

On the Effectiveness of Isogeny Walks for Extending Cover Attacks on Elliptic Curves

by

Randy Yee

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Combinatorics and Optimization

Waterloo, Ontario, Canada, 2016

© Randy Yee 2016

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Cryptographic systems based on the elliptic curve discrete logarithm problem (ECDLP) are widely deployed in the world today. In order for such a system to guarantee a particular security level, the elliptic curve selected must be such that it avoids a number of well-known attacks. Beyond this, one also needs to be wary of attacks whose reach can be extended via the use of isogenies. It is an open problem as to whether there exists a field for which the isogeny walk strategy can render all elliptic curves unsuitable for cryptographic use.

This thesis provides a survey of the theory of elliptic curves from a cryptographic perspective and overviews a few of the well-known algorithms for computing elliptic curve discrete logarithms. We perform some experimental verification for the assumptions used in the analysis of the isogeny walk strategy for extending Weil descent-type cover attacks, and explore its applicability to elliptic curves of cryptographic size. In particular, we demonstrate for the first time that the field $\mathbb{F}_{2^{150}}$ is partially weak for elliptic curve cryptography.

Acknowledgements

I would like to thank my supervisor Alfred Menezes for his mentorship and guidance during my time in the Combinatorics and Optimization Master's program. My gratitude also goes out to the administrative staff for the ever-friendly and helpful presence, and my readers David Jao and Douglas Stinson for their valuable insights.

Lastly I would like to thank my family and friends for their support and encouragement. Special thanks to Moriah Pellowe and Luis Ruiz-Lopez for their profound influence on me as I tried to find my feet in Waterloo. You made this place home.

Table of Contents

Author's Declaration	ii
Abstract	iii
List of Figures	vii
1 Introduction	1
2 Mathematical Background	5
2.1 Elliptic Curves	5
2.1.1 Torsion Points	7
2.1.2 Maps Between Curves	8
2.2 Algebraic Number Theory	11
2.2.1 Quadratic Number Fields	11
2.2.2 Ideals and Factorization	12
3 Isogenies	14
3.1 Preliminaries	14
3.2 Standard Form	15
3.3 Kernels of Isogenies	17
3.4 Endomorphism Rings	18
3.5 Isogeny Graphs	20
3.5.1 The Structure of Isogeny Graphs	22

4	Cryptographic Background	24
4.1	The Elliptic Curve Discrete Log Problem	24
4.1.1	Pollard's Rho	24
4.2	Index Calculus	26
4.3	Point Decomposition	26
4.3.1	Summation Polynomials	27
4.3.2	Decomposition Attacks in Higher Genus	29
4.4	Weil Descent	29
4.5	Cover and Decomposition Attack	30
4.6	Isogeny walks	31
4.6.1	Modeling the Walk	33
5	Distribution of Curves	35
5.1	Distribution of Isomorphism Classes	35
5.1.1	Small Fields	38
5.1.2	Medium Fields	42
5.2	Distribution of Weak Curves	42
5.2.1	Small Fields	46
5.2.2	Medium Fields	50
6	Experiments in Large Fields	55
6.1	Discrete Logarithms in \mathbb{F}_{q^3}	55
6.2	The Field $\mathbb{F}_{2^{150}}$	56
6.3	The Field $\mathbb{F}_{2^{210}}$	59
7	Concluding Remarks	62
	References	64

List of Figures

5.1	Plot of isogeny class sizes for $\mathbb{F}_{2^{18}}$	39
5.2	Plot of isogeny class sizes for $\mathbb{F}_{2^{24}}$	40
5.3	Plot of isogeny class sizes for $\mathbb{F}_{2^{30}}$	41
5.4	Plot of isogeny class sizes for $\mathbb{F}_{2^{48}}$,	43
5.5	Plot of isogeny class sizes for $\mathbb{F}_{2^{54}}$,	44
5.6	Weak orbits in $\mathbb{F}_{2^{18}}$	47
5.7	Weak orbits in $\mathbb{F}_{2^{24}}$	48
5.8	Weak orbits in $\mathbb{F}_{2^{30}}$	49
5.9	Weak orbits in $\mathbb{F}_{2^{42}}$	52
5.10	Weak orbits in $\mathbb{F}_{2^{48}}$	53
5.11	Weak orbits in $\mathbb{F}_{2^{54}}$	54

Chapter 1

Introduction

Let E be an elliptic curve defined over a finite field $K = \mathbb{F}_{q^n}$. The elliptic curve discrete logarithm problem (ECDLP) in the group $E(K)$ is the following: Given two points P and Q in $E(K)$ with $\text{ord}(P) = r$ and $Q \in \langle P \rangle$, find an integer α such that $\alpha P = Q$. The presumed intractability of this problem has been the subject of much study in recent years and has formed the basis for the security of many modern cryptosystems.

