{\mathfrak{B}}$. We spare the proof to the reader, since a much better bound will be derived (and proved) shortly.

Note that the class number n satisfies $n = |\text{disc } \mathcal{O}|^{1/2+o(1)}$; however, analytic methods can be used to derive effective, tighter bounds on n .

To find relations of the group $G = \text{Pic}(\mathcal{O})$ on B , one can use the baby-step giant-step method. It consists in splitting the basis \mathfrak{B} into a disjoint union $\mathfrak{B}_0 \sqcup \mathfrak{B}_1$ of two sets of approximately equal size, so that this splitting carries over to box B and decomposes it as a direct product $B_0 \times B_1$, where B_i is the set of vectors of B with support in \mathfrak{B}_i .

Algorithm VI.2.1.

INPUT: A box B where to look for relations under $\sigma_{\mathcal{O}} : B \rightarrow G$.

OUTPUT: A relation, that is, a vector of $\ker(\sigma_{\mathcal{O}})$.

1. Split B as the direct product $B_0 \times B_1$.
2. For vectors $x \in B_0$: store x in a table indexed by $\sigma_{\mathcal{O}}(x)$.
3. For vectors $y \in B_1$:
4. If $(\sigma_{\mathcal{O}}(y))^{-1} = \sigma_{\mathcal{O}}x$, return the relation $x + y$.

The table constructed in Step 2 is typically implemented as a hash table, so that the cost of the lookup in Step 4 is negligible. A Gray code can be used to enumerate elements of B_0 and B_1 so that each evaluation of $\sigma_{\mathcal{O}}$ just requires $O(1)$ operations. This algorithm then requires an expected $O(\sqrt{n})$ number of group operations and storage space.

Note that a space-efficient generic method for finding relations in arbitrary finite groups will be presented in the next chapter; it can be used in Picard groups in particular. For the moment, let us discuss a simple application of such generic algorithms to the computation of endomorphism rings.

RELATIONS OF THE ENDOMORPHISM RING

Let us briefly present an alternative to our approach to computing the endomorphism ring $\text{End}(\mathcal{A})$ of a simple ordinary abelian variety \mathcal{A} defined over a finite field: we first gave a method for computing $\text{End}(\mathcal{A})$ *from below* by finding principal ideals of candidate orders and testing them in the isogeny graph; then we gave a method which works *from above* by attempting to prove that $\mathcal{O} = \text{End}(\mathcal{A})$ for orders \mathcal{O} of increasing discriminant.

A more direct way of computing $\text{End}(\mathcal{A})$ *from above* is simply to *reverse* our first method which proceeds *from below*: rather than finding relations of orders and evaluating them in the isogeny graph, we can find relations in the isogeny graph and evaluate them in Picard groups. This gives the method below.

Algorithm VI.2.2.

INPUT: *A simple ordinary abelian variety \mathcal{A} over a finite field \mathbb{F}_q .*

OUTPUT: *An order isomorphic to its endomorphism ring.*

1. *Compute the Frobenius polynomial $\chi_\pi(x)$ of \mathcal{A} .*
2. *Factor the discriminant Δ and construct the order $\mathcal{O}' = \mathcal{O}_K$.*
3. *For orders \mathcal{O} directly below \mathcal{O}' :*
4. *If $\text{End}(\mathcal{A}) \subset \mathcal{O}$ set $\mathcal{O}' \leftarrow \mathcal{O}$ and go to Step 3.*
5. *Return \mathcal{O}' .*

To test whether $\text{End}(\mathcal{A})$ lies below some order \mathcal{O} , we find isogeny chains from \mathcal{A} to itself: in the baby-step giant-step algorithm above, it suffices to replace σ_θ by the map

$$x \in \mathbb{N}^{2g} \mapsto \underbrace{\phi_{p_1} \circ \cdots \circ \phi_{p_1}}_{x_{p_1} \text{ times}} \circ \underbrace{\phi_{p_2} \circ \cdots \circ \phi_{p_2}}_{x_{p_2} \text{ times}} \circ \cdots (\mathcal{A})$$

(better yet, use the Pollard approach of the next chapter); once a principal ideal of the isogeny graph is found, it suffices to check whether it is principal in the order \mathcal{O} as well.

This approach has the advantage that, quite often, only one relation of the isogeny graph suffices to rule out all orders but one, so the endomorphism ring is computed in just one shot.

As before, this is a probabilistic process: the ideal we find in $\text{End}(\mathcal{A})$ might actually also be principal in some strictly smaller order; in order to increase the probability of success, we can use several relations, but to unconditionally prove the output (henceforth transforming our method into an algorithm of Las-Vegas type), we have to rely on certificates.

SUBEXPONENTIAL ALGORITHMS

SEYSEN (1987) first gave an algorithm for finding relations of $\Lambda_{\mathcal{O}}$ when \mathcal{O} is an imaginary quadratic orders; building upon it, HAFNER and MCCURLEY (1989) proved that the full Picard group structure, that is, a generating set for $\Lambda_{\mathcal{O}}$, can be determined in proven subexponential time under the generalized Riemann hypothesis. This was later extended by BUCHMANN (1989) to arbitrary number fields, under additional heuristic assumptions.

All find relations using a classical smoothness-based technique which exploits the integer-like structure of ideals in number fields.

Algorithm VI.2.3.

INPUT: A box B where to look for relations under $\sigma_{\mathcal{O}} : B \rightarrow \text{Pic}(\mathcal{O})$.

OUTPUT: A relation, that is, a vector of $\ker(\sigma_{\mathcal{O}})$.

1. Take a random element $x \in B$ and compute $\mathfrak{a} = \sigma_{\mathcal{O}}(x)$.
2. Reduce \mathfrak{a} to an equivalent but smaller ideal \mathfrak{b} .
3. If possible, find a preimage $y \in \sigma_{\mathcal{O}}^{-1}(\mathfrak{b})$ and return $x - y$.
4. Return to Step 1.

To find preimages easily, SEYSEN (1987) takes as basis \mathfrak{B} the set of prime ideals of norm less than some bound, so that the existence of a preimage in B can be asserted by a smoothness test on the norm of the ideal, and the factorization of that norm yields the preimage. Several ingredients are needed to bound its complexity, the most important one being that a random integer in $\{1, \dots, n\}$ has a probability $L(n)^{-1/2c+o(1)}$ of being $L(n)^c$ -smooth, for any constant $c > 0$; in the case that \mathcal{O} is an imaginary quadratic orders, SEYSEN (1987) proved that norms of reduced ideals are distributed as random integers; in fact, this behavior is observed, although not proven, for orders of general number fields as well.

The next chapter will present all these arguments rigorously.

SMALLER BOXES

Since our relations (and the ideals derived from them) are expected to discriminate the endomorphism ring from other orders of the lattice, we must ensure that when we generate a relation in $\Lambda_{\mathcal{O}}$ for some order \mathcal{O} , it does not belong to $\Lambda_{\mathcal{O}'}$ for some other order \mathcal{O}' . Of course, we have seen that $\mathcal{O} \subset \mathcal{O}'$ implies that $\Lambda_{\mathcal{O}} \subset \Lambda_{\mathcal{O}'}$, and our lattice-ascending algorithm actually takes advantage of that, so we should rather require the above for orders \mathcal{O}' not above \mathcal{O} , that is, $\mathcal{O} \not\subset \mathcal{O}'$.

Note that there exist orders $\mathcal{O} \neq \mathcal{O}'$ with $\Lambda_{\mathcal{O}} = \Lambda_{\mathcal{O}'}$, but not too many: for $g = 1$, there are just three such cases, and we can easily fall back on a specific method to deal with them. Rigorous details will be given in the next chapter.

In general, to ensure that the relations z we generate belong to $\Lambda_{\mathcal{O}}$ but not another $\Lambda_{\mathcal{O}'}$, we require that they are *random relations* in the sense that, for any order \mathcal{O}' above \mathcal{O} , we have

$$\text{Prob} [z \in \Lambda_{\mathcal{O}'} | z \in \Lambda_{\mathcal{O}}] = \frac{\# \text{Pic } \mathcal{O}'}{\# \text{Pic } \mathcal{O}} + o(1);$$

in other words, the relation is quasi-uniformly distributed in the quotient $\Lambda_{\mathcal{O}'}/\Lambda_{\mathcal{O}}$.

To obtain random relations of \mathcal{O} , HAFNER and MCCURLEY (1989) used vectors z with coordinates up to n^4 , where n is the class number. In the Picard group, a double-and-add method can be used to compute each term $\mathfrak{p}^{z_{\mathfrak{p}}}$ in time linear in $\log(n)$, so that $\sigma_{\mathcal{O}}$ can be evaluated in subexponential time.

However, for the purpose of checking whether the ideal $\sigma(x-y)$ is principal in the isogeny graph, the associated isogeny needs to be evaluated. For this, there is no double-and-add technique, and the isogeny $\phi_{\mathfrak{p}}$ has to be evaluated $z_{\mathfrak{p}}$ times, which makes the bound n^4 on the coordinates quite painful. Note that since y is the exponent vector in the factorization of the norm of a reduced ideal, it is at most linear in $\log n$, so what is really needed here to keep the isogeny-computing cost low is just to find a smaller box B for which the quasi-uniform distribution of classes still holds.

A conjectural small box was first used by B. and SUTHERLAND (2009); later, CHILDS, JAO, and SOUKHAREV (2010) noted that a result of JAO, MILLER, and VENKATESAN (2009) enables to prove, under the generalized Riemann hypothesis, that such a box indeed yields random relations. We conclude with an explicit version of the general algorithm.

Algorithm VI.2.4.

INPUT: An order \mathcal{O} of discriminant D .

OUTPUT: A random relation $z \in \Lambda_{\mathcal{O}}$.

1. Form the set \mathfrak{B} of primes \mathfrak{p} of \mathcal{O} with norm less than $N = L(D)^{\gamma}$.
2. Draw uniformly at random a vector $x \in \mathbb{Z}^{\mathfrak{B}}$ with coordinates $|x_{\mathfrak{p}}| < b_D(\log^{4+\varepsilon} |D|)$ if $N(\mathfrak{p}) < b_D(\log^{2+\varepsilon} |D|)$, else $x_{\mathfrak{p}} = 0$.
3. Compute a reduced ideal \mathfrak{a} in the class $\sigma_{\mathcal{O}}(x)$.
4. If \mathfrak{a} factors over \mathfrak{B} as $\prod \mathfrak{p}^{y_{\mathfrak{p}}}$ then return the vector $x - y$.
5. Otherwise, go back to Step 2.

Here, ε stands for any fixed positive real number. Step 3 may use the LLL algorithm as we mentioned earlier; for any “good” reduction method, the probability that Step 4 is successful is $L(D)^{-1/4\gamma+o(1)}$; the overall complexity is then $L(D)^{1/4\gamma+o(1)}$ to generate a relation of length $L(D)^{\gamma}$; the longer the relation, the costlier the evaluation of the associated isogeny.

VI.3 Computing the Action of Ideals

We now consider effective means of testing whether an ideal \mathfrak{a} acts trivially on the isogeny graph of an abelian variety \mathcal{A} . Here, we focus on the case of elliptic curves, but certain bricks will be reused in the last chapter for abelian varieties of dimension two.

MODULAR EQUATIONS

Once a principal ideal \mathfrak{a} of \mathcal{O} in the form $\prod_{\mathfrak{p}} \mathfrak{p}^{z_{\mathfrak{p}}}$ is found, we wish to determine whether the associated isogeny acts trivially on \mathcal{A} ; in fact, this does not require explicitly evaluating the isogeny $\phi_{\mathfrak{a}}$, but only determining whether it maps \mathcal{A} on \mathcal{A} .

Elliptic curves isogenous to a given one with a prescribed way can be listed efficiently via modular polynomials; this uses j -invariants to identify isomorphism classes of curves, and modular polynomials $\Phi_m(X, Y)$ which we now recall.

Proposition VI.3.1. *For any $m \in \mathbb{N}$, there exists some polynomial $\Phi_m(X, Y) \in \mathbb{Q}[X, Y]$ of degree $m + 1$ such that, over fields of characteristic coprime to m , the j -invariants of elliptic curves m -isogenous to a prescribed j_0 are exactly the roots of $\Phi_m(X, j_0)$.*

COHEN (1984) proved the bit-size of Φ_m to be $O(m^{3+o(1)})$. It can be computed in quasi-linear time by the floating-point method of ENGE (2009), or by the alternative method of BRÖKER, LAUTER, and SUTHERLAND (2010) based on the Chinese remainder theorem, which offers additional advantages such as reduced memory requirements.

To test whether $\phi_{\mathfrak{a}}$ acts trivially on \mathcal{A} , we can evaluate $\Phi_{N(\mathfrak{a})}(X, Y)$ at $(j(\mathcal{A}), j(\mathcal{A}))$. If the result is non-zero, then $\phi_{\mathfrak{a}}$ cannot send \mathcal{A} to \mathcal{A} ; if the result is zero, then there exists one isogeny of degree $N(\mathfrak{a})$ from \mathcal{A} to \mathcal{A} , but it need not be $\phi_{\mathfrak{a}}$ in general.

For practical purposes, rather than seeing $\phi_{\mathfrak{a}}$ as an isogeny of degree $N(\mathfrak{a})$, we see it as a chain formed of $z_{\mathfrak{p}}$ isogenies of norm $N(\mathfrak{p})$ for each $\mathfrak{p} \in \mathfrak{B}$. Consequently, it suffices to compute the modular polynomials $\Phi_{N(\mathfrak{p})}$ and to combine them as isogeny steps. We now detail this procedure, in a manner which also addresses the issue of the previous paragraph.

CARDINALITIES

When we evaluate $\Phi_{N(\mathfrak{p})}(X, Y)$ at $X = j(\mathcal{A})$, the roots in Y are the j -invariants of the codomain of degree- $N(\mathfrak{p})$ isogenies with domain \mathcal{A} . Amongst these roots lies $\phi_{\mathfrak{p}}(\mathcal{A})$ but we have no information as to which it is.

To address this, we can explore *all* isogenies of degree $N(\mathfrak{p})$. When \mathfrak{a} has many factors, this can be costly as we might have to consider several roots of $\Phi_{N(\mathfrak{p})}$ at each step of the isogeny chain, therefore eventually exploring an exponential number of varieties in $\log N(\mathfrak{a})$.

Endomorphism rings of elliptic curves are imaginary quadratic orders, and there are therefore at most two ideals of a given prime norm: \mathfrak{p} and $\bar{\mathfrak{p}}$. In the isogeny chain

$$\mathcal{A}_0 \xrightarrow{\phi_{\mathfrak{p}}} \mathcal{A}_1 \xrightarrow{\phi_{\mathfrak{p}}} \cdots \xrightarrow{\phi_{\mathfrak{p}}} \mathcal{A}_{z_{\mathfrak{p}}}$$

corresponding to the factor $\mathfrak{p}^{z_{\mathfrak{p}}}$ of \mathfrak{a} , the conjugate prime $\bar{\mathfrak{p}}$ acts on \mathcal{A}_i as the dual isogeny of $\phi_{\mathfrak{p}} : \mathcal{A}_{i-1} \rightarrow \mathcal{A}_i$. Thus, for $i > 0$, we can determine which of the two roots of $\Phi_{N(\mathfrak{p})}(j(\mathcal{A}_i), Y)$ is *not going backward* in the chain, and the two roots need to be considered only for $i = 0$.

This helps when \mathfrak{a} does not have many prime factors but has one with high exponent: rather than just testing if $\prod_{\mathfrak{p} \in \mathfrak{B}} \mathfrak{p}^{z_{\mathfrak{p}}}$ is principal, we count how many products $\prod_{\mathfrak{p} \in \mathfrak{B}} \hat{\mathfrak{p}}^{z_{\mathfrak{p}}}$ are, where $\hat{\mathfrak{p}} \in \{\mathfrak{p}, \bar{\mathfrak{p}}\}$; this is equivalent to counting the number of endomorphisms of \mathcal{A} that are chains consisting in $z_{\mathfrak{p}}$ non-backwards isogenies of degree $N(\mathfrak{p})$, for each \mathfrak{p} .

When there are just two ideals \mathfrak{p} and $\bar{\mathfrak{p}}$ of norm $N(\mathfrak{p})$, this gives:

Definition VI.3.2. Let $\prod_{\mathfrak{p} \in \mathfrak{B}} \mathfrak{p}^{z_{\mathfrak{p}}}$ be the factorization of an ideal $\mathfrak{a} \in \mathbb{Z}[\pi, \bar{\pi}]$.

Its cardinality in \mathcal{O} is the number of vectors $(\hat{\mathfrak{p}}) \in \prod_{\mathfrak{p} \in \mathfrak{B}} \{\mathfrak{p}, \bar{\mathfrak{p}}\}$ for which $\prod_{\mathfrak{p} \in \mathfrak{B}} \hat{\mathfrak{p}}^{z_{\mathfrak{p}}}$ is trivial.

Its cardinality in the isogeny graph of \mathcal{A} is the number of chains formed by $z_{\mathfrak{p}}$ isogenies of norm $N(\mathfrak{p})$, for each $\mathfrak{p} \in \mathfrak{B}$, which map \mathcal{A} onto itself.

These two quantities are the same for $\mathcal{O} = \text{End}(\mathcal{A})$, and, for elliptic curves, we evaluate the latter via using the method below starting from the j -invariant $j_0 = j(\mathcal{A})$.

Algorithm VI.3.3.

INPUT: A j -invariant j_0 and an ideal $\prod_{\mathfrak{p} \in \mathfrak{B}} \mathfrak{p}^{z_{\mathfrak{p}}}$.

OUTPUT: The cardinality of this ideal in the isogeny graph of j_0 .

1. Let \mathcal{J}' be the list (j_0) .
2. For each $\mathfrak{p} \in \mathfrak{B}$:
 3. Set $\mathcal{J} \leftarrow \mathcal{J}'$ and let \mathcal{J}' be an empty list.
 4. For each j in \mathcal{J} :
 5. Let $\{j_+, j_-\}$ be the roots of $\Phi_{N(\mathfrak{p})}(X, j)$, and set $j'_+ \leftarrow j$ and $j'_- \leftarrow j$.
 6. Repeat $z_{\mathfrak{p}} - 1$ times:
 7. Set $(j''_+, j''_+) \leftarrow (j_+, \text{the root of } \Phi_{N(\mathfrak{p})}(X, j_+) \text{ different from } j'_+)$.
 8. Set $(j''_-, j''_-) \leftarrow (j_-, \text{the root of } \Phi_{N(\mathfrak{p})}(X, j_-) \text{ different from } j'_-)$.
 9. Append j''_+ and j''_- to \mathcal{J}' .
10. Return the multiplicity of j_0 in \mathcal{J}' .

Since we compute two branches for each prime factor of \mathfrak{a} , the overhead this *cardinality* algorithm adds on the *principal* approach is 2^w where w is the number of prime factors. When w is small, this is greatly compensated by the speed of using modular polynomials.

COMPLEX MULTIPLICATION ACTION

We briefly review results on evaluating the explicit isogeny $\phi_{\mathfrak{p}}$ associated to an ideal \mathfrak{p} .

Recall Proposition III.3.8 which states that invertible prime ideals \mathfrak{p} of \mathcal{O} written as $\ell\mathcal{O} + u(\pi)\mathcal{O}$ act on the kernel of the associated isogeny $\phi_{\mathfrak{p}}$ with characteristic polynomial u . Therefore, to tell the isogeny $\phi_{\mathfrak{p}}$ apart from other isogenies of degree $N(\mathfrak{p})$, one need just compute the action of the Frobenius endomorphism on its kernel.

To evaluate isogenies from their kernels, we use the formulas of VÉLU (1971) for elliptic curves, and their generalization to abelian varieties by LUBICZ and ROBERT (2009) together with the improvements of COSSET and ROBERT (2011). These methods take as input a subgroup \mathcal{H} of an abelian variety \mathcal{A} and output the isogeny $\mathcal{A} \rightarrow \mathcal{A}/\mathcal{H}$. Since they work with principally polarized abelian varieties, they additionally require that \mathcal{H} be a maximal isotropic subgroup with respect to the Weil pairing, and that it be isomorphic to $(\mathbb{Z}/\ell)^g$.

We thus seek ideals $\mathfrak{a} = \prod \mathfrak{q}^{e_{\mathfrak{q}}}$ where the kernel of each $\phi_{\mathfrak{q}}$ is maximal isotropic and of type $(\mathbb{Z}/\ell)^g$; to this extent, in dimension $g > 1$, we restrict to ideals \mathfrak{a} arising via the reflex type norm, on which the last chapter will say more. When we have a prime decomposition $\mathfrak{q} = \prod \mathfrak{p}$ for a specific term \mathfrak{q} , the Frobenius endomorphism must act on $\ker(\phi_{\mathfrak{q}})$ with characteristic polynomial $\prod u_{\mathfrak{p}}(x)$ where the $u_{\mathfrak{p}}(x)$ are such that $\mathfrak{p} = N(\mathfrak{p})\mathcal{O} + u_{\mathfrak{p}}(\pi)\mathcal{O}$.

Finally, we observe that, if \mathcal{A} is an ordinary abelian variety of dimension g defined over a finite field, all points of rational subgroups of type $(\mathbb{Z}/\ell)^g$ are defined over an extension of degree at most $\ell^g - 1$.

The characteristic polynomial of the action, on such a subgroup \mathcal{H} , of the Frobenius endomorphism divides $\chi_{\pi}(x) \bmod \ell$, and the multiplicative order n of x modulo this factor is precisely the extension degree over which all points of \mathcal{H} are defined. Therefore, to evaluate the degree of an extension over which all points of rational subgroups of type $(\mathbb{Z}/\ell)^g$ are defined, it suffices to compute the least common multiple of the multiplicative order of x modulo the degree- g factors of $\chi_{\pi}(x) \bmod \ell$.

DIRECT METHOD

Let \mathfrak{q} be an ideal such that $\ker(\phi_{\mathfrak{q}})$ is a maximal isotropic subgroup of order ℓ^g in \mathcal{A} . In order to compute this isogeny, we combine several classical tools into the algorithm below. It requires a basis for the ℓ -torsion of \mathcal{A} defined over a certain extension, which we will soon explain how to compute; the kernel is then identified by the polynomial $u = \prod u_{\mathfrak{p}}$ with $u_{\mathfrak{p}}$ defined as above, and we use the explicit isogeny algorithm to compute $\phi_{\mathfrak{q}}$ from it. We make this algorithm output the isogenous curve $\phi_{\mathfrak{q}}(\mathcal{A})$, so it can readily be plugged in to our endomorphism ring computing method.

Algorithm VI.3.4.

INPUT: An abelian variety $\mathcal{A} / \mathbb{F}_q$ with Frobenius polynomial χ_π and a suitable ideal \mathfrak{q} of norm ℓ^g .

OUTPUT: The isogenous variety $\phi_{\mathfrak{q}}(\mathcal{A})$.

1. Find a basis (P_i) of the $\mathcal{A}[\ell]$ over the extension of degree $\ell^g - 1$ of \mathbb{F}_q .
2. Write the matrix M of the Frobenius endomorphism on the basis (P_i) .
3. Enumerate those subspaces of dimension g stable under $M \in \text{Mat}_{2g}(\mathbb{Z}/\ell\mathbb{Z})$.
4. Determine which corresponds to \mathfrak{q} using the Frobenius action.
5. Compute the isogeny of which this eigenspace is the kernel.

For a maximal isotropic subgroup of \mathcal{A} of order ℓ^g defined over the extension of degree $\ell^g - 1$ of the base field, the method of LUBICZ and ROBERT (2009) requires $\ell^{3g+o(1)}$ operations as g is fixed and ℓ goes to infinity.

Step 2 decomposes $\pi(P_i)$ as $\sum_{j \in \{1, \dots, 2g\}} M_{ij} P_j$ for which a baby-step giant-step approach uses $O(\ell^g)$ operations over the extension field. Step 3 is classical and takes quasi-linear time in $g^\omega \log(\ell)$ where $\omega < 2.376$ is the best known exponent for matrix multiplication.

Finally, Step 1 uses Theorem 1 of COUVEIGNES (2009), where the extension is chosen so as to contain all points of rational subgroups of type $(\mathbb{Z}/\ell)^g$. The simple algorithm we give below actually computes all such points, from which a basis can easily be extracted; it works by selecting random ℓ^∞ -torsion points and lifting them along each others. Here, we let $k(P)$ denote the valuation at a fixed prime ℓ of the order of a point P .

Algorithm VI.3.5.

INPUT: An abelian variety $\mathcal{A} / \mathbb{F}_q$ with Frobenius polynomial χ_π and a prime ℓ .

OUTPUT: The ℓ -torsion subgroup of \mathcal{A} over $\mathbb{F}_{q^{\ell^g-1}}$.

1. Write $\#\mathcal{A}(\mathbb{F}_{q^{\ell^g-1}})$ as $m\ell^k$ where $\ell \nmid m$.
2. Create an empty associative array B .
3. While B has fewer than ℓ^{2g} keys:
 4. Let $P = mO$ where O is a random point of $\mathcal{A}(\mathbb{F}_{q^{\ell^g-1}})$.
 5. For j from $k(P) - 1$ down to 1, if $\ell^j P$ is a key of B :
 6. If $j > k(B[\ell^j P])$ then go to Step 9.
 7. Set $P \leftarrow P - \ell^{k(B[\ell^j P]) - j - 1} B[\ell^j P]$.
 8. If $P = 0$ then go back to Step 4.
 9. For all keys Q of B and $x \in \{1, \dots, \ell\}$, set $B[\ell^{k(xP+Q)-1}(xP+Q)] \leftarrow xP+Q$.
10. Return the keys of B .

Random points of \mathcal{A} can be drawn efficiently when \mathcal{A} is given as the Jacobian variety of a curve in Weierstrass form. Using the last two algorithms, we compute, in Mumford

coordinates, the kernel of the isogeny that we wish to evaluate; we then convert it to theta representation where the algorithm of COSSET and ROBERT (2011) is applied, and finally use the method of MESTRE (1991) to convert the codomain variety back as the Jacobian of a curve in Weierstrass form, so that the whole process can be iterated.

Since the cardinality of $\mathcal{A}(\mathbb{F}_{q^{\ell^g-1}})$ is $q^{\ell^g+o(1)}$ multiplying random points of it by m uses $O(\ell^g \log q)$ operations in $\mathcal{A}(\mathbb{F}_{q^{\ell^g-1}})$. Similarly, all orders are bounded by $k = O(\ell^g \log q)$. Finally, the probability of going back to Step 4 is $O(1/\ell)$ as proven by COUVEIGNES (2009).

Using fast field arithmetic, and representing points of \mathcal{A} in Mumford coordinates, operations in $\mathcal{A}(\mathbb{F}_{q^{\ell^g-1}})$ have a bit complexity of $(\ell^g \log q)^{1+o(1)}$; if an efficient data structure such as a red-black tree is used to store the keys of B , we have:

Proposition VI.3.6. *Let \mathcal{A}/\mathbb{F}_q be an abelian variety of known Frobenius polynomial, and \mathfrak{q} a suitable ideal of $\mathbb{Z}[\pi, \bar{\pi}]$. Algorithm VI.3.4 returns the abelian variety $\phi_{\mathfrak{q} \text{End}(\mathcal{A})}(\mathcal{A})$ in time bounded by $(\ell^g \log q)^{2+o(1)}$, as g is fixed and ℓ goes to infinity.*

Note that, in Algorithm VI.3.5, rather than storing the whole ℓ -torsion subgroup in an associative array, a pairing could be used to transport discrete logarithm problems to a finite field where they can be more efficiently solved. This technique gives a valuable speedup for large values of ℓ , although the overall complexity remains polynomial in ℓ due to the extension field arithmetic.

VI.4 Practical Computations

We now present the algorithms used and results obtained by practical runs on elliptic curves. Applying the same techniques to general abelian varieties will be the topic of the last chapter. Timings reported here were measured on a single core of a recent desktop computer, such as an AMD Opteron clocked at 2 GHz.

BALANCING THE COSTS

Let \mathcal{E} be an ordinary elliptic curve defined over a finite field \mathbb{F}_q . The first step of our algorithm is to compute the characteristic polynomial χ_π of the Frobenius endomorphism of \mathcal{E} . It is equivalent to counting the number of points of \mathcal{E} which is of the form $\chi_\pi(1) = p + 1 - t$ for a certain integer $t \in \{-2\sqrt{q}, \dots, 2\sqrt{q}\}$. Over a base field of cryptographic size, say, with q a prime of 256 bits, this takes under ten seconds on just one core of a standard desktop computer using the Schoof–Elkies–Atkin algorithm. Note that further developments by SUTHERLAND (2011) now make this possible for primes p over 5000 decimal digits.

Next, we need to find principal ideals of orders \mathcal{O} , and start by deciding which prime factors we want them to have. For maximal orders \mathcal{O} of imaginary quadratic fields, BACH




FIGURE 6. Dots plot the minimal k such that every class of $\text{Pic}(\mathcal{O})$ contains the product of a subset of S_k . Gray dots cover all imaginary quadratic orders \mathcal{O} of discriminant at least -10^8 , and black dots are for 10^4 random \mathcal{O} drawn according to a logarithmic distribution. The lines represent $k = d \log_2(\#\text{Pic } \mathcal{O})$ for $d = 1, 2$.

(1990) proved under the generalized Riemann hypothesis that the primes up to $6 \log^2 |\Delta|$ generate the Picard group, where Δ is the discriminant of \mathcal{O} . Heuristically, we find that much less are necessary, which lead to the following conjecture.

Conjecture VI.4.1. *For any $d > 1$, if \mathcal{O} is an imaginary quadratic order of sufficiently large discriminant, then any class of $\text{Pic}(\mathcal{O})$ contains the product of a subset of S_k , where S_k contains the first $k = d \log_2(\#\text{Pic } \mathcal{O})$ non-principal prime ideals.*

This is actually stronger than asking for S_k to generate the Picard group: it requires that S_k generates it *with bounded exponents* in $\{0, 1\}$. However, it is a natural conjecture to make since it asserts that the set S_k behaves as a random subset of $\text{Pic}(\mathcal{O})$ would in the sense of Proposition 1.1 of IMPAGLIAZZO and NAOR (1996). Our empirical verifications have not found a single order for which the conjecture does not hold with $d = 2$; for values of d closer to 1, we found this to be true for many orders above a certain lower bound, as can be seen on Figure 6.

The above is most useful when generating relations using generic methods: it states that only slightly more primes than a cardinality argument would require actually suffice. This yields short associated isogenies (which are a must in dimension two).

However, as we strive to balance the cost of finding a principal ideal in \mathcal{O} with that of evaluating the associated isogeny, generic methods do not scale well: for discriminants of more than 128 bits, a generic method would require above 32 operations in $\text{Pic}(\mathcal{O})$; from there on it is therefore advisable to switch to the subexponential method of SEYSEN (1987). Note that by the conjecture we can use a box with support in S_k .

Since our principal ideals rarely have more than 10 prime factors, it is really worth using the cardinality approach: modular polynomials permit one to compute isogenous curves quickly, and they can be precomputed and reduced modulo p for all the primes ℓ we consider, whereas computing the torsion would have to be done from scratch at each step.

HARDCODING CERTIFICATES IN DIMENSION ONE

So far, our endomorphism ring computing method tested whether $\mathcal{O} \subset \text{End}(\mathcal{A})$ for various orders \mathcal{O} ; since this process has a small probability of failure, we then certified the candidate order so as to unconditionally verify our result.

In B. and SUTHERLAND (2009), we used a quite different approach which simultaneously finds \mathcal{O} and verifies it. It exploits the particular structure of the lattice of orders for elliptic curves; we start by recalling this structure.

Let w denote the index of $\mathbb{Z}[\pi]$ in $\mathcal{O}_{\mathbb{Q}(\pi)}$ where π denotes the Frobenius endomorphism of an ordinary elliptic curve defined over a finite field. Orders \mathcal{O} of $K = \mathbb{Q}(\pi)$ have the form $\mathbb{Z} + f\mathcal{O}_K$ where f is the integer that generates their conductor over \mathcal{O}_K ; therefore, inclusion of orders corresponds to divisibility of conductors, so that orders containing $\mathbb{Z}[\pi]$ are in bijection with divisors f of w .

Let p^i be a prime power dividing w , and consider the problem of deciding whether p^i divides the conductor u of $\text{End}(\mathcal{E})$. Here, a certificate for p^i needs only consist of one ideal \mathfrak{a} which is principal in the order of conductor $w/p^{\text{val}_p(w-i+1)}$ but not in that of conductor p^i : if \mathfrak{a} is principal in the isogeny graph of \mathcal{E} , then we necessarily have $p^i | u$. Indeed, in that situation, $\text{End}(\mathcal{E})$ does not contain the order with conductor $w/p^{\text{val}_p(w-i+1)}$, which means its conductor u divides w without dividing $w/p^{\text{val}_p(w-i+1)}$, in other words, p^i divides u .

In number fields of degree greater than two, it does not seem to be possible to certify orders in a nice way as above, using just one ideal; that is why we needed to develop a more general method for arbitrary abelian varieties.

GENERIC EXAMPLE

Let \mathcal{E} be the elliptic curve with Weierstrass equation

$$Y^2 = X^3 - 3X + 2728849899765998058103612158899570741955717345$$

over \mathbb{F}_q with $q = 2872801286401014961877470682093858455400487431$

This curve is ordinary and it has $q + 1 - t$ points, for a trace t of 1868. The discriminant of π is $4q - t^2$ and its factors as $-7w^2$ where

$$w = 2 \cdot 127 \cdot 524287 \cdot 304250263527209.$$

We first compute the endomorphism locally at 2 using the method of EISENTRÄGER and LAUTER (2009), which is nearly instantaneous; it finds that the order with conductor 2 does not contain $\text{End}(\mathcal{E})$.

For the prime 127, we use the local method of KOHEL (1996): since $\Phi_{127}(j(\mathcal{E}), \cdot)$ proves to have multiple roots, 127 does not divide the conductor of $\text{End}(\mathcal{E})$. This also takes negligible time.

Since $\text{End}(\mathcal{E})$ was found to be maximal locally at 2 and 127, we can now simply set $w \rightarrow w/254$ and work with this new w . For the bigger primes, we turn to finding principal ideals with a generic method and verifying them in the isogeny graph. We choose to work with the degrees 11, 23, 29, 37, and 43, meaning that we look for a principal ideal as a product of prime ideals of norm these numbers. Note that it is interesting to use only a few primes here since then few modular polynomials have to be computed, and can be reused many times.

Using hardcoded certificates, it is easier to deal with bigger primes, so let us start with the biggest one p . Using the baby-step giant-step method, we find in just a second that the relation $23^4 \cdot 29^7 \cdot 37^{17} \cdot 43^4$ has cardinality 4 in the order with conductor w/p , and zero in that with conductor p . Computing the associated tree of isogenies took 9 seconds; as it turns out, the cardinality of the relation in the isogeny graph is 4 as well, therefore p does not divide the conductor of $\text{End}(\mathcal{E})$.

We finish with $p = 524287$: again, we look for a relation with the baby-step giant-step method; it takes roughly 20 minutes to uncover the relation $11^{47} \cdot 23^{707} \cdot 29^{540} \cdot 37^{103} \cdot 43^{197}$ which has cardinality 2 in the order with conductor w/p , and zero in that with conductor p . The associated tree of isogenies took 6 minutes to compute, and since the relation has cardinality zero in the isogeny graph, we conclude that $\text{End}(\mathcal{E})$ has conductor p .

In less than half an hour, we therefore established that $\text{End}(\mathcal{E}) = \mathbb{Z} + 524287\mathcal{O}_K$.

The runtime of this generic method is bounded by $q^{1/4+o(1)}$ but if we had computed $\text{End}(\mathcal{E})$ from above by searching for relations in its isogeny graph, the bound would have been $(\text{disc } \text{End}(\mathcal{E}))^{1/4+o(1)}$. However, since our curve was generated with the complex multiplication method (to give it an interesting endomorphism ring), it would not have been fair: we would have found $\text{End}(\mathcal{E})$ much too quickly!

SUBEXPONENTIAL EXAMPLE

Let \mathcal{E} be the elliptic curve with Weierstrass equation

$$Y^2 = X^3 - 3X + 660897170071025494489036936911 \backslash \\ 196131075522079970680898049528 \\ \text{over } \mathbb{F}_q \text{ with } q = 160693804425899027555081234320 \backslash \\ 6050075546550943415909014478299$$

where the backslash symbol denotes that a number has been wrapped over to the next line. Again, the curve is ordinary and it has trace $t = 212$ (which it takes just a few seconds to compute). Factoring the discriminant $4q - t^2$ of $\mathbb{Z}[\pi]$, we find that

$$w = 2 \cdot 127 \cdot \underbrace{524287}_{p_1} \cdot \underbrace{7195777666870732918103}_{p_2}.$$

As before, the primes 2 and 127 can be dealt with by climbing the local volcano. None of them divides the conductor u of $\text{End}(\mathcal{E})$; this only takes a few seconds.

To determine whether p_1 divides u , we use the algorithm of SEYSEN (1987) with the smoothness bound 600 to find a relation with non-zero cardinality in the order of conductor w/p_1 . It takes about four minutes to find the relation

$$2^{1798} \cdot 23^3 \cdot 29^1 \cdot 37^2 \cdot 53^{29} \cdot 137^1 \cdot 149^1 \cdot 233^1 \cdot 263^2 \cdot 547^1$$

whose cardinality in the order with conductor p_1 is zero. Computing the relevant modular polynomials via the method of BRÖCKER, LAUTER, and SUTHERLAND (2010) requires under four minutes and the associated tree of isogenies is found to have cardinality zero within just a minute; as a consequence, we deduce that p_1 is a factor of u . Note that, here, we made use of the prime 2 although it divides the index w ; this process is described in Section 4.2 of SUTHERLAND (2011).

For the prime p_2 , this is, as expected, much faster: the relation $2^{23} \cdot 11^5 \cdot 43^1 \cdot 71^2$ is found to have positive cardinality in the order with conductor w/p_2 but not that with conductor p_2 . It is found that p_2 does not divide u and the whole process takes just a few seconds.

In about 5 minutes, we have thus proved that $\text{End}(\mathcal{E})$ has conductor 524287, but note that this computation was much more difficult than the previous one due to the larger size of p_2 here: it could not have been achieved with generic methods.

References

1971. Jacques VÉLU.
 “Isogénies entre courbes elliptiques”.
 In: *Comptes Rendus de l’Académie des Sciences de Paris*. A 273. Pages 238–241.

1984. Paula COHEN.
“On the coefficients of the transformation polynomials for the elliptic modular function”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 95. Pages 389–402. DOI: 10.1017/S0305004100061697.
1985. René SCHOOF.
“Elliptic curves over finite fields and the computation of square roots mod p ”.
In: *Mathematics of Computation* 44.170. Pages 483–494.
DOI: 10.2307/2007968.
1987. Martin SEYSEN.
“A probabilistic factorization algorithm with quadratic forms of negative discriminant”. In: *Mathematics of Computation* 48.178. Pages 757–780.
DOI: 10.1090/S0025-5718-1987-0878705-X.
1989. Johannes BUCHMANN.
“A subexponential algorithm for the determination of class groups and regulators of algebraic number fields”. In: *Séminaire de Théorie des Nombres, Paris*.
Edited by Catherine GOLDSTEIN. Volume 91. Progress in Mathematics. Birkhäuser.
Pages 27–41.
1989. James L. HAFNER and Kevin S. MCCURLEY.
“A rigorous subexponential algorithm for computation of class groups”.
In: *Journal of the American Mathematical Society* 2.4. Pages 837–850.
DOI: 10.2307/1990896.
1990. Eric BACH.
“Explicit bounds for primality testing and related problems”.
In: *Mathematics of Computation* 55.191. Pages 355–380.
DOI: 10.1090/S0025-5718-1990-1023756-8.
1990. Jonathan PILA.
“Frobenius maps of abelian varieties and finding roots of unity in finite fields”.
In: *Mathematics of Computation* 55.192. Pages 745–763.
DOI: 10.2307/2008445.
1991. Jean-François MESTRE.
“Construction de courbes de genre 2 à partir de leurs modules”.
In: *Effective methods in algebraic geometry — MEGA '90*.
Edited by Teo MORA and Carlo TRAVERSO. Volume 94. Progress in Mathematics.
Birkhäuser. Pages 313–334.