The security level of an ECDLP-based cryptosystem is dictated by the expected number of operations the fastest known algorithm would need in order to solve arbitrary ECDLP instances in $E(K)$. In most cases, this can be determined by considering the computational effort needed by Pollard's rho method [30, 39]. However one must select the parameters of the elliptic curve carefully in order to avoid certain known attacks. The value r should be a large prime so that the Pohlig-Hellman [29] and Pollard's rho attacks are both infeasible (we assume $r \approx q^n$). To avoid the Weil and Tate pairing attacks [8, 22], r must be selected so that it does not divide $q^{ni} - 1$ for $1 \leq i \leq C$, where C is big enough for it to be computationally infeasible to find discrete logarithms in $\mathbb{F}_{q^{nC}}$. Moreover, the curve selected should not be \mathbb{F}_q -anomalous ($\#E(K) \neq q^n$) so that the attacks of [31, 33, 35] are avoided.

In 2001, Frey [7] introduced the idea of Weil descent, a technique that may be used to map instances of the ECDLP in elliptic curves over extension fields \mathbb{F}_{q^n} into the DLP of a higher genus covering curve ($g \geq n$) defined over \mathbb{F}_q . This idea was fleshed out for binary fields by Gaudry, Hess and Smart in [13] and later for odd characteristic by Diem [3]. They showed how one might apply this technique to reduce ECDLP instances into those of hyperelliptic curves defined over a smaller field. In particular, the problem would be embedded into the *Jacobian* of a hyperelliptic curve \mathcal{C} of genus g over \mathbb{F}_q , denoted $J_{\mathcal{C}}(\mathbb{F}_q)$. This reduction is commonly referred to as the *GHS attack*. The existence of subexponential

time algorithms in this new setting can allow the DLP to be solved more quickly.

In general there are many elliptic curves over a fixed field K , so one can expect to be able to select one that avoids each of the attacks described above. However, an additional pitfall can occur if a curve's susceptibility to an attack is not invariant within its isogeny class (that is, those curves E' which satisfy $\#E(K) = \#E'(K)$). In this case, it may be possible to use a series of small degree isogenies to map the DLP of one curve into another for which such an attack is effective. This technique can drastically increase the range of a particular attack. For example, Jacobson, Menezes and Stein [17] showed that for the field $\mathbb{F}_{2^{155}}$, only 2^{33} out of a possible 2^{156} isomorphism classes of elliptic curves defined over $\mathbb{F}_{2^{155}}$ would yield a covering curve of small enough genus for the GHS to be feasible. However in [11], it was shown that through isogeny walks roughly 2^{104} classes were vulnerable to the GHS attack; a significant increase in its range.

With this novel approach to attacking the ECDLP, concern arose over whether there existed fields in which most if not all elliptic curves would be vulnerable to some attack that performed better than Pollard's rho. This question motivated Menezes, Teske and Wang [25] to define formally the notion of *bad* and *weak* fields for elliptic curve cryptography (ECC).

Definition 1.0.1. *A finite field \mathbb{F}_q is said to be bad for elliptic curve cryptography if the following conditions are satisfied:*

1. *for some elliptic curves E/\mathbb{F}_q , solving the ECDLP in $E(\mathbb{F}_q)$ using Pollard's rho method is intractable using existing computer technology; and*
2. *algorithms are known that can feasibly solve (using existing computer technology) any ECDLP instance for any elliptic curve over \mathbb{F}_q .*

Definition 1.0.2. *A finite field \mathbb{F}_q is said to be weak for elliptic curve cryptography if the following conditions are satisfied:*

1. *for some elliptic curve E over \mathbb{F}_q , solving the ECDLP in $E(\mathbb{F}_q)$ using Pollard's rho method is intractable using existing computer technology; and*
2. *algorithms are known for which any ECDLP instance for any elliptic curve over \mathbb{F}_q can be solved in significantly less time than it takes Pollard's rho method to solve the hardest ECDLP instances over \mathbb{F}_q .*

A bad field should certainly be avoided for use in cryptography, and one ought to be wary of weak fields since this may be evidence towards badness. In the same vein, [24]

defined the notion of a *partially weak* field, which only requires that for a non-negligible portion of all elliptic curves, the ECDLP may be solved significantly faster than it takes Pollard’s rho to solve the hardest instances. One may define a *partially bad* field similarly.