1996. Russel IMPAGLIAZZO and Moni NAOR.
“Efficient cryptographic schemes provably as secure as subset sum”.
In: *Journal of Cryptology* 9.4. Pages 236–241. DOI: 10.1109/SFCS.1989.63484.
1996. David R. KOHEL.
“Endomorphism rings of elliptic curves over finite fields”.
PhD thesis. University of California at Berkeley.
URL: <http://echidna.maths.usyd.edu.au/kohel/pub/thesis.pdf>.
1996. Jin NAKAGAWA.
“Orders of a quartic field”.
In: *Memoirs of the American Mathematical Society* 122.583.
2009. Gaetan BISSON and Andrew V. SUTHERLAND.
“Computing the endomorphism ring of an ordinary elliptic curve over a finite field”.
In: *Journal of Number Theory* 131.5. Edited by Neal KOBLITZ and Victor S. MILLER.
Special Issue on Elliptic Curve Cryptography. Pages 815–831.
DOI: 10.1016/j.jnt.2009.11.003.
2009. Jean-Marc COUVEIGNES.
“Linearizing torsion classes in the Picard group of algebraic curves over finite fields”.
In: *Journal of Algebra* 321.8. Pages 2085–2118.
DOI: 10.1016/j.jalgebra.2008.09.032.
2009. Kirsten EISENTRÄGER and Kristin E. LAUTER.
“A CRT algorithm for constructing genus 2 curves over finite fields”.
In: *Arithmetic, Geometry and Coding Theory — AGCT’10*.
Edited by François RODIER and Serge VLADUT. Volume 21. Séminaires et Congrès.
Société Mathématique de France. Pages 161–176.
2009. Andreas ENGE.
“Computing modular polynomials in quasi-linear time”.
In: *Mathematics of Computation* 78.267. Pages 1809–1824.
DOI: 10.1090/S0025-5718-09-02199-1.
2009. David JAO, Stephen D. MILLER, and Ramarathnam VENKATESAN.
“Expander graphs based on GRH with an application to elliptic curve cryptography”.
In: *Journal of Number Theory* 129.6. Pages 1491–1504.
DOI: 10.1016/j.jnt.2008.11.006.
2009. David LUBICZ and Damien ROBERT.
Computing isogenies between abelian varieties. arXiv.org: 1001.2016.

2009. Markus WAGNER.
“Über Korrespondenzen zwischen algebraischen Funktionenkörper”.
PhD thesis. Technische Universität Berlin.
URL: <http://www.math.tu-berlin.de/~wagner/Diss.pdf>.
2010. Reinier BRÖKER, Kristin LAUTER, and Andrew V. SUTHERLAND.
Modular polynomials via isogeny volcanoes.
To appear in *Mathematics of Computation*. arXiv.org: 1001.0402.
2010. Andrew M. CHILDS, David JAO, and Vladimir SOUKHAREV.
Constructing elliptic curve isogenies in quantum subexponential time.
arXiv.org: 1012.4019.
2011. Gaetan BISSON.
Computing endomorphism rings of elliptic curves under the GRH.
arXiv.org: 1101.4323.
2011. Romain COSSET and Damien ROBERT.
Computing (ℓ, ℓ) -isogenies in polynomial time on Jacobians of genus 2 curves.
IACR ePrint: 2011/143.
2011. Andrew V. SUTHERLAND.
“Computing Hilbert class polynomials with the Chinese remainder theorem”.
In: *Mathematics of Computation* 80. Pages 501–538.
DOI: 10.1090/S0025-5718-2010-02373-7.
2011. Andrew V. SUTHERLAND.
Genus 1 point counting in quadratic space and essentially quartic time. In preparation.

Complexity Analysis

This chapter is devoted to a rigorous analysis of the method that we have just presented; the main result is a proof, under the generalized Riemann hypothesis, that our algorithm indeed computes endomorphism rings of ordinary elliptic curves in subexponential time.

Most of material used has already appeared in B. (2011) for elliptic curves; here, when this can be done, we state our results for general varieties. Polarization issues are deferred to the next chapter, which will therefore also cover practical computational aspects in dimension $g > 1$.

As usual, let \mathcal{A} be a simple ordinary abelian variety defined over a finite field \mathbb{F}_q .

VII.1 Orders from Picard Groups

We first prove that if we can identify the structure of the Picard group of the endomorphism ring of \mathcal{A} , then we can determine $\text{End}(\mathcal{A})$ unambiguously.

PRELIMINARIES

Recall that the first step is to compute the characteristic polynomial χ_π of the Frobenius endomorphism π of \mathcal{A} . For this, we use the method of PILA (1990) and more precisely the improved algorithm of ADLEMAN and HUANG (2001) which, when \mathcal{A} is the Jacobian variety of a genus- g hyperelliptic curve, has a complexity of

$$(\log q)^{O(g^2 \log g)}.$$

Even if it were not for cryptographic reasons, we would avoid non-Jacobian varieties since our algorithm requires to efficiently draw points at random, which we cannot do when \mathcal{A} is expressed in a more general form (such as theta constants).

The number of points of \mathcal{A} defined over the extension of degree e is then

$$\#\mathcal{A}(\mathbb{F}_{q^e}) = \text{Res}_u(\chi_\pi(u), u^e - 1)$$

which means that our algorithm for computing the ℓ -torsion does not have to count the number of points over a new extension every time a new prime ℓ is considered.

To navigate the lattice of orders of the complex multiplication field $K = \mathbb{Q}[X]/(\chi_\pi(X))$, that is, compute $\mathfrak{M} = \mathcal{O}_K$, $\mathfrak{m} = \mathbb{Z}[\pi, \bar{\pi}]$ and the factorization of $[\mathfrak{M} : \mathfrak{m}]$, we need to factor the discriminant Δ of χ_π which satisfies

$$|\Delta| \leq (2\sqrt{q})^{2g(2g-1)}.$$

For this, the unconditional method of LENSTRA and POMERANCE (1992) uses $L(|\Delta|)^{1+o(1)}$ operations; assuming unproved hypotheses, we might also use the number field sieve of COPPERSMITH (1993) with conjectured runtime

$$L_{1/3}^{\text{NFS}}(|\Delta|) \quad \text{where } c_{\text{NFS}} = \frac{1}{3} \sqrt[3]{92 + 26\sqrt{13}} \approx 1.902.$$

For elliptic curves, we were able to prove the correctness and complexity of the rest of our method only assuming the generalized Riemann hypothesis. In that case, the complexity is

$$L(q)^{1/\sqrt{2+o(1)}},$$

so the cost of factoring via the unconditionally proven method dominates; we found it curious that no known factoring algorithm achieves a better exponent assuming solely the generalized Riemann hypothesis: there seems to be a gap in the hypothesis required as, in terms of asymptotically fastest methods, we go straight from an unconditionally proven method to one which relies on many non-standard heuristics.

In dimension two, we will see that additional unproved hypotheses, other than the generalized Riemann hypothesis, are necessary.

ORDERS AND IDEALS

Let us briefly address the complexity of the algorithms used for navigating the lattice and computing with ideals of arbitrary orders in it.

The algorithms used greatly differ from dimension one to dimension two: in dimension one, the lattice is simply the set of divisors of $[\mathfrak{M} : \mathfrak{m}]$ while in higher dimension its structure has no such special form; again in dimension one, ideals can be dealt with extremely efficiently as binary quadratic forms while in higher dimension only general methods involving Hermite normal form and LLL reduction can be used.

In fact, we find that, in the realm of elliptic curves, many problems can be solved in *essentially linear* time, that is, with a complexity asymptotically equivalent to the size of the output, up to an exponent of $1 + o(1)$; but those problems become suddenly much harder with higher-dimensional abelian varieties and no such satisfying algorithm is known. This is for instance the case for the generation of Hilbert class polynomials. Our own endomorphism ring computing algorithm will not be an exception to this rule, as many simple and easy to analyze aspects of it are lost when going from dimension $g = 1$ to $g = 2$.

Regardless of the dimension, since we use the building blocks for orders and ideals on inputs of size for which their complexity is polynomial in $\log(q)$, we need not worry too much about them: as our overall expected complexity is superpolynomial, the cost of all these subroutines disappears within the $o(1)$ term of the exponent. This might seem a little too rough, so we refer to COHEN (1993) for more careful statements regarding the complexity of these standards calculations.

ORDERS WITH SAME PICARD GROUP

Our relation method uses the Picard group structure to characterize an order. This section and the next are devoted to proving the correctness of this approach: here, we will see that there are not many orders with the same Picard group structure, and there, we will describe a workaround technique for distinguishing these rare orders from each other.

We first consider the one-dimensional case, as the ideal structure of non-maximal orders is much better understood in this case. If \mathcal{O} is an order of an imaginary quadratic field K , we let \mathfrak{B} be a generating set of ideals for $\text{Pic}(\mathcal{O})$, and denote by $\Lambda_{\mathcal{O}}$ the relations of $\text{Pic}(\mathcal{O})$ for this basis \mathfrak{B} ; in other words, we assume that $\text{Pic}(\mathcal{O}) \simeq \mathbb{Z}^{\mathfrak{B}} / \Lambda_{\mathcal{O}}$.

Proposition VII.1.1. *Let \mathcal{O} and \mathcal{O}' be two orders in an imaginary quadratic field K . The lattice $\Lambda_{\mathcal{O}'}$ contains $\Lambda_{\mathcal{O}}$ if and only if the order \mathcal{O}' contains \mathcal{O} or if one of the following holds:*

1. $K = \mathbb{Q}(\sqrt{-4})$ and \mathcal{O}' has conductor 2;
2. $K = \mathbb{Q}(\sqrt{-3})$ and \mathcal{O}' has conductor 2 or 3;
3. The prime 2 splits in K and \mathcal{O}' has index 2 in some order above \mathcal{O} of odd conductor.

Proof. Denote by $S_{\mathcal{O}}$ (resp. $S_{\mathcal{O}'}$) the set of primes ℓ that split into principal ideals in \mathcal{O} (resp. \mathcal{O}'). Using relations formed of a single prime ideal, we see that $\Lambda_{\mathcal{O}} \subseteq \Lambda_{\mathcal{O}'}$ implies $S_{\mathcal{O}} \subseteq S_{\mathcal{O}'}$. Now $S_{\mathcal{O}}$ (resp. $S_{\mathcal{O}'}$) is also the set of primes that split completely in the ring class field $L_{\mathcal{O}}$ of \mathcal{O} (resp. $L_{\mathcal{O}'}$). By Chebotarev's density theorem $S_{\mathcal{O}} \subseteq S_{\mathcal{O}'}$ thus implies $L_{\mathcal{O}'} \subseteq L_{\mathcal{O}}$ which means that the class field theory conductor $f(L_{\mathcal{O}'}/K)$ of $L_{\mathcal{O}'}$ divides $f(L_{\mathcal{O}}/K)$.

This conductor $f(L_{\mathcal{O}}/K)$ is related to that $f_{\mathcal{O}}$ of \mathcal{O} in the following manner (see Exercises 9.20–9.23 of COX (1989)).

$$f(L_{\mathcal{O}}/K) = \begin{cases} \mathcal{O}_K, & \text{when } K = \mathbb{Q}(\sqrt{-4}) \text{ and } f_{\mathcal{O}} = 2, \\ \mathcal{O}_K, & \text{when } K = \mathbb{Q}(\sqrt{-3}) \text{ and } f_{\mathcal{O}} = 2 \text{ or } 3, \\ f', & \text{when } 2 \text{ splits in } K \text{ and } f_{\mathcal{O}} = 2f' \text{ with } f' \text{ odd,} \\ f_{\mathcal{O}}, & \text{otherwise.} \end{cases}$$

Naturally, the same stands for \mathcal{O}' . In the latter case, the fact that $f(L_{\mathcal{O}}/K)$ divides $f(L_{\mathcal{O}'}/K)$ implies that $f_{\mathcal{O}'}$ divides $f_{\mathcal{O}}$, in other words $\mathcal{O} \subseteq \mathcal{O}'$; the three other cases correspond, in order, to the exceptions listed in the proposition. \square

Intuitively, this means that identifying orders by their Picard groups has a single blind spot locally at 2 and 3 where the two largest orders cannot be distinguished.

For orders in higher-degree number fields, we were unable to prove a similar result, but have observed that pairs of orders with identical Picard group structure follow a similar pattern to what the proposition above describes for imaginary quadratic orders; therefore, we will assume:

Assumption VII.1.2. *Fix $g \in \mathbb{N}$; there exists an integer B such that, if any two orders \mathcal{O} and \mathcal{O}' of a complex multiplication field K of degree $2g$ have identical Picard group structure, then one is contained in the other with index a divisor of B , and both orders are maximal at all primes but the factors of B .*

For instance, in the case of quartic complex multiplication fields, our computations support

$$B = 2^6 \cdot 3^4 \cdot 5^3 \cdot 7^2 \cdot 11^2 \cdot 13 \cdot 17 \cdot 19 \cdot 23 \cdot 31 \cdot 41 \cdot 83 \cdot 127 \cdot 131 \cdot 151$$

This bound B could be reduced by excluding finitely many number fields.

Even if this assumption turns out to be wrong, our algorithms will still be functional as they do not need to know in advance which orders have the same Picard group structure: it can always be tested, as we ascend the lattice of orders and generate certificates, if an order has the same Picard group structure as some order directly above or below it. This is naturally quite expensive, but retains the unconditional correctness of our output.

LOCAL WORKAROUND

As we have seen, two distinct orders of a complex multiplication field K can have identical Picard group structure, in a limited number of cases. Those orders cannot be distinguished

using the complex multiplication action, so we need another method to tell them apart from each other.

To tackle these cases, we apply our lattice-ascending and order-testing procedures normally and fall back on a second method when the endomorphism ring is found to be one of these. This amounts to ascending the lattice of orders quotiented by classes of orders with identical Picard group structure; when the class of $\text{End}(\mathcal{A})$ is identified, we determine precisely which order $\text{End}(\mathcal{A})$ is using the following algorithm.

Algorithm VII.1.3.

INPUT: A simple ordinary abelian variety \mathcal{A} over the finite field with q elements, and an order \mathcal{O} with the same Picard group structure as $\text{End}(\mathcal{A})$.

OUTPUT: An order isomorphic to $\text{End}(\mathcal{A})$.

1. Compute the Frobenius polynomial $\chi_\pi(x)$, and factor $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ as $\prod \ell^{v_\ell}$.
2. For all prime factors ℓ with $\ell^{2q^{v_\ell}} < L(|\Delta|)$:
3. Determine $\text{End}(\mathcal{A})$ locally at ℓ .
4. For other prime factors ℓ :
5. Compute various ℓ -isogenies and see if they change the Picard group structure of the endomorphism ring.
6. Deduce $\text{End}(\mathcal{A})$.

The condition in Step 2 ensures that the complexity of determining the endomorphism ring locally at ℓ via the method of EISENTRÄGER and LAUTER (2009) in Step 3 is bounded subexponentially. Basically, since orders with identical Picard group structure only differ by smooth indices (as we saw in the previous section), only small primes ℓ will be of interest here (for others, \mathcal{O}_ℓ is the only possibility for $\text{End}(\mathcal{A})_\ell$); for these small primes, the condition means that the depth v_ℓ of the local lattice is not too large.

When v_ℓ is large, this method is too costly. On the other hand, since only the first few top orders have identical Picard group structure, we can compute random chains of ℓ -isogenies and count the minimal number of isogenies it takes to reach a variety whose endomorphism ring has a different Picard group structure (which we determine using our subexponential method). Since we can compute exactly which orders have identical Picard group structure, this gives us some information as to which order our endomorphism ring is.

This is obviously a rather poor approach. Best would be to use a higher-dimensional analog to the method of IONICA and JOUX (2010) and generalize the algorithm of KOHEL (1996) to compute the endomorphism ring locally at ℓ in time $\ell^{O(1)}$ rather than $\ell^{O(v_\ell)}$.

As the complexity of our fall back method depends not only on the prime ℓ at which we want to locally compute $\text{End}(\mathcal{A})$, but on the entire factor ℓ^{v_ℓ} of the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$, and we found no satisfying way of patching it, we simply rule out deep lattices.

Assumption VII.1.4. Let $\mathcal{O} \subset \mathcal{O}'$ be two orders containing $\mathbb{Z}[\pi, \bar{\pi}]$ with identical Picard group structures. If ℓ is a prime factor of the index $[\mathcal{O}' : \mathcal{O}]$, we assume that the valuation v_ℓ of $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ at ℓ is such that $\ell^{2g_{v_\ell}} < L(q)$.

In dimension one, the method of KOHEL (1996) computes $\text{End}(\mathcal{A})_\ell$ locally at ℓ by climbing the ℓ -isogeny volcano in time $v_\ell \ell^{2+o(1)}$, so the assumption above is not required in that case.

VII.2 Picard Groups from Relations

RELATION SETTING

We recall the standard “generator and relations” setting based on prime ideals to study the structure of Picard groups of orders in number fields.

Throughout this section, \mathcal{O} will be an order in an algebraic number field, and \mathfrak{B} a generating set of ideals for its Picard group; for computational reasons we assume that \mathfrak{B} consists of prime ideals. We denote by $\Lambda_{\mathcal{O}}$ the lattice of relations amongst elements of \mathfrak{B} seen as vectors of $\mathbb{Z}^{\mathfrak{B}}$, so that we have

$$\text{Pic}(\mathcal{O}) \simeq \mathbb{Z}^{\mathfrak{B}} / \Lambda_{\mathcal{O}}.$$

Our first task will be to bound the norm of primes contained in \mathfrak{B} ; this is the purpose of the following section which describes various Chebotarev theorems that have been used over the years — this application being just one specific use of them.

Next, we will consider bounding the diameter of the lattice $\Lambda_{\mathcal{O}}$ which plays a crucial role in the generation of relations that characterizes \mathcal{O} . More explicitly, HAFNER and MCCURLEY (1989) proved that any bound on the diameter of the lattice $\Lambda_{\mathcal{O}}$ yields a box B whose pushforward distribution by $\sigma_{\mathcal{O}}$ is quasi-uniform; in other words, products of random elements of this box give quasi-random elements of the Picard group of \mathcal{O} .

This property is crucial to ensure that the relations we obtain permit us to distinguish a lattice from strictly smaller ones.

Originally, a bound elementarily derived from the theorem of SIEGEL (1935) was used by HAFNER and MCCURLEY (1989); later, BUCHMANN (1989) adapted their algorithm to general number fields, therefore relying on the theorem of BRAUER (1947). We will here give, as a consequence of the generalized Riemann hypothesis, a better bound which we will derive from a more general result of JAO, MILLER, and VENKATESAN (2009).

CHEBOTAREV THEOREMS

Let us first recall the classical *density theorem* of TSCHEBOTAREFF (1926).

Theorem VII.2.1. *Let L/K be a finite normal extension of number fields, and denote by $\pi(\mathfrak{p})$ the Frobenius element in $\text{Gal}(L/K)$ which corresponds to a given prime \mathfrak{p} of K . Such Frobenius elements are asymptotically uniformly distributed in the sense that, for any conjugacy class \mathcal{C} of the Galois group,*

$$\#\{\mathfrak{p} : \pi(\mathfrak{p}) \in \mathcal{C}, N(\mathfrak{p}) < x\} \underset{x \rightarrow \infty}{\sim} \frac{\#\mathcal{C}}{\#\text{Gal}(L/K)} \text{Li}(x)$$

where $\text{Li}(x) = \int_2^x \frac{dt}{\log t}$ is asymptotically equal to the number of prime ideals of norm less than x .

This theorem has countless applications; for instance, if L is the splitting field of a polynomial $f \in K[x]$, it gives the density of primes \mathfrak{p} of K modulo which f has prescribed splitting patterns.

In our setting, we are mostly interested in the case where $K = \mathbb{Q}$ and L is the ring class field $\mathcal{H}_{\mathcal{O}}$ of an order \mathcal{O} in some complex multiplication number field. Via the Artin map, the Chebotarev density theorem descends to ideals of the order \mathcal{O} and asserts that the density of prime ideals which belong to a prescribed ideal class of $\text{Pic}(\mathcal{O})$ is $1/\#\text{Pic}(\mathcal{O})$; this implies in particular that each ideal class can be represented by a prime ideal, from which we can conclude that it is indeed possible to have a generating set \mathcal{B} for $\text{Pic}(\mathcal{O})$ made of prime ideals.

More generally, so-called *effective Chebotarev theorems* give upper bounds on elements generating number theoretic groups. Historically, interest first lied in bounding the least quadratic non-residue modulo n : GAUSS first established the bound $2\sqrt{n} + 1$ (for $n > 2$) elementarily and, to date, the best known unconditional bound of BURGESS (1957) is still exponential — the proof mixes arguments of VINOGRADOV (1919) with the Hasse–Weil bound on the number of points of hyperelliptic curves.

Assuming the Riemann hypothesis for the zeta function of certain fields L , more precise results can be derived. Most often, authors simply assume the extended Riemann hypothesis (ERH), or even the generalized Riemann hypothesis (GRH) for convenience. Under this assumption, ANKENY (1952) proved that the bound above can be made $O(\log^2 n)$.