Currently no bad or partially bad fields for ECC are known and it is an open problem as to whether such fields exist. The authors of [25] provided a convincing argument that fields of the form \mathbb{F}_{2^N} , where $N \in [185, 600]$ is divisible by 5 are weak, and proposed $\mathbb{F}_{2^{210}}$ as a strong candidate for being a bad field.

In 2012, Joux and Vitse [19] proposed the use of a combination attack on elliptic curves defined over composite extension fields which utilizes both the Weil descent cover attack and the decomposition technique. The rough idea is to make use of the tower of extensions $\mathbb{F}_{q^n}/\mathbb{F}_q$ and $\mathbb{F}_q/\mathbb{F}_p$, where $q = p^m$ for some integer m . One would first use Weil descent on the extension $\mathbb{F}_{q^n}/\mathbb{F}_q$ to transfer the DLP into the Jacobian of a (hyperelliptic) curve defined over \mathbb{F}_q , and then apply the point decomposition attack of Nagao [28] with the second extension.

In particular, [19] analyzed certain elliptic curves E defined over composite extension fields of the form \mathbb{F}_{p^6} and \mathbb{F}_{p^4} . Using a modified decomposition technique, they found that their attack was more effective than Pollard’s rho in both scenarios, with the improvements being most pronounced in the case of \mathbb{F}_{p^6} using the tower $\mathbb{F}_{p^6} - \mathbb{F}_{p^2} - \mathbb{F}_p$. With this method, an elliptic curve defined over the 149-bit extension field \mathbb{F}_{p^6} where $p = 33554467$ was successfully attacked, whereas no preexisting algorithms would have been practical.

In order for the attack to outperform generic algorithms, the genus of the intermediate cover must not be too large. When this not the case, it may be possible to use the isogeny walk strategy to transfer the DLP from E to a more vulnerable curve E' . Depending on the size of the isogeny class and the number of weak curves contained in it, the isogeny walk can be the dominating factor of the algorithm’s complexity.

A natural extension of the above work is to verify experimentally the effectiveness of their attack in fields of characteristic 2. In particular, it seems reasonable to look at fields of the form $\mathbb{F}_{2^{6m}}$ for some integer m , since degree-6 extensions were the most effective scenario considered in [19]. The equivalent approach in this setting would be to use Weil descent to find a genus-3 hyperelliptic covering curve defined over $\mathbb{F}_{2^{2m}}$, and then apply Nagao’s decomposition over the extension $\mathbb{F}_{2^{2m}}/\mathbb{F}_{2^m}$. We would also like to consider a field of equivalent size to Joux and Vitse’s main example, namely $\mathbb{F}_{2^{150}}$, and determine whether the isogeny walk is effective enough to consider this field partially bad for ECC.

The purpose of this work will be to test some of the assumptions made in previous work pertaining to the distribution of curves vulnerable to the GHS attack amongst isogeny classes, and to determine whether the field $\mathbb{F}_{2^{150}}$ is indeed partially bad for use in elliptic

curve cryptography. In particular, we would like to ascertain the following (note that all experiments are done in fields of the form $\mathbb{F}_{2^{6m}}$):

1. Whether it is reasonable to treat isogeny classes as close to equal in size;
2. if it is reasonable to assume that as the field size grows most isogeny classes contain curves vulnerable to Weil descent;
3. whether or not the distribution of vulnerable curves takes on a uniformly random distribution amongst isogeny classes; and
4. whether the ECDLP in a field of cryptographic size may be effectively attacked using the isogeny walk strategy.

The remainder of this thesis is organized as follows: First we recall in Chapter 2 some of the essential mathematical facts needed to understand elliptic curves and the isogenies between them. In Chapter 3, we discuss isogenies in more detail and the structure of the so-called isogeny graph. Chapter 4 deals with the cryptographic background related to solving the ECDLP, providing an overview of the relevant attacks including the work of Joux and Vitse as well as the details of the isogeny walk strategy. In Chapter 5 we describe the various assumptions and results on the distribution of curves which are used in the analyses of the isogeny walk strategy, and provide some experimental evidence to back these assumptions in Chapter 6. Finally we end with directions for future work in Chapter 7.