LAGARIAS and ODLYZKO (1977) later generalized this to general number fields: they proved that if L is a finite nontrivial extension of an algebraic number field K , the least prime ideal of K that does not split completely in L is bounded by $O(\log^2(\text{disc}(K)^2 N(f(L/K))))$.

BACH (1990) gave explicit constants O for these results: he showed that in the result of ANKENY (1952) we have $O \leq 2$, and that $O \leq 3$ for the generalized result. He derived the following:

Theorem VII.2.2. *Assuming the Riemann hypothesis for the zeta function of the number field K , its Galois group $\text{Gal}(K/\mathbb{Q})$ is generated by the Frobenius elements of its prime ideals of norm less than $12 \log^2 |\text{disc}(K)|$.*

DISTRIBUTION OF SHORT PRODUCTS

As we have already pointed out, knowing that the set \mathfrak{B} of prime ideals of norm less than $12 \log^2 |\Delta|$ generates the Picard groups of orders \mathcal{O} containing $\mathbb{Z}[\pi, \bar{\pi}]$ is not sufficient. Indeed, evaluating isogenies associated to ideals \mathfrak{a} which involve large exponents is costly, so it is not sufficient to write \mathfrak{a} as a product of primes of \mathfrak{B} : we also want this product to be short. In other words, we ask that $\mathfrak{a} = \prod_{\mathfrak{p} \in \mathfrak{B}} \mathfrak{p}^{n_{\mathfrak{p}}}$ for a small exponent vector n .

Obviously, its norm $\|n\|_1 = \sum |n_{\mathfrak{p}}|$ is less than the class number. In their Lemma 1, HAFNER and MCCURLEY (1989) proved that any bound on the diameter of the lattice $\Lambda_{\mathcal{O}}$ yields a box B suitable to search for relations, and as a bound they used the latter elementary result on the norm of n . BUCHMANN (1989) did the same in his Lemma 3.4 for arbitrary orders.

However, assuming the generalized Riemann hypothesis, a much better bound can be derived from Corollary 1.3 of JAO, MILLER, and VENKATESAN (2009), which implies:

Theorem VII.2.3 (GRH). *For all $g \in \mathbb{N}$ and $\varepsilon > 0$, there exists $c > 1$ such that, if \mathcal{O} is an order of dimension $2g$ and discriminant Δ , then for random vectors x drawn from the box*

$$B = \left\{ x \in \mathbb{Z}^{\{\mathfrak{p}: N(\mathfrak{p}) < \log^{2+\varepsilon} |\Delta|\}} : \sum_{\mathfrak{p}} |x_{\mathfrak{p}}| = c \frac{\log |\Delta|}{\log \log |\Delta|} \right\}$$

the probability that $\sigma_{\mathcal{O}}(x)$ falls in any fixed ideal class of $\text{Pic}(\mathcal{O})$ is at least $1/2\#\text{Pic}(\mathcal{O})$.

In terms of distribution, this states that the pushforward distribution by $\sigma_{\mathcal{O}}$ of the uniform distribution \mathbb{U}_X on the set X of vectors of norm $c \log |\Delta| / \log \log |\Delta|$ is within variation distance $1/2$ from the uniform distribution on the Picard group. Essentially, this says that products of randomly selected primes of quadratic norm behave as uniformly-drawn elements of the Picard group.

DIAMETER OF RELATION LATTICES

The above theorem implies that each element of $\text{Pic}(\mathcal{O})$ has a preimage of small norm, from which we can easily derive a bound on the diameter of $\Lambda_{\mathcal{O}}$. Recall that the *diameter* of a lattice is the smallest value $\text{diam}(F)$ where F ranges over its fundamental domains.

Corollary VII.2.4 (GRH). *Fix any positive number ε . If \mathcal{O} is an order of discriminant Δ and \mathfrak{B} denotes its set of primes of norm less than $\log^{2+\varepsilon} |\Delta|$, the diameter of the lattice $\Lambda_{\mathcal{O}}$ is $o(\log^{4+\varepsilon} |\Delta|)$.*

Proof. To prove this, we construct a generating set for $\Lambda_{\mathcal{O}}$ formed by $O(\log^{2+\varepsilon} |\Delta|)$ relations of norm $o(\log^2 |\Delta|)$. BRAUER (1947) showed that $\text{Pic}(\mathcal{O})$ is an abelian group of order $\Delta^{1/2+o(1)}$ so there exist $O(\log |\Delta|)$ ideal classes α_i such that $\mathbb{Z}^{\mathfrak{B}}/\Lambda_{\mathcal{O}} \simeq \prod \langle \alpha_i \rangle$; we fix these and proceed to write a generating set for $\Lambda_{\mathcal{O}}$ consisting of:

- relations expressing that $\alpha_i^{\text{ord}(\alpha_i)} = 1$;
- relations expressing the primes $\mathfrak{p} \in \mathfrak{B}$ in terms of the α_i .

First define a map $\sigma_{\mathcal{O}}^{-1}$ by fixing a preimage of norm at most $c \log |\Delta| / \log \log |\Delta|$ for each ideal class; it exists by Theorem VII.2.3. Now use a double-and-add approach to ensure that norms remain small: for each i , express that $\alpha_i^{\text{ord}(\alpha_i)} = 1$ by the relations

$$(i) \quad \sigma_{\mathcal{O}}^{-1}(\alpha_i^{2^j}) - 2\sigma_{\mathcal{O}}^{-1}(\alpha_i^{2^{j-1}}) \text{ for } j \in \{1, \dots, \lfloor \log_2 \text{ord}(\alpha_i) \rfloor\};$$

$$(ii) \quad \sum_j b_j \sigma_{\mathcal{O}}^{-1}(\alpha_i^{2^j}) \text{ where } b_j \text{ denotes the } j^{\text{th}} \text{ least significant bit of } \text{ord}(\alpha_i).$$

Now write each $\mathfrak{p} \in \mathfrak{B}$ on the α_i by decomposing its class as a product $\prod \alpha_i^{n_i}$ where $n_i \in \{0, \dots, \text{ord}(\alpha_i)\}$; noting $\delta_{\mathfrak{p}}$ the vector with coordinate one at \mathfrak{p} and zero elsewhere, this gives the relations:

$$(iii) \quad \delta_{\mathfrak{p}} - \sum_i \sum_j c_{ij} \sigma_{\mathcal{O}}^{-1}(\alpha_i^{2^j}) \text{ where } c_{ij} \text{ is the } j^{\text{th}} \text{ least significant bit of } n_i.$$

Preimages by $\sigma_{\mathcal{O}}$ have length $o(\log |\Delta|)$ and there are at most $\sum \lfloor \log_2 \text{ord}(\alpha_i) \rfloor = O(\log |\Delta|)$ terms, therefore each such relation has length $o(\log |\Delta|)^2$. \square

VII.3 Relations from Smooth Ideals

Let us now give the mathematical background required to prove the complexity of the subexponential method for finding smooth relations in Picard groups.

INTEGER SMOOTHNESS

We start by reviewing fundamental properties of smooth numbers; these are the base on which most subexponential algorithms are built upon (for instance, we have already mentioned factoring algorithms). First recall their definition.

Definition VII.3.1. *An integer x is said to be y -smooth if it has no prime factor larger than y . The number of y -smooth integers less than x is denoted $\Psi(x, y)$.*

Bounding the value of the Ψ function for particular ranges of x and y is an important problem. For instance, for any fixed $u \geq 1$, we have

$$\Psi(x, x^{1/u}) \underset{x \rightarrow \infty}{\sim} x \rho(u)$$

where the constant $\rho(u)$ is the Dickman function. This function was extensively studied by DE BRUIJN who gave many ways to evaluate it. To use such smoothness results in index-calculus methods, we need more than a polynomial relation of the form $y = x^{1/u}$: we would like to consider the case where $u \rightarrow \infty$ as $x \rightarrow \infty$. The specific result we rely on is due to CANFIELD, ERDŐS, and POMERANCE (1983).

Theorem VII.3.2. *For $u \geq 3$ we have*

$$\Psi(x, x^{1/u}) \geq x \exp(-u(\log u + \log \log u - 1 + o(1)))$$

Corollary VII.3.3. *The probability for a random number of $\{1, \dots, x\}$ to be $L(x)^\gamma$ -smooth is equivalent to $L(x)^{-1/2\gamma+o(1)}$ as $x \rightarrow \infty$.*

Proof. Apply the theorem above to $u = \frac{1}{\gamma} \sqrt{\frac{\log x}{\log \log x}}$ and combine it with the upper bound in Theorem 2 of BRUIJN (1966). \square

See GRANVILLE (2008) for a survey of this topic.

IDEAL SMOOTHNESS

Our algorithms do not exactly work with integers: they work with ideals. Via the norm, the structure of the ring of ideals resembles that of integers; for our particular goal, it suffices to say that ideals are smooth if and only if their norms are. However, not all results are easy to generalize from integers to ideals.

In fact, our first algorithm for computing endomorphism rings of elliptic curves, from B. and SUTHERLAND (2009), relied on the assumption that certain ideals we generated had a uniformly distributed norm, so that we could directly apply the result of the previous section. We now explain how this assumption can, in some setting, be rigorously proven.

Let us first recall the relevant part of our algorithm: for an order \mathcal{O} of discriminant Δ , we first select a vector x uniformly at random from the box $B = \{0, \dots, \log^{4+\epsilon} |\Delta|\}^{\mathfrak{B}}$ where \mathfrak{B} is the set of prime ideals of norm less than $\log^{2+\epsilon} |\Delta|$; we then look for a small representative \hat{x} of the class $\sigma_{\mathcal{O}}(x) \in \text{Pic}(\mathcal{O})$ and attempt to factor it over the base consisting of all the prime ideals of norm less than $L(|\Delta|)^\gamma$.

To rigorously bound the number of times random vectors $x \in B$ have to be selected before one with smooth reduction is found, we need to show that the norm of \hat{x} behaves like a random integer in a certain interval.

For imaginary quadratic orders, SEYSEN (1987) used the standard reduction of binary quadratic forms; to obtain a result on the smoothness probability of \hat{x} , he proceeds in two steps: Proposition 4.3 and 4.4:

Proposition VII.3.4. *Ideal classes $\sigma_{\mathcal{O}}(x)$ of randomly selected vectors $x \in \mathbb{B}$ are quasi-uniformly distributed in the Picard group of \mathcal{O} .*

By *quasi-uniformly distributed*, we mean that the probability for $\sigma_{\mathcal{O}}(x)$ to belong to a prescribed subset S of $\text{Pic}(\mathcal{O})$ is

$$(1 + o(1)) \frac{\#S}{\#\text{Pic}(\mathcal{O})}$$

in other words, the pushforward distribution $\sigma_{\mathcal{O}*} \mathbb{U}_{\mathbb{B}}$ is within variation distance $o(1)$ of the uniform distribution on $\text{Pic}(\mathcal{O})$.

Note that SEYSEN (1987) started from a much bigger box \mathbb{B} than ours; it was, back then, the best possible under the generalized Riemann hypothesis; however, here, we make use of Corollary 1.3 of JAO, MILLER, and VENKATESAN (2009) and of the smaller box \mathbb{B} it proves to suffice.

When we know that $\sigma_{\mathcal{O}}(x)$ is quasi-random, it remains to see whether the element \widehat{x} of each $\sigma_{\mathcal{O}}(x)$ has a smoothness probability comparable to integers of $\{1, \dots, \sqrt{|\Delta|}/3\}$.

Proposition VII.3.5. *The number of reduced ideals whose norm is $L(|\Delta|)^{\gamma}$ -smooth is at least $n/L(|\Delta|)^{1/2\gamma+o(1)}$ where $n = \#\text{Pic}(\mathcal{O})$ is the total number of reduced ideals.*

The proof of SEYSEN (1987) involves calculations which are specific to the arithmetic of binary quadratic forms. This makes it challenging to generalize this proposition in higher-dimensional orders, and another issue is that there is no canonical notion of reduction there. The method of BUCHMANN (1989) for arbitrary orders relies on the following assumption, and we do as well.

Assumption VII.3.6. *The norms of reduced ideals used by the smooth relation finding algorithm are as likely to be smooth as random integers of $\{1, \dots, \sqrt{|\Delta|}\}$.*

RANDOMNESS OF RELATIONS

To obtain a generating set for the lattice $\Lambda_{\mathcal{O}}$ by finding relations of it, we must ensure that those relations do not lie in some particular subset. For instance, if the order \mathcal{O} contains \mathcal{O}' , then we have $\Lambda_{\mathcal{O}'} \subset \Lambda_{\mathcal{O}}$, and we must prove that our relations have no predisposition of actually lying in $\Lambda_{\mathcal{O}'}$. Whence the following definition.

Definition VII.3.7. *Let P be a probabilistic procedure which, on input an order \mathcal{O} containing $\mathbb{Z}[\pi, \overline{\pi}]$ for some Weil number π , returns a relation $x \in \Lambda_{\mathcal{O}}$, which we see as a random variable.*

We say that P generates quasi-uniformly distributed relations of \mathcal{O} if, for any order \mathcal{O}' containing $\mathbb{Z}[\pi, \overline{\pi}]$, the projection of x in the quotient group $\Lambda_{\mathcal{O}}/\Lambda_{\mathcal{O} \cap \mathcal{O}'}$ is within variation distance $o(1)$ from the uniform distribution, as the discriminant of π goes to infinity.

Proving that the method of SEYSEN (1987) does indeed generate quasi-uniformly distributed relations was done by HAFNER and MCCURLEY (1989) in their Lemma 2.

Proposition VII.3.8. *If \mathcal{O}' is an order contained in \mathcal{O} , relations found by the method of SEYSEN (1987) are quasi-uniformly distributed in $\Lambda_{\mathcal{O}}/\Lambda_{\mathcal{O}'}$ when $B = \{0, \dots, \#\mathfrak{B}d^{1+\varepsilon}\}^{\mathfrak{B}}$, where d is a bound on the diameter of $\Lambda_{\mathcal{O}}$.*

The proof is pretty simple and involves looking at the geometry of the lattices in a fairly elementary way. We reproduce it below, in the more general context of an unspecified bound d on $\text{diam } \Lambda_{\mathcal{O}}$.

Proof. Let x be a random variable with uniform distribution on $B_t = \{0, \dots, t\}^{\mathfrak{B}}$, let $\hat{x} \in \sigma_{\mathcal{O}}(x)$ denote its reduction, and note \mathcal{S} the set of ideals with \mathcal{L} -smooth norms. We want to prove that

$$\text{Prob} [x - \sigma^{-1}(\hat{x}) \in \omega | \hat{x} \in \mathcal{S}] = [\Lambda_{\mathcal{O}} : \Lambda_{\mathcal{O}'}]^{-1} (1 + o(1))$$

for any fixed class $\omega \in \Lambda_{\mathcal{O}}/\Lambda_{\mathcal{O}'}$. We can rewrite the left-hand side as

$$\frac{\#\{x \in B_t : \hat{x} \in \mathcal{S}, x - \sigma^{-1}(\hat{x}) \in \omega\}}{\#\{x \in B_t : \hat{x} \in \mathcal{S}\}}$$

and by summing over all possible reduced ideals y we further obtain

$$\frac{\sum_{y \in \mathcal{S}} \#\{x \in B_t : x \in \sigma^{-1}(y) + \omega\}}{\sum_{y \in \mathcal{S}} \#\{x \in B_t : x \in \sigma^{-1}(y) + \Lambda_{\mathcal{O}}\}}.$$

Now, to evaluate each term of these sums, let us count the number of points of $\bar{B}_t = [0, t+1]^{\mathfrak{B}}$ which lie in the translation $z + \Lambda$ of some lattice Λ . To this extent, let \mathcal{F} be a fundamental domain for Λ : each point of $z + \Lambda$ corresponds to a cell in the tiling of $\mathbb{R}^{\mathfrak{B}}$ by \mathcal{F} ; if $\text{diam } \mathcal{F} \leq d$ we therefore have

$$\bar{B}_{t-d} \subset (z + \Lambda) \cap \bar{B}_t + \mathcal{F} \subset \bar{B}_{t+d}$$

which gives, in terms of volumes,

$$(t-d)^{\#\mathfrak{B}} \leq \det \Lambda \cdot \#((z + \Lambda) \cap B_t) \leq (t+d)^{\#\mathfrak{B}}$$

so as soon as $\#\mathfrak{B}d = o(t)$, the sandwich theorem proves that

$$\#((z + \Lambda) \cap B_t) = \frac{t^{\#\mathfrak{B}}}{\det \Lambda} (1 + o(1));$$

by substituting this in the probability sought expressed as a quotient of sums, we obtain

$$\# \mathcal{S} \frac{t^{\#\mathfrak{B}}}{\det \Lambda_{\mathcal{O}'}} (1 + o(1)) \bigg/ \# \mathcal{S} \frac{t^{\#\mathfrak{B}}}{\det \Lambda_{\mathcal{O}}} (1 + o(1)) ;$$

Choosing $t = \#\mathfrak{B}d^{1+\varepsilon}$ satisfies the requirement $\#\mathfrak{B}d = o(t)$ and gives the result. □

Recall that if \mathcal{O} is an order of discriminant Δ and \mathfrak{B} consists of all prime ideals of norm less than $\log^{2+\varepsilon} |\Delta|$, then the diameter of $\Lambda_{\mathcal{O}}$ is $o(\log^{4+\varepsilon} |\Delta|)$. Therefore, when \mathcal{O} is imaginary quadratic, the above proposition shows that the algorithm of SEYSEN (1987) generates quasi-uniformly distributed relations of $\Lambda_{\mathcal{O}}$ when drawing its random vectors uniformly from the box $B = \{0, \dots, \log^{2+\varepsilon} |\Delta|\}^{\#\mathfrak{B}}$.

When \mathcal{O} is an order in a complex multiplication of degree four or more, as we have mentioned before, we do not know of similar results and believe that they might be quite difficult to establish. However, we can still amend the algorithm of BUCHMANN (1989) to make use of this type of bound. This gives a conjectural running time, but the result can in any case be unconditionally proven by certificates, so we have a Las-Vegas algorithm.

GENERATING ENOUGH RELATIONS

To prepare for the jump to the next chapter, let us put together the results that we have established so far. Here, we let π be the Frobenius endomorphism of an abelian variety of dimension g defined over a finite field \mathbb{F}_q , and recall from Lemma III.2.5 that $\text{disc}(\mathbb{Z}[\pi, \bar{\pi}]) = q^{g^2+o(1)}$ so that via the theorem of BRAUER (1947) the class number is $q^{g^2/2+o(1)}$.

Proposition VII.3.9. *Let \mathcal{O} be an order of discriminant Δ in a number field of degree $2g$; random relations of \mathcal{O} involving polynomially many ideals in $\log |\Delta|$ of norm up to $L(|\Delta|)^\gamma$ can be found in probabilistic time $L(|\Delta|)^\gamma + L(|\Delta|)^{1/4\gamma+o(1)}$.*

This assumes the generalized Riemann hypothesis for $g = 1$, and Assumption VII.3.6 for $g > 1$.

Unlike HAFNER and MCCURLEY (1989), we do not seek to compute the full group structure of $\text{Pic}(\mathcal{O})$ — this would be costly since a subexponential number of relations is required to eliminate all factors of the factor base. Here, we just aim at distinguishing orders containing $\mathbb{Z}[\pi, \bar{\pi}]$ from one another.

If \mathcal{O}' is an order such that $\Lambda_{\mathcal{O}'}$ is strictly contained in $\Lambda_{\mathcal{O}}$, a quasi-uniformly distributed relation has probability at most $1/2 + o(1)$ of also holding in \mathcal{O}' . Therefore, since we have a polynomial number of orders in $\log |\Delta|$ to discriminate from, it is sufficient to only generate polynomially many orders in $\log \log |\Delta|$ to ensure that the relations characterize the lattice $\Lambda_{\mathcal{O}}$ with probability $1 - o(1)$.

Combining the above with our earlier notes on the complexity of isogeny computation, we have proved the following.

Theorem VII.3.10. *Let \mathcal{A} be a simple ordinary abelian variety of dimension g defined over the finite field with q elements. Under the generalized Riemann hypothesis, we can compute $\text{End}(\mathcal{A})$:*

- if $g = 1$, in $L(q)^{1+o(1)} + L(q)^{1/\sqrt{2}+o(1)}$ operations;
- if $g = 2$, in $L(q)^g \sqrt{3g/2+o(1)}$ operations, under Assumptions VII.1.2, VII.1.4, VII.3.6, and VIII.1.1.

For $g = 2$, details will be given in the next chapter.

VII.4 Relations from Thin Air

As a supplement to this chapter, we shall now see how to generate relations in a generic manner, that is, not using any extrinsic information about the underlying group. For Picard groups, such methods are much slower than smoothness-based ones but yield much shorter relations; this will be an important ingredient for making practical use of our method in dimension two.

GENERIC SHORT PRODUCTS

Let S be a sequence of elements in a finite group G of order n , written multiplicatively, and consider the problem of writing a prescribed element $z \in G$ as the product of a subsequence of S ; we call such a subsequence a *short product representation* of z on S .

If G were a commutative group, we could have noted it additively, let S be a multiset of elements of it, and look for a sub-multiset which adds up to z ; in the case that S has no repeated elements, this is known as the *subset sum problem*. However, since for our approach it makes absolutely no difference whether G is commutative, we have chosen to use the more general formalism of non-necessarily-commutative groups.

Consider the product map $\pi : \mathfrak{P}(S) \rightarrow G$ where $\mathfrak{P}(S)$ denotes the set of all subsequences of S . For all elements of G to admit short product representations, the map π needs to be surjective which, by a counting argument, implies $k \geq \log_2 n$ where k is the length of S .

In the case that G is commutative, ERDŐS and RÉNYI (1965) showed that this bound is not far from being sufficient: they prove that a random sequence S of length

$$k = \log_2 n + \log_2 \log n + \omega_n$$

satisfies $\pi(\mathfrak{P}(S)) = G$ with probability approaching 1 as $n \rightarrow \infty$, provided that $\omega_n \rightarrow \infty$.

For finding short product representations via generic means, just knowing the existence of a preimage by π for all $z \in G$ is not enough: we need to know the distribution of such preimages. IMPAGLIAZZO and NAOR (1996) proved the following result on the inverse distribution.

Theorem VII.4.1. *Fix some real number d . For groups G of order n large enough, we have*

$$\text{Prob}_S \left[\left\| \pi_* \mathbb{U}_{\mathfrak{P}(S)} - \mathbb{U}_G \right\| \geq n^{-c} \right] \leq n^{-c}$$

where $c = (d-1)/2$ and the sequence S is drawn uniformly at random from the set of sequences of G with length $k = (d + o(1)) \log_2 n$.

Recall that \mathbb{U}_X denotes the uniform distribution on the (finite) set X , and that the *push-forward distribution* $f_*\sigma$ of a distribution σ on X by a function $f: X \rightarrow Y$ is defined as

$$f_*\sigma(y) = \sigma(\{x \in X : f(x) \in y\}),$$

for any subset y of Y . Finally, the *variation distance* $\|\sigma - \sigma'\|$ between two distributions on Y is the maximum value of $|\sigma(y) - \sigma'(y)|$ as y ranges over all subsets of Y .

In other words, the theorem means that, for a random sequence S of *density* $d > 1$, the distribution of subsequence products almost surely converges to the uniform distribution on G as n goes to infinity.

In some particular cases, finding short product representations is a well-known problem. For instance, when G is the Picard group of some order and S contains all prime powers p^a with $p < L(|\Delta|)$ and $a < \log_p |\Delta|$, then this is exactly the problem of finding relations which we have studied extensively. Now this problem does not have a “constant” density, as the quantity $k/\log_2 n$ goes to infinity pretty quickly with n .

For instances of constant density in the group $G = \mathbb{Z}/n\mathbb{Z}$, the best algorithm has a time and space complexity of $O(n^{0.3113})$; it consists in lifting the instance to k subset sum problems in \mathbb{Z} , also known as knapsack problems, which can be solved efficiently by a method of HOWGRAVE-GRAHAM and JOUX (2010). Again, this algorithm is tailored for a specific group representation.

Algorithms that only perform multiplications and inversions (which return uniquely identified group elements), draw elements at random from G , and test their equality, are called *generic algorithms*. In essence, they are not tied to any specific group and apply to any effective group. SHOUP (1997) proved that solving discrete logarithm problems generically has a lower bound of $\Omega(\sqrt{p})$ where p is the largest prime factor of n ; since this is a special case of short product representation, this means that generic short product representation algorithms cannot have a faster-than-square-root complexity in the worst case.

BABY-STEP GIANT-STEP

Let us first review classical approaches to the problem of finding a short product representation of an element $z \in G$ on a sequence S .

A brute-force algorithm would exhaustively enumerate the set $\mathfrak{P}(S)$ and for each element y of it test whether $\pi(y) = z$.

The standard baby-step giant-step approach splits the search space as a direct product $\mathfrak{P}(S) = \mathfrak{P}(A) \times \mathfrak{P}(B)$ simply by writing S as the concatenation of two smaller sequences A and B ; then, it aims at finding a pair of elements $(x, y) \in \mathfrak{P}(A) \times \mathfrak{P}(B)$ which *collide* in the sense that $\pi(x) = z\pi(y)^{-1}$. This can be implemented efficiently by first precomputing and storing a table of all $\pi(x)$ for $x \in \mathfrak{P}(A)$, and then checking whether each $z\pi(y)^{-1}$ for $y \in \mathfrak{P}(B)$ is in this table; the lookup can be done in time $O(\log n)$ using an efficient data structure.

For convenience, we define an application μ which maps any sequence $y = (y_1, \dots, y_m)$ to $\mu(y) = (y_m^{-1}, \dots, y_1^{-1})$, so that $\pi(y)$ and $\pi(\mu(y))$ are inverses in G . The baby-step giant-step algorithm then amounts to the following procedure.

Algorithm VII.4.2.

INPUT: A finite sequence S and a target $z \in G$.

OUTPUT: If it exists, a subsequence of S whose product is z .

1. Split S as a concatenation AB of sequences of roughly equal sizes.
2. For each $x \in \mathfrak{P}(A)$, store x in a table indexed by $\pi(x)$.
3. For each $y \in \mathfrak{P}(B)$:
4. If $\pi(z\mu(y)) = \pi(x)$ for some x , then return xy .
5. Return that z has no preimage by π in $\mathfrak{P}(S)$.

As each element of $\mathfrak{P}(A)$ can be represented by $k/2$ bits (which is a constant factor away from the size of a group element, when the density d is fixed), the total memory consumed by this algorithm is $O(2^{k/2})$. By enumerating elements of $\mathfrak{P}(A)$ and $\mathfrak{P}(B)$ in a suitable order (for instance, using a Gray code), only one group operation is required per step, so that the total runtime is $O(2^{k/2})$.

SCHROEPEL and SHAMIR (1981) gave a more specialized generic method for solving knapsack problems, which improves the space complexity of the baby-step giant-step algorithm to $O(2^{k/4})$.

POLLARD RHO

In order to apply the Pollard ρ approach to the problem of finding short product representations, we simply need a notion of collision on a certain domain \mathcal{C} and an iteration map

$\phi : \mathcal{C} \rightarrow \mathcal{C}$ which preserves collisions in the sense that if x and y collide, then $\phi(x)$ and $\phi(y)$ also collide.

Here, we use the same domain that was used by the baby-step giant-step algorithm: split S as a concatenation AB of two sequences of roughly equal size, and let the domain be the disjoint union $\mathcal{C} = \mathcal{A} \sqcup \mathcal{B}$ where $\mathcal{A} = \mathfrak{P}(A)$ and $\mathcal{B} = z\mu(\mathfrak{P}(B))$. Now collisions are defined with respect to the product map $\pi : \mathcal{C} \rightarrow G$; when an element $x \in \mathcal{A}$ collides with an element $y \in \mathcal{B}$, that is, $\pi(x) = \pi(y)$, then we have a short product representation of z as xy' where $y' = z\mu(y')$.

Now since the iteration map ϕ must respect collisions, it must factor through the product map π so we can write $\phi = \eta \circ \pi$ for some $\eta : G \rightarrow \mathcal{C}$. Since we have no requirement on η , we simply take it to be a hash function from G to \mathcal{C} , that is, an effective map which behaves as if it were drawn uniformly at random from \mathcal{C}^G .

In practice, to compute $\eta(g)$ we can take the unique bit-string representation of g , hash it using a strong cryptographic hash function, and use the resulting bit-string $g_0g_1g_2\dots$ to dictate an element of \mathcal{C} ; for instance, the first bit g_0 can be used to decide whether $\phi(g)$ lies in $\mathfrak{P}(A)$ or $z\mu(\mathfrak{P}(B))$, the second bit g_1 to decide whether the first element of A (resp. B) belongs to $\phi(g)$, etc. (Note that η cannot be surjective since G is smaller than \mathcal{C} .)

This gives the following algorithm.

Algorithm VII.4.3.

INPUT: A finite sequence S and a target $z \in G$.

OUTPUT: A subsequence of S whose product is z .

1. Split S as a concatenation AB of sequences of roughly equal sizes.
2. Pick a random element $w \in \mathcal{C}$ and a hash function $\eta : G \rightarrow \mathcal{C}$.
3. Find the least $i > 0$ and $j \geq 0$ such that $\phi^{(i+j)}(w) = \phi^{(j)}(w)$.
4. If $j = 0$ then return to Step 1.
5. Let $s = \phi^{(i+j-1)}(w)$ and let $t = \phi^{(j-1)}(w)$.
6. If $\pi(s) \neq \pi(t)$ then return to Step 1.
7. If $s \in \mathcal{A}$ and $t = z\mu(y) \in \mathcal{B}$ for some y , output sy and terminate.
8. If $t \in \mathcal{A}$ and $s = z\mu(y) \in \mathcal{B}$ for some y , output ty and terminate.

Basically, we start from a random point w and compute iterates $\phi^{(i)}(w)$ until we find two which are equal: once we have the first such collision, that is, $\phi(s) = \phi(t)$ with $s \neq t$, we first make sure it is not due to the hash function, so that the collision must arise in the product map. Then, if it is a collision between an element of \mathcal{A} and one of \mathcal{B} , which happens with expected probability $1/2$, we have a short product representation.

Step 2 can be implemented by Floyd's algorithm, by the method of distinguished points, or any other collision-detection technique (which reduces by a constant factor the number of expected evaluations of the map ϕ before finding a collision).

This gives an algorithm with constant storage space and a time complexity of $O(k\sqrt{n})$. We refer the reader to B. and SUTHERLAND (2011) for a rigorous proof (and also for details regarding this whole section) and now turn to applications.

APPLICATIONS

This method actually has a broad range of applications; in particular, it can be used to find isogenies between two ordinary elliptic curves defined over a finite field having the same endomorphism ring in square-root time and without storage requirements. This application can be found in B. and SUTHERLAND (2011). Here, we will present a different one, maybe not as important, but which applies directly to the topic of computing endomorphism rings.

As usual, we fix an ambient finite base field \mathbb{F}_q and let \mathcal{A} denote a simple ordinary abelian variety. Consider the set G of isomorphism classes of abelian varieties whose endomorphism ring is the same as that of \mathcal{A} ; as we have seen before, it is a principal homogeneous space for the Picard group $\text{Pic}(\text{End } \mathcal{A})$ whose cardinality we denote n (in the worst case, it is exponential in $\log(q)$ and the dimension g of \mathcal{A}).

Our method for computing $\text{End}(\mathcal{A})$ has so far been to compute relations in the Picard group of the possible orders (those that contain $\mathbb{Z}[\pi, \bar{\pi}]$) and checking whether they hold in the isogeny graph. Here, we take the inverse approach: we will look for relations in the isogeny graph, and then rule out from the list of possibilities those orders in which the relations do not hold.

Of course, since the only algorithms we have at our disposal for finding relations in the isogeny graph are generic, this is much slower than looking for relations in Picard groups. However, this gives a runtime which mostly depends on the output: the closer to \mathcal{O}_K the endomorphism ring of \mathcal{A} , the faster it is found.

To look for relations in the isogeny graph of \mathcal{A} , a baby-step giant-step approach is simple to use: let S be a set of prime ideals of \mathcal{O}_K which are coprime to the conductor of $\mathbb{Z}[\pi, \bar{\pi}]$, split it as a concatenation AB , let $\mathcal{A} = \mathfrak{P}(A)$ and $\mathcal{B} = \mathfrak{P}(B)$, and define $\mathcal{C} = \mathcal{A} \sqcup \mathcal{B}$. We view an element $x = (\mathfrak{p}_1, \mathfrak{p}_2, \dots, \mathfrak{p}_m)$ of \mathcal{C} as the isogeny

$$\phi_{\mathfrak{p}_1 \mathfrak{p}_2 \dots \mathfrak{p}_m}(\mathcal{A}) = \phi_{\mathfrak{p}_1} \circ \phi_{\mathfrak{p}_2} \circ \dots \circ \phi_{\mathfrak{p}_m}(\mathcal{A})$$

and we define the map $\pi : \mathcal{C} \rightarrow G$ as sending x to the variety which is the codomain of this isogeny.

Now it is straightforward to adapt the Pollard ρ method to this context as we have done before: it suffices to take a hash function $\eta : G \rightarrow \mathcal{C}$ and to iterate the map $\phi = \eta \circ \pi$ enough times to find a collision. Recall from Chapter III that, in the worst case, we might have

$$\#G = \# \text{Pic}(\text{End } \mathcal{A}) = q^{(1/2+o(1))g^2}$$

so that if we take a sequence S of length at least

$$(d + o(1))g^2 \log_2 q$$

for some $d > 1$, we can effectively find a relation of the isogeny graph in probabilistic time $q^{(1/4+o(1))g^2}$ using virtually no memory, assuming the quasi-uniform distribution of products of S in the Picard group; this assumption can be replaced by the generalized Riemann hypothesis by substituting $\log_2(q)$ by $\log^{2+\varepsilon}(q)$ above, via a result of JAO, MILLER, and VENKATESAN (2009) — note however that this has little effect on the runtime: although the products to be computed have more terms, the collision probability is unchanged.

By finding relations in the isogeny graph of \mathcal{A} , we can test whether a given order \mathcal{O} contains $\text{End}(\mathcal{A})$ in time $\text{disc}(\text{End}(\mathcal{A}))^{1/4+o(1)}$ up to polynomial factors in $\log(q)$ and g . Therefore, locating the endomorphism ring takes just as much time using the “reversed” lattice-ascending procedure of the previous chapter for computing $\text{End}(\mathcal{A})$ from above.

Note that certificates that are generated with such generic methods have a length polynomial in the size of the base field $\log q$, which is much smaller than what subexponential methods can generate. More precisely, this length can essentially be quadratic if we require that the runtime of the generation algorithm be bounded under the generalized Riemann hypothesis (via Theorem VII.2.3), or linear if the heuristic Conjecture VI.4.1 is used instead.

Verifying the certificate then just requires polynomial time in its size: it suffices to verify the number of points on the variety and compute the isogenies associated to the ideals in the relation.

Here again, we have made use of isogenies between isomorphism classes of abelian varieties, not involving any polarizations, which is not an effective notion in dimension $g > 1$. We thus devote the next chapter to describing the changes required for making effective use of our endomorphism computing method on abelian varieties of dimension $g > 1$.

References

1919. Ivan M. VINOGRADOV.
 “On the distribution of quadratic residues and non-residues”.
 In: *Journal of the Physico-Mathematical Society of Perm* 2. Pages 1–16.
1926. Nikolai TSCHEBOTAREFF.
 “Die Bestimmung der Dichtigkeit einer Menge von Primzahlen, welche zu einer gegebenen Substitutionsklasse gehören”.
 In: *Mathematische Annalen* 95.1. Pages 191–228. DOI: 10.1007/BF01206606.

1935. Carl L. SIEGEL.
“Über die Classenzahl quadratischer Zahlkörper”.
In: *Acta Arithmetica* 1. Pages 83–86.
1947. Richard BRAUER.
“On the zeta-functions of algebraic number fields”.
In: *American Journal of Mathematics* 69.2. Pages 243–250.
DOI: 10.2307/2371849.
1952. Nesmith C. ANKENY.
“The least quadratic non residue”. In: *Annals of Mathematics* 55.1. Pages 65–72.
DOI: 10.2307/1969420.
1957. David A. BURGESS.
“The distribution of quadratic residues and non-residues”.
In: *Mathematika* 4. Pages 106–112. DOI: 10.1112/S0025579300001157.
1965. Paul ERDŐS and Alfréd RÉNYI.
“Probabilistic methods in group theory”.
In: *Journal d'Analyse Mathématique* 14.1. Pages 127–138.
DOI: 10.1007/BF02806383.
1966. Nicolaas G. de BRUIJN.
“On the number of positive integers $\leq x$ and free of prime factors $> y$, II”.
In: *Koninklijke Nederlandse Akademie van Wetenschappen. Proceedings, Series A*
69.3. Pages 239–247.
1977. Jeffrey C. LAGARIAS and Andrew M. ODLYZKO.
“Effective versions of the Chebotarev density theorem”.
In: *Algebraic number fields: L-functions and Galois properties*.
Proceedings of a Symposium held at the University of Durham in 1975.
Academic Press. Pages 409–464.
1981. Richard SCHROEPEL and Adi SHAMIR.
“A $T = O(2^{n/2})$, $S = O(2^{n/4})$ algorithm for certain NP-complete problems”.
In: *SIAM Journal of Computing* 10.3. Pages 456–464. DOI: 10.1137/0210033.
1983. Earl CANFIELD, Paul ERDŐS, and Carl POMERANCE.
“On a problem of Oppenheim concerning ‘factorisatio numerorum’”.
In: *Journal of Number Theory* 17.1. Pages 1–28.
DOI: 10.1016/0022-314X(83)90002-1.
1987. Martin SEYSEN.
“A probabilistic factorization algorithm with quadratic forms of negative

- discriminant". In: *Mathematics of Computation* 48.178. Pages 757–780.
DOI: 10.1090/S0025-5718-1987-0878705-X.
1989. Johannes BUCHMANN.
"A subexponential algorithm for the determination of class groups and regulators of algebraic number fields". In: *Séminaire de Théorie des Nombres, Paris*.
Edited by Catherine GOLDSTEIN. Volume 91. Progress in Mathematics. Birkhäuser.
Pages 27–41.
1989. David A. COX.
Primes of the form $x^2 + ny^2$. John Wiley & Sons. ISBN: 0-471-19079-9.
1989. James L. HAFNER and Kevin S. MCCURLEY.
"A rigorous subexponential algorithm for computation of class groups".
In: *Journal of the American Mathematical Society* 2.4. Pages 837–850.
DOI: 10.2307/1990896.
1990. Eric BACH.
"Explicit bounds for primality testing and related problems".
In: *Mathematics of Computation* 55.191. Pages 355–380.
DOI: 10.1090/S0025-5718-1990-1023756-8.
1990. Jonathan PILA.
"Frobenius maps of abelian varieties and finding roots of unity in finite fields".
In: *Mathematics of Computation* 55.192. Pages 745–763.
DOI: 10.2307/2008445.
1992. Hendrik W. LENSTRA and Carl POMERANCE.
"A rigorous time bound for factoring integers".
In: *Journal of the American Mathematical Society* 5.3. Pages 483–516.
DOI: 10.1090/S0894-0347-1992-1137100-0.
1993. Henri COHEN.
A course in computational algebraic number theory. Volume 138.
Graduate Texts in Mathematics. Springer. ISBN: 3-540-55640-0.
1993. Don COPPERSMITH.
"Modifications to the number field sieve".
In: *Journal of Cryptology* 6.3. Pages 169–180. DOI: 10.1007/BF00198464.
1996. Russel IMPAGLIAZZO and Moni NAOR.
"Efficient cryptographic schemes provably as secure as subset sum".
In: *Journal of Cryptology* 9.4. Pages 236–241. DOI: 10.1109/SFCS.1989.63484.