Chapter 2

Mathematical Background

2.1 Elliptic Curves

Let p be a prime and $k = \mathbb{F}_q$ a finite field of size $q = p^n$ for some positive integer n . We define $K = \mathbb{F}_{q^m}$ to be an extension field of degree m over k , noting that K will have composite extension degree over \mathbb{F}_p . An elliptic curve E over k (sometimes denoted E/k) can be thought of simply as an equation of the form

$$E : Y^2 + a_1XY + a_3Y = X^3 + a_2X^2 + a_4X + a_6, \quad (2.1)$$

where the $a_i \in k$ satisfy a non-singularity condition. We will occasionally refer to k as the *field of definition*. If k is not of characteristic 2 or 3, we can without loss of generality use a linear change of variables to transform the above equation into the form

$$E : Y^2 = X^3 + aX + b, \quad a, b \in k. \quad (2.2)$$

The above is known as the *short Weierstrass equation* of E . Here the non-singularity condition is met when $4a^3 + 27b^2 \neq 0$. When k is of characteristic 2, an *ordinary* elliptic curve may be transformed into the form

$$E : Y^2 + XY = X^3 + aX^2 + b, \quad (2.3)$$

where $a, b \in k$. Non-singularity is met when $b \neq 0$.

Before we can explain what we mean by an *ordinary* elliptic curve, a few definitions will be required.

Definition 2.1.1. Let E/k be an elliptic curve and $X, Y \in k$. The ordered pair (X, Y) is called a **k -rational point of E** if X and Y satisfy the equation of E . For technical reasons, we also include a special point \mathcal{O}_E , called the point at infinity. The set of all k -rational points of E is denoted by $E(k)$.

Remark 2.1.1. It is worth noting that if E is defined over k , we are not confined to only looking at subsets of $E(k)$. In many cases we will need to consider points whose coordinates lie in an extension field of k such as its algebraic closure \bar{k} . In these instances, if L is an extension field of k , we will use the notation $E(L)$ to clarify the points of interest.

It is well-known that one may define a binary operation on $E(\bar{k})$ such that the conditions of a group are satisfied. This operation is known as the *elliptic curve addition law*. We note that for any intermediate field $k \subseteq k' \subseteq \bar{k}$, the set $E(k')$ is closed under the above operation and thus may be regarded as a subgroup of $E(\bar{k})$ (and hence a group in its own right).

Traditionally, the cryptographic use of elliptic curves has been centred around the group of k -rational points of E with the security of these types of cryptosystems being more or less dictated by the size or **order** of E over k , which we denote by $\#E(k)$. With this application in mind, we will begin to explore a few of the basic results pertaining to elliptic curves and their use in cryptography. Our major references will be [10], [40] and [15].

Theorem 2.1.2. [15, Theorem 3.7] Let E/k be an elliptic curve where $k = \mathbb{F}_q$. Then

$$q + 1 - 2\sqrt{q} \leq \#E(k) \leq q + 1 + 2\sqrt{q}.$$

The above result is known as Hasse's Theorem, and the interval $[q+1-2\sqrt{q}, q+1+2\sqrt{q}]$ is called the Hasse interval. An alternative formula for the group order is given by

$$\#E(k) = q + 1 - t,$$

where $|t| \leq 2\sqrt{q}$. The value t is the **trace of Frobenius**, which will be described in greater detail in the next sections.

Definition 2.1.2. Let $p = \text{char}(k)$. An elliptic curve E/k is called **supersingular** if $p|t$. If this is not the case then we say E is an **ordinary** elliptic curve.

The distinction between ordinary and supersingular curves will be made less superficial in Section 2.1.1, but for the time being Definition 2.1.2 provides a simple means of distinguishing the two types of curves.

It is worth mentioning that computing $\#E(k)$ may be done efficiently using techniques such as the Schoof-Elkies-Atkins (SEA) algorithm; so determining t is straightforward in practice. We now present a theorem that is useful for determining the possible values for $\#E(k)$.

Theorem 2.1.3. [15, Theorem 3.8] *Let k be a field of order $q = p^n$. There exists an elliptic curve E/k with $\#E(k) = q + 1 - t$ if and only if one of the following conditions holds:*

(i) $t \not\equiv 0 \pmod{p}$ and $t^2 \leq 4q$.