1996. David R. KOHEL.
“Endomorphism rings of elliptic curves over finite fields”.
PhD thesis. University of California at Berkeley.
URL: <http://echidna.maths.usyd.edu.au/kohel/pub/thesis.pdf>.
1997. Victor SHOUP.
“Lower bounds for discrete logarithms and related problems”.
In: *Advances in Cryptology — EUROCRYPT ’97*. Edited by Walter FUMY.
Volume 1233. Lecture Notes in Computer Science. Springer. Pages 256–266.
DOI: 10.1007/3-540-69053-0_18.
2001. Leonard M. ADLEMAN and Ming-Deh HUANG.
“Counting points on curves and abelian varieties over finite fields”.
In: *Journal of Symbolic Computation* 23.3. Pages 171–189.
DOI: 10.1006/jscs.2001.0470.
2008. Andrew GRANVILLE.
“Smooth numbers: computational number theory and beyond”.
In: *Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography*.
Edited by Joseph P. BUHLER and Peter STEVENHAGEN. Volume 44.
Mathematical Sciences Research Institute Publications. Cambridge University Press.
Pages 267–323.
2009. Gaetan BISSON and Andrew V. SUTHERLAND.
“Computing the endomorphism ring of an ordinary elliptic curve over a finite field”.
In: *Journal of Number Theory* 131.5. Edited by Neal KOBLITZ and Victor S. MILLER.
Special Issue on Elliptic Curve Cryptography. Pages 815–831.
DOI: 10.1016/j.jnt.2009.11.003.
2009. Kirsten EISENTRÄGER and Kristin E. LAUTER.
“A CRT algorithm for constructing genus 2 curves over finite fields”.
In: *Arithmetic, Geometry and Coding Theory — AGCT ’10*.
Edited by François RODIER and Serge VLADUT. Volume 21. Séminaires et Congrès.
Société Mathématique de France. Pages 161–176.
2009. David JAO, Stephen D. MILLER, and Ramarathnam VENKATESAN.
“Expander graphs based on GRH with an application to elliptic curve cryptography”.
In: *Journal of Number Theory* 129.6. Pages 1491–1504.
DOI: 10.1016/j.jnt.2008.11.006.
2010. Nick HOWGRAVE-GRAHAM and Antoine JOUX.
“New generic algorithms for hard knapsacks”.
In: *Advances in Cryptology — EUROCRYPT ’10*. Edited by Henri GILBERT.

-
- Volume 6110. Lecture Notes in Computer Science. Springer. Pages 235–256.
DOI: 10.1007/978-3-642-13190-5_12.
2010. Sorina IONICA and Antoine JOUX.
“Pairing the volcano”. In: *Algorithmic Number Theory — ANTS-IX*.
Edited by Guillaume HANROT, François MORAIN, and Emmanuel THOMÉ.
Volume 6197. Lecture Notes in Computer Science. Springer. Pages 201–218.
DOI: 10.1007/978-3-642-14518-6_18.
2011. Gaetan BISSON.
Computing endomorphism rings of elliptic curves under the GRH.
arXiv.org: 1101.4323.
2011. Gaetan BISSON and Andrew V. SUTHERLAND.
“A low-memory algorithm for finding short product representations in finite groups”.
In: *Designs, Codes and Cryptography*. To appear.
DOI: 10.1007/s10623-011-9527-8.

EIGHT

Polarized Method

To make practical use of our subexponential method for computing endomorphism rings of ordinary abelian varieties in dimension higher than one, polarizations must be taken into account. This requires certain modifications to be made on our framework, algorithms, and implementation, which we now describe. We also need to rely on more unproven assumptions.

We focus on the case of Jacobian varieties of genus-two hyperelliptic curves, since the availability of certain computational tools (such as the method of MESTRE (1991)) is limited in higher dimensions. Notwithstanding those issues, we believe most of the differences that higher-dimensional varieties have in comparison to elliptic curves are addressed here.

The modified algorithm will be presented before the computation of isogenies; we then give actual computation results and finally discuss vertical isogenies.

VIII.1 Algorithm

COMPLEX MULTIPLICATION FRAMEWORK

We start by recalling some of the theory on which our approach relies.

Let \mathcal{A} be a simple ordinary principally polarized abelian variety of dimension g defined over a finite field. We assume that an embedding of its complex multiplication field $K = \mathbb{Q}(\pi)$ into $\text{End}(\mathcal{A}) \otimes \mathbb{Q}$ has been fixed, which is equivalent to fixing a type Φ on K .

As we saw in Chapter III, ideals of the reflex field K^r act on isomorphism classes of principally polarized abelian varieties \mathcal{A} via the reflex type norm (see Figure 7):

$$\tau \in \mathfrak{I}(K^r) : \mathbb{C}^g / \Phi(\mathfrak{a}), E_{\Phi}^{\xi} \mapsto \mathbb{C}^g / \Phi(N_{\Phi^r}(\tau)^{-1} \mathfrak{a}), E_{\Phi}^{N_{K^r/\mathbb{Q}}(\tau)\xi}$$




FIGURE 7. Complex multiplication field extensions and their reflex counterparts.

The main difference to the preceding chapter is that, when the dimension g is two or more, this action only gives certain elements of the polarized class group $\mathcal{C}(\mathcal{O}_K)$; in other words, it describes certain, but not all, isogenies. Therefore, a rigorous analysis of our algorithm in this setting would be much more involved than in the utopian case where polarizations were disregarded: one would need to assert the existence of short relations arising via the reflex type norm, which we see no simple way of doing. Therefore we assume:

Assumption VIII.1.1. *Under the map $(\mathfrak{a}, \ell) \mapsto (\mathfrak{a}\mathcal{O}, \ell)$, composed to the right with the reflex type norm, ideals of the ring of integers of the reflex field act faithfully on the set $\text{AV}_{\mathcal{O}}(k)$ of principally polarized abelian varieties with endomorphism \mathcal{O} over the base field k .*

This comes on top of the generalized Riemann hypothesis, and Assumptions VII.1.2, VII.1.4, and VII.3.6, which state respectively:

- Orders $\mathcal{O} \subset \mathcal{O}'$ for which the above action is identical have bounded index $[\mathcal{O}' : \mathcal{O}]$.
- The method of EISENTRÄGER and LAUTER (2009) computes $\text{End}(\mathcal{A})_{\ell}$ in $\ell^{O(1)}$ time.
- The norms of reduced ideals are as smooth as random integers.

The first assumption is a helpful heuristic, the third comes from BUCHMANN (1989), and the second deliberately rules out cases where the local lattice of orders is deep. They were all largely verified in the range of practical problems that we considered, except in certain rare cases.

We also require the ability to draw points at random from \mathcal{A} and other varieties of its isogeny class; for $g = 2$, this is always the case using Weierstrass forms, to which any variety can be put using the method of MESTRE (1991). Therefore we additionally impose $g = 2$.

Under all these assumptions, the expected runtime is, as we mentioned before:

$$L(q)^g \sqrt{3g/2+o(1)}$$

OVERVIEW

Let \mathcal{A} be the input polarized abelian variety, given as the Jacobian variety of a hyperelliptic curve \mathcal{C} defined over the finite field with q elements. First, we compute the characteristic polynomial χ_π of its Frobenius endomorphism π , which the algorithm of PILA (1990) does in polynomial time. In practice, we relied on the point-counting routines of the MAGMA (1997) computational algebra system, which use the techniques of GAUDRY and HARLEY (2000); larger base fields could be reached using the state-of-the-art implementation and optimizations of GAUDRY and SCHOST (2010).

In the lattice of orders, we find $\text{End}(\mathcal{A})$ from below using the following algorithm from Chapter VI — we also proposed a way of finding $\text{End}(\mathcal{A})$ from above which is suited to varieties constructed via the complex multiplication method (rather than at random, as below); however, at the time of this writing, only abelian varieties with maximal endomorphism rings can be generated in this way, except in the one-dimensional case.

Algorithm VIII.1.2.

INPUT: A simple ordinary principally polarized abelian variety \mathcal{A} over a finite field \mathbb{F}_q .

OUTPUT: An order isomorphic to its endomorphism ring.

1. Compute the Frobenius polynomial $\chi_\pi(x)$ of \mathcal{A} .
2. Factor the discriminant Δ and construct the order $\mathcal{O}' = \mathbb{Z}[\pi, \bar{\pi}]$.
3. For orders \mathcal{O} directly above \mathcal{O}' :
4. If $\mathcal{O} \subset \text{End}(\mathcal{A})$ set $\mathcal{O}' \leftarrow \mathcal{O}$ and go to Step 3.
5. Return \mathcal{O}' .

To determine whether a specific order \mathcal{O} is contained in the endomorphism ring of \mathcal{A} , we selected several relations of it (typically logarithmically many in the number of orders of containing $\mathbb{Z}[\pi, \bar{\pi}]$, although doubly logarithmically many should theoretically be enough), and checked whether these relations hold in the isogeny graph. The latter step requires us to evaluate isogenies and is the bottleneck of the whole algorithm.




FIGURE 8. Galois groups of the complex multiplication fields tower.

DENSITY OF SPLITTING PATTERNS

To study the splitting pattern of rational primes ℓ in complex multiplication field K , let us first present the setting to which Theorem VII.2.1 can be applied. We are mostly interested in the splitting of primes in the reflex field K' of the field by which our variety has complex multiplication, but it makes no difference for this analysis.

Denote by K any complex multiplication field of degree $2g$, and write K^c for its normal closure. Similarly, denote by K_+ the normal closure of its totally real subfield K_+ . This gives a tower of fields as displayed on Figure 8.

In the typical case of non-Galois number fields, DODSON (1984) established the isomorphisms $\text{Gal}(K_+/\mathbb{Q}) \simeq \mathfrak{S}_g$ and $\text{Gal}(K^c/K_+) \simeq (\mathbb{Z}/2)^\nu$ for some integer ν in $\{1, \dots, g\}$, and described the action of the former on the latter so that we have an explicit description of the Galois structure of K^c/\mathbb{Q} as

$$\text{Gal}(K^c/\mathbb{Q}) \simeq (\mathbb{Z}/2)^\nu \rtimes \mathfrak{S}_g.$$

Note that, when a principally polarized abelian variety \mathcal{A} is absolutely simple (as we assume here), its complex multiplication field K is primitive and we have $\nu = g$. In dimension $g = 2$, the Galois group of K^c/\mathbb{Q} is then isomorphic to the dihedral group $D_4 = \mathbb{Z}/4 \rtimes \mathbb{Z}/2$, and we obtain the densities of Figure 9 as a consequence.

SPLITTING PATTERN	(1, 1, 1, 1)	(1, 1, 2)	(1, 3)	(2, 2)	(4)
DENSITY OF PRIMES	1/8	1/4	0	3/8	1/4

FIGURE 9. Density of rational primes p splitting in a fixed non-normal quartic complex multiplication field as $\prod p_i$ with pattern $(N(p_i))$.

FINDING RELATIONS

Finding relations is a quite standard step. We have already mentioned that the computability of the algebraic structures we deal with has been well studied. Here, in fact, we do not even need to compute the polarized class group of Shimura: since we are restricted to using isogenies which arise under the reflex type norm, we are in fact seeking for relations of the class group of \mathcal{O}^v . To obtain a subexponential asymptotic runtime, we use the generalization of the algorithm of HAFNER and MCCURLEY (1989) by BUCHMANN (1989).

Remark. As a practical optimization, since evaluating isogenies is so costly, more time may be dedicated to finding a shorter relation. For the range of input sizes we considered, it was well worth using the exponential algorithm below which is essentially a baby-step giant-step approach borrowing ideas of COHEN, DIAZ Y DIAZ, and OLIVIER (1997) for the effective ideal arithmetic; it finds the shortest possible relation, therefore improving greatly the speed of the isogeny step, and reducing the overall runtime.

Notation. Recall that $b_x(f(x))$ may denote any function satisfying the inequalities $f(x) < b_x(f(x)) < f(x)^{1+o(1)}$ and computable in essentially linear time in $f(x)$.

Algorithm VIII.1.3.

INPUT: An order \mathcal{O} of discriminant Δ in a number field K .

OUTPUT: Relations of \mathcal{O} .

1. Let \mathfrak{B} consist of all prime ideals with norm up to $b_\Delta(12 \log^2 |\Delta|)$.
2. Create a hash table H .
3. Compute the product \mathfrak{a} of a random subset of \mathfrak{B} .
4. Let \mathfrak{b} be an LLL reduction of \mathfrak{a} .
5. If H has an entry for \mathfrak{b} , output $H(\mathfrak{b}) - \mathfrak{a}$.
6. Otherwise, set $H(\mathfrak{b}) \leftarrow \mathfrak{a}$ and go back to Step 3.

Step 4 means that \mathfrak{b} is the ideal generated over \mathcal{O} by an LLL basis of the ideal \mathfrak{a} , where the LLL reduction can be computed along any direction as described by COHEN, DIAZ Y DIAZ, and OLIVIER (1997). The ideals \mathfrak{b} act as class representatives and we do not require

that they are unique: it is enough that they are small so that, by the pigeonhole principle, classes are identified after a few more trials than what would be required otherwise.

The use of such an exponential algorithm also has an additional benefit: it allows us to choose which primes we want to include in our relations, which subexponential smoothness-based methods do not permit.

For instance, we can choose to only use primes which split as $p\bar{p}$, hence allowing for a cardinality-based approach and sparing the need to compute the characteristic polynomial of the Frobenius polynomial acting on the kernel of the isogeny. This is not a tremendous improvement since our current isogeny-evaluating technique actually requires computing the kernel, but it would be helpful if a modular-polynomial-based method was used.

More importantly, we can restrict to primes which are congruent to one modulo four; this avoids the need for an additional quadratic extension to compute torsion points, and lowers the complexity of level-change formulas from $\ell^{2g+o(1)}$ to $\ell^{g+o(1)}$.

VIII.2 Computing Isogenies

To determine the endomorphism ring of a principally polarized abelian variety by exploiting the complex multiplication action, we need to evaluate the isogeny $\phi_{\mathfrak{a}}$ corresponding to a prescribed ideal \mathfrak{a} . In dimension one, this uses the formulas of VÉLU (1971) and Stage 3 of the algorithm by GALBRAITH, HESS, and SMART (2002).

The work of RICHELLOT (1836) was interpreted by BOST and MESTRE (1988) to compute isogenies of type $(\mathbb{Z}/2)^2$ between Jacobian varieties of genus-two hyperelliptic curves. Later, CARLS, KOHEL, and LUBICZ (2008) obtained relations describing pairs of abelian surfaces related by an isogeny of type $(\mathbb{Z}/3)^2$; this was implemented and publicly released in the ECHIDNA (2007) package.

This section gives a brief overview of the evaluation of general isogenies between abelian varieties as implemented in the library of B., COSSET, and ROBERT (2010); for most of the mathematical aspects, we refer to LUBICZ and ROBERT (2009) and COSSET and ROBERT (2011). We evaluate $\phi_{\mathfrak{a}}$ in four steps: we first find its kernel, convert it into theta coordinates, then perform the actual isogeny computation, and finally express the result as absolute invariants.

FINDING KERNELS

Kernels of isogenies of type $(\mathbb{Z}/\ell)^g$ that respect the polarization are maximal isotropic rational subgroups of \mathcal{A} isomorphic to $(\mathbb{Z}/\ell)^g$ and defined over an algebraic closure of the base field \mathbb{F}_q .

Since each has order ℓ^g , is rational over the base field, and contains the neutral element, they are all defined over an extension of degree $\ell'(\ell)$ at most $\ell^g - 1$. We will thus simply enumerate all such subgroups of the ℓ -torsion group of $\mathcal{A}(\mathbb{F}_{q^{\ell'(\ell)}})$ and then find which one corresponds to the ideal \mathfrak{a} as mentioned above.

To find these, first compute a basis of the ℓ -Sylow subgroup of \mathcal{A} over the extension field, which we denote by

$$\mathcal{A}(\mathbb{F}_{q^{\ell'(\ell)}})[\ell^\infty];$$

for this, we use the method of COUVEIGNES (2009) which we have discussed before: it amounts to taking random points of \mathcal{A} (this is easy, for instance, when it has a Weierstrass form), multiplying them by the cofactor of ℓ^∞ in $\#\mathcal{A}(\mathbb{F}_{q^{\ell'(\ell)}})$, and “lifting” these points along each other until a basis of the ℓ -torsion group is obtained.

We then derive a symplectic basis of $\mathcal{A}(\mathbb{F}_{q^{\ell'(\ell)}})[\ell]$ for the Weil pairing. For simplicity, fix an ℓ^{th} root of unity and consider the problem additively under the corresponding logarithm $\log : \mu_\ell(\mathbb{C}) \rightarrow \mathbb{Z}/\ell$. On the basis we are looking for, (the logarithm of) the Weil pairing is given by the matrix

$$\begin{pmatrix} 0 & \mathbf{I}_g \\ -\mathbf{I}_g & 0 \end{pmatrix}.$$

To obtain such a basis $(e_1, \dots, e_g, f_1, \dots, f_g)$ satisfying

$$\begin{cases} \log \Psi_{\text{Weil}}(e_i, f_j) = \delta_{ij} \\ \log \Psi_{\text{Weil}}(e_i, e_j) = 0 \\ \log \Psi_{\text{Weil}}(f_i, f_j) = 0 \end{cases}$$

we use an elementary, orthogonalization-like algorithm, similar to the classical algorithm for computing Smith normal forms.

This basis allows us to enumerate all symplectic subgroups easily and, amongst these, we select those that are rational, that is, stable under the Frobenius endomorphism, and find which is acted upon with characteristic polynomial u (given by the ideal \mathfrak{a}).

Note that when ℓ is congruent to one modulo four, finding random points of \mathcal{A} is faster by a factor of two since computing the square root of the Weierstrass polynomial evaluated at x in order to get the y -coordinate simply amounts to a modular exponentiation.

MAPPING TO THETA COORDINATES

Recall that if $\mathcal{A} \simeq \mathbb{C}^g / (\mathbb{Z}^g + \Omega\mathbb{Z}^g)$ is a complex torus with period matrix Ω in \mathbb{H}^g , then the set of theta functions

$$\Theta_{a,b}^{\mathcal{A}} : z \in \mathbb{C}^g \mapsto \sum_{(u+a) \in \mathbb{Z}^g} \exp\left(i\pi \left(\frac{1}{n} \widehat{u} \Omega u + 2\widehat{u}(z+b)\right)\right),$$

where $a = 0$ and b is a vector of $\frac{1}{n}(\mathbb{Z}/n)^g$, forms the *theta coordinate system of level n* . It is a coordinate system for abelian varieties (and also incorporates information about the n -torsion), but can represent points of such varieties too. It has an algebraic counterpart which is applicable to varieties defined over finite fields.

The points P of the kernel of the isogeny we wish to evaluate, as output by the method of COUVEIGNES (2009), are expressed in Mumford coordinate on a Weierstrass model for the hyperelliptic curve $\mathcal{C} : y^2 = f(x)$ of which \mathcal{A} is the Jacobian variety. As a first step towards mapping these points to theta coordinates, we extend the base field so as to make f split completely; then, by a homographic transformation (also known as Möbius transformation) of the x coordinate, we derive its Rosenhain normal form

$$y^2 = x(x-1) \prod_{i=1}^{2g-1} (x-a_i)$$

which might require working in an extension of the base field.

The formulas of THOMAE (1870), then give theta coordinates of level two or four corresponding to the variety $\mathcal{A} = \text{Jac}(\mathcal{C})$. In order to map points from Mumford representation to theta coordinates, we need equations derived by VAN WAMELEN (1998).

Note that theta coordinates of level two actually represent the Kummer surface of an abelian variety, that is, identify a variety $\mathcal{A} = \text{Jac}(\mathcal{C} : y^2 = f(x))$ with its twist $\text{Jac}(\widetilde{\mathcal{C}} : \omega y^2 = f(x))$ where ω is a non-quadratic residue in the base field. This is not too much of an issue for us since the isogeny class of \mathcal{A} is identified by the characteristic polynomial of its Frobenius endomorphism, so there is no ambiguity on which of an abelian variety \mathcal{B} or its twist an isogeny $\phi_{\mathfrak{a}}$ with domain \mathcal{A} maps to.

However, for a cleaner approach, we prefer to use level four theta coordinates which identify the variety \mathcal{A} uniquely; this comes at the expense of speed, but the slow down is minor, especially as finding the ℓ -torsion remains the overall bottleneck.

ISOGENIES VIA LEVEL CHANGING

LUBICZ and ROBERT (2009) described isogenies as projections from higher-level theta coordinate systems to lower-level ones; they also described the associated machinery (addi-




FIGURE 10. Evaluating isogenies of type $(\mathbb{Z}/\ell^2)^g$ via two theta level changes.

tion laws, etc.) required to make effective use of this result. Before discussing how it applies to our setting, let us briefly recall their result.

Theorem VIII.2.1. *Let \mathcal{H} be a subgroup isomorphic to $(\mathbb{Z}/\ell)^g$ of an abelian variety \mathcal{A} of dimension g , and let n be any integer coprime to ℓ . The theta functions of level n on \mathcal{A}/\mathcal{H} are a subset of the theta functions of level ℓn on \mathcal{A} .*

This introduces a change of level; to address this, LUBICZ and ROBERT (2009) noted that subsets of the Fourier transform of theta functions of level ℓn on \mathcal{A} correspond to theta functions of level n for abelian varieties obtained by dual isogenies of degree ℓ ; this allows them to compute isogenies of type $(\mathbb{Z}/\ell^2)^g$ between abelian varieties expressed by level- n theta functions; see Figure 10.

Our framework for computing endomorphism rings can be adapted to this setting: relations can be constrained to only involve squares of ideals, so that the associated isogenies are all of type $(\mathbb{Z}/\ell^2)^g$. However, this implies losing all the information regarding the 2-torsion of the reflex class group $\mathcal{C}(\mathcal{O}^n)$. COHEN and LENSTRA (1984) showed that class groups typically have a large 2-torsion subgroup, so it is not likely that all pairs of class groups that are identical up to 2-torsion can be distinguished efficiently using the local method of EISENTRÄGER and LAUTER (2009).

COSSET and ROBERT (2011) then derived from earlier work of KOIZUMI (1976) and KEMPF (1989) formulas which allow to map points from level- ℓn theta coordinates to level- n theta coordinates, avoiding the need to evaluate an additional isogeny. They apply these formulas to evaluating isogenies of type $(\mathbb{Z}/\ell)^g$ between abelian varieties expressed in theta coordinates of level n .

MAPPING BACK TO INVARIANTS

In order to determine whether a relation holds in the isogeny graph of an abelian variety (to eventually determine its endomorphism ring), we need to compose many isogenies of type $(\mathbb{Z}/\ell)^g$ for various primes ℓ . We have explained how to compute an isogeny $\mathcal{A} \rightarrow \mathcal{A}'$ of prescribed kernel where \mathcal{A} is given in Weierstrass form and \mathcal{A}' is given as theta coordinates of level n . To iterate this construction, it remains to explain how we can obtain a Weierstrass equation for \mathcal{A}' .

In fact, this can be done elementarily by inverting the formulas of THOMAE (1870). However, the theta coordinates of \mathcal{A} that we used in the isogeny computation are defined over a large extension of the base field which contains the roots of the Weierstrass polynomial of the curve, certain n -torsion points (recall that $n = 2$ or 4) and certain ℓ -torsion points; the theta coordinates of \mathcal{A}' , and therefore also its Weierstrass equation that we derive, are consequently defined over that large extension.

When we know that \mathcal{A}' is actually defined over the base field (for instance, because the chosen isogeny is rational), we recover a rational Weierstrass equation by first computing the absolute invariants of \mathcal{A}' and then using the algorithm of MESTRE (1991) to reconstruct a curve \mathcal{C}' whose Jacobian variety $\text{Jac}(\mathcal{C}')$ is \mathcal{A}' .

As an optimization to the algorithm for finding the ℓ -torsion of the new curve \mathcal{A}' and then compute the next isogeny step, ROBERT noticed that part of the ℓ -torsion of \mathcal{A} can be reused: indeed, we have $\mathcal{A}[\ell] \simeq (\mathbb{Z}/\ell)^{2g}$ and the isogeny $\mathcal{A} \rightarrow \mathcal{A}'$ only kills half of it; therefore, we can map the remaining points all the way from \mathcal{A} to \mathcal{A}' and start the search for a basis of $\mathcal{A}'(\mathbb{F}_{q^{1/\ell}})[\ell]$ knowing already half the solution. This can speed up the search for rational torsion subgroups of type $(\mathbb{Z}/\ell)^g$ by a factor of two.

Abelian varieties of dimension strictly greater than two are not necessarily Jacobian varieties of hyperelliptic curves, and from dimension four on they might not even be Jacobian varieties at all. Therefore, two of our building blocks fail:

- the selection of random points (to find a basis for $\mathcal{A}[\ell]$);
- the method of MESTRE (1991) (to reduce the field of definition of \mathcal{A}').

The former can easily be addressed by assuming that our abelian varieties come equipped with an efficient algorithm for obtaining random points. The latter is a more delicate issue: the isogeny computation requires working in an extension field, and for $g > 2$ we do not know how to go back to the base field afterwards.




FIGURE 11. Average time for finding the ℓ -torsion of an abelian variety of dimension two over the field with 251 elements, for $\ell \in \{2, 3, 5, 7, 11, 13, 17, 19\}$ and all possible $e'(\ell)$.

VIII.3 Practical Computations

All computations were realized using the library of B., COSSET, and ROBERT (2010).

FINDING TORSION

The bottleneck of our algorithm is typically to find a basis for the ℓ -torsion subgroup of \mathcal{A} over an extension where all points of rational subgroups of type $(\mathbb{Z}/\ell)^\mathcal{G}$ are defined. The cost is twofold:

- computing over an extension of degree $e'(\ell)$ of the base field;
- multiplying points by the cofactor of ℓ^∞ in $\#\mathcal{A}(\mathbb{F}_{q^{e'(\ell)}}) \sim q^{g e'(\ell)}$.

In the worst case, $e'(\ell)$ can be as large as $\ell^g - 1$, so that the overall complexity is $\ell^{2g+o(1)}$ disregarding logarithmic factors in q , which quickly becomes prohibitive. As argued before, exponential methods for finding relations offer the advantage that specific primes ℓ can be chosen for which $e'(\ell)$ is small.

Figure 11 shows the time it takes, on average for 10 randomly chosen abelian surfaces defined over the field \mathbb{F}_{251} , to compute the ℓ -torsion over an extension of degree e' .




FIGURE 12. Number of iterations the latter algorithm requires before finding a relation in a quartic complex multiplication field with certain class number (also known as Picard number). The lines plot $y = x$ and $y = \sqrt{x}$.

As can be expected, this runtime is slightly more than linear in the extension degree, and does not highly depend on ℓ . However, we observe that for a prescribed ℓ the torsion of varieties with a certain $e'(\ell)$ is sometimes faster than those of varieties with a smaller $e'(\ell)$; this is likely due to the internal representation of the extensions as tower fields in MAGMA (1997), and also possibly to special features of the varieties.

FINDING RELATIONS

We implemented in MAGMA (1997) the simple baby-step giant-step method that we described above and found that it behaves well: in most cases, the number of iterations required to find a collision is not so far from the \sqrt{b} (where b denotes the class number) that would be expected if each ideal class contained a unique reduced ideal.

Figure 12 shows the number of iterations our algorithm goes through before the first relation is found; we use the order $\mathcal{O} = \mathbb{Z}[\pi, \bar{\pi}]$ for a thousand Jacobian varieties of random hyperelliptic curves of genus two. The class number displayed is actually the approximation $\sqrt{|\Delta|/R}$ given by the Brauer–Siegel theorem.

We observe that the iteration count lies somewhere in between \sqrt{b} and b . Although in

some cases this number went slightly above the class number, the runtime was always acceptable: it was never more than two seconds when the class number was less than a thousand, and always less than a hundred seconds in our range of parameters.

BEST-CASE SCENARIO

Let us first present an example where our algorithm performs much better than all other alternatives. The *conductor gap* is the largest prime factor of the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$; here, we consider a case of large conductor gap, which makes the method of EISENTRÄGER and LAUTER (2009) impractical. Unfortunately, we were unable to compare our method with that of WAGNER (2009), as we did not have an implementation of the latter at our disposal.

To find an abelian variety with a large conductor gap, we generated genus-two hyperelliptic curves at random until one whose Jacobian variety has the desired property was found; we obtained the hyperelliptic curve with equation

$$y^2 = 80742x^5 + 56078x^4 + 76952x^3 + 134685x^2 + 60828x + 119537$$

defined over the field with 161983 elements; let \mathcal{A} denote its Jacobian variety. The characteristic polynomial of its Frobenius endomorphism is

$$z^4 - 144z^3 + 10368z^2 - 144 \cdot 161983z + 161983^2$$

and it defines a quartic complex multiplication field $K = \mathbb{Q}(\pi)$ in which the ring of integers contains the minimal order $\mathbb{Z}[\pi, \bar{\pi}]$ with prime index $\ell = 156799$.

Since the full ℓ -torsion of \mathcal{A} lies in an extension of degree $e(\ell) = 78399$, it is challenging to try to compute $\text{End}(\mathcal{A})$ using the method of EISENTRÄGER and LAUTER (2009).

However, the Picard group of $\mathfrak{M} = \mathcal{O}_K$ has order 460; this is not surprising as a large part of $\Delta = \text{disc}(\pi)$ contributes to the conductor gap so little is left to build up $\text{disc}(K)$. It is thus easy to find relations in the associated polarized class group $\mathfrak{C}(\mathcal{O}_K)$. For instance, one easily verifies that the element $(\mathfrak{a}, 3)$ has order 115, where \mathfrak{a} can be any ideal of norm 9 (there are just two such elements, inverses of each others).

The action of $(\mathfrak{a}, 3)^{115}$ on \mathcal{A} is computed easily, as the 3-torsion of \mathcal{A} lives over an extension of degree 8. Using just one core of an Intel Xeon E5440 processor clocked at 2.83 GHz, our humble Magma implementation computes it in just over four minutes. Since it finds that $\phi_{(\mathfrak{a}, 3)^{115}} \mathcal{A} \neq \mathcal{A}$, we deduce that $\text{End}(\mathcal{A}) = \mathbb{Z}[\pi, \bar{\pi}]$. Note that, since the action of \mathcal{O}_K on $\text{AV}_{\mathcal{O}_K}(k)$ is always faithful, and there are only two orders in the lattice, this result holds unconditionally regardless of the assumptions.




FIGURE 13. Lattice of orders with \mathcal{O}_K on top and $\mathbb{Z}[\pi, \bar{\pi}]$ at the bottom; lines indicate that the order below is contained in the order above with index the label right off the line.

WORST-CASE SCENARIO

Now let us consider an abelian variety that is expectedly suited to the method of EISENTRÄGER and LAUTER (2009), namely, one for which the conductor gap $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ is short. We take the Jacobian variety \mathcal{A} of the hyperelliptic curve

$$y^2 = 2987x^5 + 1680x^4 + 3443x^3 + 1918x^2 + 2983x + 489$$

defined over the field with 3499 elements. The characteristic polynomial of π is

$$z^4 + 48z^3 + 1152z^2 + 48 \cdot 3499z + 3499^2$$

and we find that there are 2^4 orders containing (or equal to) $\mathbb{Z}[\pi, \bar{\pi}]$; their indices in the maximal order divide $13^2 \cdot 37 \cdot 79$ as displayed on Figure 13.

We use $\alpha_\ell = (\mathfrak{a}_\ell, \ell) \in \mathfrak{C}$ for $\ell \in \{3, 5, 7\}$ where \mathfrak{a}_ℓ is an arbitrary ideal of norm ℓ^2 ; the full ℓ -torsion is defined over an extension of degree 8, 24, and 24, respectively, so it takes on average 1, 3.5, and 5.5 seconds to evaluate one ℓ -isogeny.

We used the relation $\alpha_3^5 \alpha_7^7 = 1$ for the yellow square order, $\alpha_5^{10} = 1$ for the blue triangle order, and $\alpha_3^2 \alpha_5^{16} \alpha_7^{-2} = 1$ for both the red circle and green diamond order. Checking these relations in the isogeny graph took only slightly more than two minutes, and since none was found to hold, our algorithm returned that $\text{End}(\mathcal{A}) = \mathbb{Z}[\pi, \bar{\pi}]$.

Even in this case, which would *a priori* favor the method of EISENTRÄGER and LAUTER (2009) (the full 37 and 79-torsion are defined over extensions of degree 1332 and 948, respectively), our algorithm performs well while still leaving some room for improvement.

VIII.4 Isogeny Volcanoes

Let us now fix a prime ℓ and study the structure of the connected component of the graph of isogenies of type $(\mathbb{Z}/\ell)^\xi$ containing a prescribed principally polarized simple ordinary abelian variety \mathcal{A} defined over some finite field.

GENERAL STRUCTURE

KOHEL (1996) and later FOUQUET and MORAIN (2002) depicted the structure of such graphs in dimension one as volcanoes, containing a *crater* formed by varieties whose endomorphism ring is locally maximal. Horizontal isogenies arrange these varieties in a (possibly degenerated) circle, and from each vertex on it hang complete ℓ -ary trees; their number and depth are entirely determined by ℓ and the isogeny class.

In dimension two or more, most specific details are lost, but the general structure remains the same; most important for our algorithms is that craters are still Cayley graphs.

Let $G = (V, E)$ be such an isogeny graph: vertices V correspond to abelian varieties and edges E (a symmetric subset of V^2) to isogenies of type $(\mathbb{Z}/\ell)^\xi$ between them. We start by partitioning G into *layers* $G_\mathcal{O}$ for each order \mathcal{O} above $\mathbb{Z}[\pi, \bar{\pi}]$: each layer contains the vertices whose associated varieties have an endomorphism ring isomorphic to \mathcal{O} .

Note that, in a connected component, certain layers can be empty as not all isogenous varieties might be reachable by sequences of isogenies of type $(\mathbb{Z}/\ell)^\xi$. We say that a layer $G_\mathcal{O}$ is maximal when there is no non-empty $G_{\mathcal{O}'}$ with $\mathcal{O} \subsetneq \mathcal{O}'$; typically, this means that when $G_{\mathcal{O}_k}$ is non-empty, it is the unique maximal layer.

Our observations of isogeny volcanoes will be split in three parts:

- the *core*: the union of maximal layers and their horizontal isogenies;
- the *branches*: the vertical isogenies;
- the *covering*: the horizontal isogenies in non-maximal layers.

Often, the graph has the familiar picture of a core, out of which branches hang, and there is no covering. However, we will see that unusual phenomena can occur, such as part of the branches substituting to the core structure.

At any rate, we must warn the reader that our description of branches (which are the key to understanding the relationship of ℓ -isogeny volcanoes and the structure of endomorphism rings locally at ℓ) will be short and qualitative, as this thesis focuses on using horizontal isogenies and does not pretend to add any insight on the structure of vertical isogenies.

CORES

Assume the core consists of a single layer $G_{\mathcal{O}}$ (we will consider the case where there are two or more below).

At least in the case that \mathcal{O} is a maximal order, the theory of complex multiplication proves that the set of horizontal isogenies of type $(\mathbb{Z}/\ell)^g$ in $G_{\mathcal{O}}$ corresponds to a certain subgroup of $\mathcal{C}(\mathcal{O})$ formed of ideals of norm ℓ^g . Therefore, the core is a Cayley graph. We shall denote by $C(X|Y)$ the Cayley graph of X in the free abelian group generated by X with relations Y .

When $g = 2$, the order \mathcal{O} is quartic, and the possible unramified splitting patterns of a prime ℓ in \mathcal{O} are $(1, 1, 1, 1)$, $(1, 1, 2)$, $(1, 3)$, $(2, 2)$, and (4) . The third case never happens in complex multiplication fields (it is incompatible with complex conjugation) and the latter is that of inert primes which act trivially on the isogeny graph, so we disregard both.

In the second case where ℓ splits as $\mathfrak{p}\bar{\mathfrak{p}}\mathfrak{q}$ with $N(\mathfrak{q}) = \ell^2$ there are, in general, no ideals \mathfrak{a} of norm ℓ^2 such that $\mathfrak{a}\bar{\mathfrak{a}}$ is principal, which means there are no corresponding elements in the polarized class group $\mathcal{C}(\mathcal{O})$ and no isogenies of type $(\mathbb{Z}/\ell)^g$.

In the fourth case where ℓ splits as $\mathfrak{p}\bar{\mathfrak{p}}$, both \mathfrak{p} and $\bar{\mathfrak{p}}$ lift to $\mathcal{C}(\mathcal{O})$ as $\alpha = (\mathfrak{p}, \ell)$ and $\beta = (\bar{\mathfrak{p}}, \ell)$. The core of the isogeny graph $G_{\mathcal{O}}$ is then the Cayley graph $C(\alpha, \beta | \alpha\beta, \alpha^{\text{ord}\alpha})$, where the orders implied are those of the corresponding ideals as elements of the Picard group. This gives a cycle structure as Figure 14 displays.

In the first case where ℓ splits as $\mathfrak{p}\bar{\mathfrak{p}}\mathfrak{q}\bar{\mathfrak{q}}$, there are four ideals of norm ℓ^2 whose product with their conjugate is principal, namely $\mathfrak{p}\bar{\mathfrak{p}}$, $\bar{\mathfrak{p}}\mathfrak{q}$, $\mathfrak{p}\mathfrak{q}$, and $\bar{\mathfrak{p}}\bar{\mathfrak{q}}$; if we denote the corresponding elements of $\mathcal{C}(\mathcal{O})$ by α, β, γ , and δ , we obtain that the core $G_{\mathcal{O}}$ is the Cayley graph $C(\alpha, \beta, \gamma, \delta | \alpha\beta, \gamma\delta, \alpha^{\text{ord}\alpha}, \gamma^{\text{ord}\gamma})$; this is a quadrangulation of a torus, as can be seen on Figure 15.

Although we were unable to compute actual isogeny graphs for $g > 2$, primes ℓ which completely split as $\prod_{\mathfrak{p} \in \mathfrak{P}} \mathfrak{p}\bar{\mathfrak{p}}$ (with $\#\mathfrak{P} = g$) would then yield the 2^g elements of $\mathcal{C}(\mathcal{O})$

$$\alpha_{\mathfrak{F}} = \left(\prod_{\mathfrak{p} \in \mathfrak{F}} \mathfrak{p} \prod_{\mathfrak{p} \notin \mathfrak{F}} \bar{\mathfrak{p}}, \ell^g \right)$$

for each subset \mathfrak{F} of \mathfrak{P} ; the core would then be the Cayley graph

$$C \left((\alpha_{\mathfrak{F}})_{\mathfrak{F} \subset \mathfrak{P}} \left| \left(\prod_{\mathfrak{F} \in \mathfrak{G}} \alpha_{\mathfrak{F}} \right), (\alpha_{\mathfrak{F}}^{\text{ord}\alpha_{\mathfrak{F}}})_{\mathfrak{F} \subset \mathfrak{P}} \right. \right)$$

where the middle sequence ranges over all sets \mathfrak{G} of subsets of \mathfrak{P} which satisfy $\#\{\mathfrak{F} \in \mathfrak{G} : \mathfrak{p} \in \mathfrak{F}\} = \#\{\mathfrak{F} \in \mathfrak{G} : \mathfrak{p} \notin \mathfrak{F}\}$ for all $\mathfrak{p} \in \mathfrak{P}$. Topologically, this is the 1-skeleton of a simplicial complex homeomorphic to the g -torus (the product of g copies of the 1-sphere).




FIGURE 14. Graph of isogenies of type $(\mathbb{Z}/3)^2$ containing the Jacobian variety of the curve $y^2 = 3x^5 + 15x^4 + 11x^3 + 3x^2 + 11x + 12$ over the field with 19 elements.




FIGURE 15. Graph of isogenies of type $(\mathbb{Z}/7)^2$ containing the Jacobian variety of the curve $y^2 = 106x^6 + 83x^5 + 18x^4 + 52x^3 + 49x^2 + 11x + 41$ over the field with 109 elements.

Note that all the above holds over an algebraic closure, as not all isogenies corresponding to ideals of norm ℓ^s of $\mathcal{C}(\mathcal{O})$ need to be rational.

BRANCHES

Let us now consider two-dimensional ℓ -isogeny graphs in the case that $\ell\mathcal{O}_K$ is not coprime with the conductor of $\mathbb{Z}[\pi, \bar{\pi}]$. Although our algorithms for computing endomorphism rings prefer to avoid such situations, they are an interesting application of our isogeny-computing library.

Each figure contains two parts: the isogeny graph to the left, and the lattice of orders to the right. Vertices of the isogeny graph are colored the same way as the endomorphism rings of the corresponding abelian varieties are in the lattice of orders.

Recall that in dimension one, a certain number of complete n -ary trees of uniform depth hang from each vertex of the core. This might also happen in higher dimension, but other scenarios are possible. For instance, BRÖKER, GRUENEWALD, and LAUTER (2009) observed in their Example 8.3 that trees hanging from the core might have different depths. Figure 16 shows the same phenomenon in a more generic-looking graph. This unbalance shows that not all isogenies of type $(\mathbb{Z}/\ell)^s$ need be uniformly rational.

Figure 16 also features isogeny of type $(\mathbb{Z}/\ell)^2$ between abelian varieties whose endomorphism rings have index ℓ^2 in each other, more specifically between the green diamond and cyan octagon dots. This can lead to disturbing graphs such as that of Figure 17 where the endomorphism rings of varieties $(\mathbb{Z}/\ell)^2$ -isogenous to varieties with maximal ones are the order of index 3^2 , some order of index 3, but not the maximal order itself. Going from one variety with maximal endomorphism ring to another is however possible by first going through a non-maximal one and then going up again.

In such cases, the partitioning of the features of isogeny graphs into a core, branches, and coverings is somewhat flawed. Although with our definition, the core of Figure 17 consists of both curves with red circle (maximal) and yellow square (index 3) endomorphism rings.

This illustrates another obstruction to climbing higher-dimensional volcanoes: sometimes, steps can only be climbed in pairs, which prevents one to fully enumerate an isogeny class just by following isogenies of type $(\mathbb{Z}/\ell)^s$. Naturally, we see (hypothetical) isogenies of type (\mathbb{Z}/ℓ) as the answer to this problem.

COVERINGS

We call covering the outer layers of the isogeny graph; those are horizontal isogenies arising as complex multiplication “residues.” Although there are no ideals of norm ℓ^s in imaginary quadratic orders whose conductors are divisible by ℓ , this sometimes happen in higher




FIGURE 16. Graph of isogenies of type $(\mathbb{Z}/3)^2$ containing the Jacobian variety of the curve $y^2 = 44x^6 + 36x^5 + 48x^4 + 29x^3 + 3x^2 + 44x + 34$ over the field with 61 elements.




FIGURE 17. Graph of isogenies of type $(\mathbb{Z}/3)^2$ containing the Jacobian variety of the curve $y^2 = 13x^6 + 5x^5 + 37x^4 + 31x^3 + x^2 + 5x + 3$ over the field with 43 elements.

dimension; by complex multiplication, such ideals give rise to horizontal isogenies amongst varieties with non-maximal endomorphism rings.

This was first noted by BRÖKER, GRUENEWALD, and LAUTER (2009) in their Example 8.2 as an obstruction to a straightforward generalization of the endomorphism-ring-computing algorithm of KOHEL (1996). Indeed, the presence of cycles other than at the core, such as seen in Figures 18 and 19, makes it difficult to obtain useful information about endomorphism rings by exploring the isogeny graph blindly.

In arbitrary dimension g , when a prime ℓ is completely split in the maximal order, we have argued before that the core of the isogeny graph is the 1-skeleton of a g -torus. In orders \mathcal{O} of conductor not coprime to ℓ , since not all prime ideals of norm ℓ can be invertible (otherwise ℓ itself would be), there are at most $g - 1$ of them. The construction of the covering as a Cayley graph is then identical to the maximal case except for two differences:

- \mathfrak{P} now consists of $g - 1$ ideals at the most;
- its action on $G_{\mathcal{O}}$ need not be transitive.

Since we defined our isogeny graphs as being connected components, the subgraph of horizontal isogenies in the core was always connected (in this case where we assume that ℓ completely splits and that all elements of $\mathfrak{C}(\mathcal{O})$ of norm ℓ^g arise as rational isogenies); however, there is no reason for this to happen in the cover where we have a smaller \mathfrak{P} , which is the reason for the second difference.

The graph of horizontal isogenies of $G_{\mathcal{O}}$ therefore has the topological structure of several copies of the 1-skeleton of a simplicial complex homeomorphic to the $u_{\mathcal{O}}$ -torus, for some integer $u_{\mathcal{O}} < g$. Obviously, the integer $u_{\mathcal{O}}$ is non-decreasing with respect to the order \mathcal{O} (for the inclusion order).

In the case $g = 2$, when the subgroup generated by the invertible ideals of norm ℓ^2 in $\mathfrak{C}(\mathcal{O})$ is small, we obtain an isogeny graph such as that of Figure 18. On the other hand, when it is large, its shape is similar to Figure 19.

To compute endomorphism rings, such ideals can be allowed in our relations as long as they are invertible in $\mathbb{Z}[\pi, \bar{\pi}]$. Although this has no effect on the asymptotic complexity of our method, it provides a valuable practical optimization: since computing isogenies is the bottleneck, not using *any* ideal of norm ℓ^g just because *some* are not invertible would be a loss, especially if the full ℓ -torsion conveniently lies in a small extension of the base field.




FIGURE 18. Graph of isogenies of type $(\mathbb{Z}/3)^2$ containing the Jacobian variety of the curve $y^2 = 8x^6 + 3x^5 + 7x^4 + 5x^3 + 12x^2 + 5x + 5$ over the field with 23 elements.




FIGURE 19. Graph of isogenies of type $(\mathbb{Z}/3)^2$ containing the Jacobian variety of the curve $y^2 = 10x^6 + 18x^5 + 24x^4 + 3x^3 + 33x^2 + 26x + 25$ over the field with 41 elements.

References

1836. Friedrich J. RICHELLOT.
 “Essai sur une méthode générale pour déterminer la valeur des intégrales ultra-elliptiques, fondée sur des transformations remarquables de ces transcendentes”.
 In: *Comptes Rendus de l'Académie des Sciences de Paris* 2. Pages 622–627.
1870. Carl J. THOMAE.
 “Beitrag zur Bestimmung von $\vartheta(0, 0, \dots, 0)$ durch die Klassenmoduln algebraischer Functionen”.
 In: *Journal für die reine und angewandte Mathematik* 1870.71. Pages 201–222.
 DOI: 10.1515/crll.1870.71.201.
1971. Jacques VÉLU.
 “Isogénies entre courbes elliptiques”.
 In: *Comptes Rendus de l'Académie des Sciences de Paris*. A 273. Pages 238–241.
1976. Shoki KOIZUMI.
 “Theta relations and projective normality of abelian varieties”.
 In: *American Journal of Mathematics* 98.4. Pages 865–889.
 DOI: 10.2307/2374034.
1984. Henri COHEN and Hendrik W. LENSTRA.
 “Heuristics on class groups of number fields”.
 In: *Number Theory Noordwijkerhout — Journées Arithmétiques '83*.
 Edited by Hendrik JAGER. Volume 1068. Lecture Notes in Mathematics. Springer.
 Pages 33–62. DOI: 10.1007/BFb0099440.
1984. Bruce DODSON.
 “The structure of Galois groups of CM-fields”.
 In: *Transactions of the American Mathematical Society* 283.1. Pages 1–32.
 DOI: 10.1090/S0002-9947-1984-0735406-X.
1988. Jean-Benoît BOST and Jean-François MESTRE.
 “Moyenne arithmético-géométrique et périodes des courbes de genre 1 et 2”.
 In: *Gazette des Mathématiciens* 38. Pages 36–64.
1989. Johannes BUCHMANN.
 “A subexponential algorithm for the determination of class groups and regulators of algebraic number fields”. In: *Séminaire de Théorie des Nombres, Paris*.
 Edited by Catherine GOLDSTEIN. Volume 91. Progress in Mathematics. Birkhäuser.
 Pages 27–41.

1989. James L. HAFNER and Kevin S. McCURLEY.
“A rigorous subexponential algorithm for computation of class groups”.
In: *Journal of the American Mathematical Society* 2.4. Pages 837–850.
DOI: 10.2307/1990896.
1989. George R. KEMPF.
“Linear systems on abelian varieties”. In: *American Journal of Mathematics* 111.1.
DOI: 10.2307/2374480.
1990. Jonathan PILA.
“Frobenius maps of abelian varieties and finding roots of unity in finite fields”.
In: *Mathematics of Computation* 55.192. Pages 745–763.
DOI: 10.2307/2008445.
1991. Jean-François MESTRE.
“Construction de courbes de genre 2 à partir de leurs modules”.
In: *Effective methods in algebraic geometry — MEGA '90*.
Edited by Teo MORA and Carlo TRAVERSO. Volume 94. Progress in Mathematics.
Birkhäuser. Pages 313–334.
1996. David R. KOHEL.
“Endomorphism rings of elliptic curves over finite fields”.
PhD thesis. University of California at Berkeley.
URL: <http://echidna.maths.usyd.edu.au/kohe1/pub/thesis.pdf>.
1997. Wieb BOSMA, John CANNON, and Catherine PLAYOUST.
“The Magma algebra system: the user language”.
In: *Journal of Symbolic Computation* 24.3–4. Pages 235–265.
DOI: 10.1006/jSCO.1996.0125.
1997. Henri COHEN, Francisco DIAZ Y DIAZ, and Michel OLIVIER.
“Subexponential algorithms for class group and unit computations”.
In: *Journal of Symbolic Computation* 24.3–4. Special issue on computational algebra
and number theory: proceedings of the first MAGMA conference. Pages 433–441.
DOI: 10.1006/jSCO.1996.0143.
1998. Paul VAN WAMELEN.
“Equations for the Jacobian of a hyperelliptic curve”.
In: *Transactions of the American Mathematical Society* 350.8. Pages 3083–3106.
DOI: 10.1090/S0002-9947-98-02056-X.
2000. Pierrick GAUDRY and Robert HARLEY.
“Counting points on hyperelliptic curves over finite fields”.

- In: *Algorithmic Number Theory — ANTS-IV*. Edited by Wieb BOSMA.
Volume 1838. Lecture Notes in Computer Science. Springer. Pages 313–332.
DOI: 10.1007/10722028_18.
2002. Mireille FOUQUET and François MORAIN.
“Isogeny volcanoes and the SEA algorithm”.
In: *Algorithmic Number Theory — ANTS-V*.
Edited by Claus FIEKER and David R. KOHEL. Volume 2369.
Lecture Notes in Computer Science. Springer. Pages 47–62.
DOI: 10.1007/3-540-45455-1_23.
2002. Steven D. GALBRAITH, Florian HESS, and Nigel P. SMART.
“Extending the GHS Weil descent attack”.
In: *Advances in Cryptology — EUROCRYPT ’02*. Edited by Lars R. KNUDSEN.
Volume 2332. Lecture Notes in Computer Science. Springer. Pages 29–44.
DOI: 10.1007/3-540-46035-7_3.
2007. David R. KOHEL.
ECHIDNA: Algorithms for elliptic curves and higher dimensional analogues.
URL: <http://echidna.maths.usyd.edu.au/>.
2008. Robert CARLS, David R. KOHEL, and David LUBICZ.
“Higher dimensional 3-adic CM construction”.
In: *Journal of Algebra* 319.3. Pages 971–1006.
DOI: 10.1016/j.jalgebra.2007.11.016.
2009. Reinier BRÖKER, David GRUENEWALD, and Kristin E. LAUTER.
Explicit CM-theory in dimension 2. arXiv.org: 0910.1848.
2009. Jean-Marc COUVEIGNES.
“Linearizing torsion classes in the Picard group of algebraic curves over finite fields”.
In: *Journal of Algebra* 321.8. Pages 2085–2118.
DOI: 10.1016/j.jalgebra.2008.09.032.
2009. Kirsten EISENTRÄGER and Kristin E. LAUTER.
“A CRT algorithm for constructing genus 2 curves over finite fields”.
In: *Arithmetic, Geometry and Coding Theory — AGCT ’10*.
Edited by François RODIER and Serge VLADUT. Volume 21. Séminaires et Congrès.
Société Mathématique de France. Pages 161–176.
2009. David LUBICZ and Damien ROBERT.
Computing isogenies between abelian varieties. arXiv.org: 1001.2016.

-
2009. Markus WAGNER.
“Über Korrespondenzen zwischen algebraischen Funktionenkörper”.
PhD thesis. Technische Universität Berlin.
URL: <http://www.math.tu-berlin.de/~wagner/Diss.pdf>.
2010. Gaetan BISSON, Romain COSSET, and Damien ROBERT.
AVIsogenies, a library for computing isogenies between abelian varieties.
Registered at the Agence pour la Protection des Programmes under reference
IDDN.FR.001.440011.000.R.P.2010.000.10000.
URL: <http://avisogenies.gforge.inria.fr/>.
2010. Pierrick GAUDRY and Éric SCHOST.
Genus 2 point counting over prime fields. HAL-INRIA: 00542650.
2011. Romain COSSET and Damien ROBERT.
Computing (ℓ, ℓ) -isogenies in polynomial time on Jacobians of genus 2 curves.
IACR ePrint: 2011/143.

Index

- algorithm
 - generic, 11, 147
 - Las Vegas, 13
 - probabilistic, 13
- attack
 - brute-force, 5
 - side-channel, 4
- baby-step giant-step method, 12
- birthday paradox, 13
- cipher, 3
 - block, 4
 - stream, 4
 - text, 3
- class number, 50
- collision, 12, 148
 - proper, 13
- complex multiplication, 46
 - action
 - plain, 54
 - polarized, 54
 - field, 46
 - reflex, 58
 - method, 70
 - type, 47, 58
 - induced, 58
 - norm, 58
 - primitive, 58
 - trace, 58
- conductor, 50
 - gap, 169
- coordinate
 - affine, 26
 - projective, 26
 - ring, 27
- correspondence, 100
- curve
 - algebraic, 32
 - anomalous, 18
 - elliptic, 32
 - supersingular, 49, 102
 - family, 75
 - hyperelliptic, 32
- degree
 - embedding, 34, 65
 - morphism, 27
- density theorem, 138
- dimension, 17, 27
- distance
 - order, 88
 - variation, 147
- distribution

- pushforward, 147
- quasi-uniform, 143
- encryption
 - asymmetric, 8
 - homomorphic, 10
 - symmetric, 3
- endomorphism
 - Frobenius, 29
 - monoid, 27
 - Verschiebung, 48
- essentially linear, 135
- field
 - base, 25
 - constant extension, 100
 - function, 27, 100
 - of definition, 26
 - ring class, 56
- floor of rationality, 92
- function
 - hash, 5
 - one-way, 5
 - pseudorandom, 13
 - zeta, 29
- genus, 32
- group
 - algebraic, 27
 - coding, 11
 - generic, 11
- hypothesis
 - extended Riemann, 15
 - generalized Riemann, 15
- ideal
 - cardinality, 121
 - constant, 100
 - invertible, 50
 - kernel, 56
 - principal, 111
 - short, 115
- invariant, 56, 60
 - j -invariant, 38, 71
 - Igusa, 39
- isogenous, 37
- isogeny, 37
 - ℓ -isogeny, 37
 - dual, 37
 - horizontal, 87
 - separable, 37
 - type, 37
 - vertical, 87
- key, 3
 - generation, 8
 - private, 8
 - public, 8
 - symmetric, 3
- lattice, 30
 - diameter, 140
 - of relations, 51
- map
 - distortion, 69
 - rational, 27
 - scalar multiplication, 33
- morphism, 26, 28
- number field sieve, 16
- one-time pad, 4
- order, 48
 - certificate, 111
 - lattice, 110
 - localization, 93
 - maximal, 48, 93
 - minimal, 48, 49, 93

- relation, 115
- test, 110, 112
- verification, 111
- pairing, 10, 34
 - Weil, 35
- Picard group, 50
- plaintext, 3
- point, 25
- polarization, 31
 - principal, 31
- polynomial
 - class, 71
 - Hilbert, 71
 - modular, 39
 - Weil, 38
- problem
 - Diffie–Hellman, 7
 - discrete logarithm, 6
 - knapsack, 8
 - pairing inversion, 35
 - short product, 8, 146
 - subset sum, 8, 146
- quasi-linear, 53
- random oracle, 6
- reduction
 - at a prime, 30
 - good, 48
 - cryptographic, 6
 - divisor, 32, 33
 - ideal, 52
- secrecy
 - computational, 4
 - perfect, 3
- sequence density, 147
- short product, 8
- signature, 8
- space
 - affine, 26
 - moduli, 33
 - projective, 26
 - Siegel upper half, 30
- string, 3
- subgroup, 28
 - p -Sylow, 12
 - torsion
 - full, 34
- theta
 - constant, 38, 39
 - coordinate, 38
 - level n , 164
 - function, 38, 39
- variety, 26
 - abelian, 17, 28
 - ordinary, 34
 - principally polarized, 31
 - supersingular, 68
 - absolutely irreducible, 26
 - absolutely simple, 29
 - affine, 25
 - Albanese, 32
 - Jacobian, 32
 - nonsingular, 27
 - projective, 26
 - quasiprojective, 26
 - superspecial, 61
- volcano, 90
 - branch, 171
 - core, 171
 - covering, 171
 - crater, 95, 171
 - layer, 171

Summary

ENDOMORPHISM RINGS IN CRYPTOGRAPHY

Modern communications heavily rely on cryptography to ensure data integrity and privacy. Over the past two decades, very efficient, secure, and featureful cryptographic schemes have been built on top of abelian varieties defined over finite fields. This thesis contributes to several computational aspects of ordinary abelian varieties related to their endomorphism ring structure.

This structure plays a crucial role in the construction of abelian varieties with desirable properties. For instance, pairings have recently enabled many advanced cryptographic primitives; generating abelian varieties endowed with efficient pairings requires selecting suitable endomorphism rings, and we show that more such rings can be used than expected.

We also address the inverse problem, that of computing the endomorphism ring of a prescribed abelian variety, which has several applications of its own. Prior state-of-the-art methods could only solve this problem in exponential time, and we design several algorithms of subexponential complexity for solving it in the ordinary case.

For elliptic curves, our algorithms are very effective and we demonstrate their practicality by solving large problems that were previously intractable. Additionally, we rigorously bound the complexity of our main algorithm assuming solely the extended Riemann hypothesis. As an alternative to one of our subroutines, we also consider a generalization of the subset sum problem in finite groups, and show how it can be solved using little memory.

Finally, we generalize our method to higher-dimensional abelian varieties, for which we rely on further heuristic assumptions. Practically speaking, we develop a library enabling the computation of isogenies between abelian varieties; using this important building block in our main algorithm, we apply our generalized method to compute several illustrative and record examples.

Research Prospects

In this thesis, we effectively exploited complex multiplication theory to compute the endomorphism ring structure of a prescribed ordinary abelian variety defined over a finite field. For elliptic curves, we were additionally able to rigorously analyze our algorithms, and we believe their asymptotic complexity leaves little room for improvement.

On the other hand, although we described a practical method for varieties of dimension $g = 2$, several topics remain to be explored for $g \geq 2$:

- We dealt with orders having identical Picard groups locally, using the method of Eisen-träger and Lauter. As its complexity is exponential in the valuation of the conductor gap, this is however impractical in certain cases. It would be interesting to address this by developing a generalization of Kohel's techniques to dimension two and more.
- Having a deeper insight on the structure of isogeny graphs would certainly help solving the above, and we note that recent work on elliptic curves by Joux and Ionica offers promising perspectives of developments on this matter in higher dimension.
- Besides the extended Riemann hypothesis, heuristics we relied on should be further analyzed, such as the assumption that norms of LLL-reduced ideals are as smooth as random integers, or that complex multiplication applies to non-maximal orders.
- The convenient structure of Jacobian varieties was used to draw points at random, and to uniquely identify isomorphism classes. Using our method beyond dimension three would require to solely work in theta-coordinates, using the Heisenberg group for the latter, and finding an efficient way of doing the former.

Closely connected topics include the computation of class polynomials and of modular polynomials; it is only natural that they should benefit from further exploiting complex multiplication theory as well. For elliptic curves, this was done successfully for both problems by Sutherland, and by Bröker, Lauter, and Sutherland, respectively.

However, similar work remains to be done in higher dimension: although substantive improvements have been made on it over the past few years, the computation of class polynomials remains a topic of active study, albeit particularly unexplored in the case of non-maximal orders. On the other hand, modular polynomials have not attracted many research, due to their prohibitive height; it would be challenging to improve on this and compute more such polynomials, as an alternative to explicit isogeny computation.

Finally, more of the code written during this thesis should be optimized, fully automated, and cleaned up for inclusion in open software packages, as experimentation using efficient computer routines becomes increasingly important to research activities in many fields.

Curriculum Vitæ

Gaetan BISSON was born on the 4th of November 1984 in Les Ulis, France. After obtaining the *diplôme national du brevet* in 1999 at the collège Paul Arène (Peymeinade, France) and the *baccalauréat* in 2002 at the lycée Amiral (Grasse, France), he attended *classes préparatoires aux grandes écoles*, majoring in mathematics, at the lycée Masséna (Nice, France).

In August 2004 he was admitted to the École Normale Supérieure (Paris, France) where he received a *licence* and a *maîtrise* of mathematics in July 2005. After passing the *agrégation* of mathematics in July 2006, he pursued a master in analysis, arithmetic, and geometry at the Université Paris-Sud XI (Orsay, France) which he completed in October 2007 together with the *magistère* of fundamental and applied mathematics, and computer science; his master's thesis, entitled "On the generation of pairing-friendly elliptic curves," was realized in the number theory group of the Tokyo Institute of Technology (Tokyo, Japan).

Funded by a grant from the French Ministry of Research, he started a joint PhD project at the Institut National Polytechnique de Lorraine (Nancy, France) and at the Technische Universiteit Eindhoven (Eindhoven, The Netherlands) in February 2008, of which the research results are presented in this dissertation.