(ii) n is odd and either

(a) $t = 0$,

(b) $t^2 = 2q$ and $p = 2$, or

(c) $t^2 = 3q$ and $p = 3$.

(iii) n is even and either

(a) $t^2 = 4q$,

(b) $t^2 = q$ and $p \not\equiv 1 \pmod{3}$, or

(c) $t = 0$ and $p \not\equiv 1 \pmod{4}$.

We will mostly be concerned with ordinary elliptic curves over binary fields (i.e. \mathbb{F}_{2^m} for some positive integer m). In this context, Theorem 2.1.3 tells us that for all even values l that lie in the Hasse interval, there exists an elliptic curve E/k with $\#E(k) = l$. For the remainder of this thesis, we will largely be neglecting the supersingular case. We may state some results that pertain to supersingular curves for completeness, but it should be understood that such curves are not the focus of this work.

2.1.1 Torsion Points

Let E/k be an elliptic curve and n a positive integer. We define the set of **n -torsion points** to be the set

$$E[n] = \{P \in E(\bar{k}) : nP = \mathcal{O}_E\}.$$

A simple check will verify that the above is a group under addition. The n -torsion points will play an important role in the discussion of isogenies in Chapter 3, so to this end we present some results regarding their structure as groups.

Theorem 2.1.4. [37, Chapter 7] *Let E be an elliptic curve over a field k of characteristic p , and suppose l is a prime and e is a positive integer. If $\gcd(l, p) = 1$, then*

$$E[l^e] \simeq \mathbb{Z}/l^e\mathbb{Z} \oplus \mathbb{Z}/l^e\mathbb{Z}$$

as groups. If $\gcd(l, p) \neq 1$, then

$$E[l^e] \simeq \mathbb{Z}/l^e\mathbb{Z} \quad \text{or} \quad \{0\}.$$

Using this theorem, one can subsequently use the classification of finite abelian groups to determine the structure of $E[n]$ for any n . We state this result formally below.

Theorem 2.1.5. [40, Theorem 3.2] *Let E be an elliptic curve over a field k of characteristic $p > 0$ and suppose that n is a positive integer. If $\gcd(n, p) = 1$, then*

$$E[n] \simeq \mathbb{Z}/n\mathbb{Z} \oplus \mathbb{Z}/n\mathbb{Z}$$

as groups. If $\gcd(n, p) \neq 1$, then write $n = p^r n'$ where $p \nmid n'$. Then

$$E[n] \simeq \mathbb{Z}/n'\mathbb{Z} \oplus \mathbb{Z}/n'\mathbb{Z} \quad \text{or} \quad \mathbb{Z}/n\mathbb{Z} \oplus \mathbb{Z}/n'\mathbb{Z}.$$

If one considers the group $E[p]$, the two structural possibilities are consistent with our earlier definition of supersingular/ordinary curves. By this we mean the following:

Theorem 2.1.6. *Using the same notation as above, E is supersingular if and only if $E[p] \simeq \{0\}$.*

The group of points $E[n]$ ties in with an important class of functions defined on E known as **multiplication-by- n -maps**. They are denoted $[n] : E(\bar{k}) \rightarrow E(\bar{k})$ and defined such that

$$[n](P) = nP,$$

for all $P \in E(\bar{k})$. The main detail to note is that $E[n]$ is the kernel of $[n]$.

2.1.2 Maps Between Curves

Suppose E_1 and E_2 are two elliptic curves defined over k . A map $\varphi : E_1(\bar{k}) \rightarrow E_2(\bar{k})$ is called a **rational map over k** if there exist rational functions f and g with coefficients in k such that

$$\varphi(P) = (f(P), g(P))$$

for all $P \in E_1(\bar{k})$. We will write $\varphi : E_1 \rightarrow E_2$ to mean a function $E_1(\bar{k}) \rightarrow E_2(\bar{k})$ unless otherwise stated.

In algebraic geometry, the concept of a rational map may be specialized to a class of functions known as **morphisms**, which respect the geometry of the objects on which they

act. A rough explanation is that rational maps need not be ‘defined’ for all $P \in E(\bar{k})$, and morphisms are those rational maps which are. When dealing with maps between elliptic curves, Lemma 7.3.6 of [10] tells us that every non-constant rational map over k is a morphism. In the interest of simplification, any apparently undefined values (those points that evaluate to zero in the denominator of f or g) are mapped to \mathcal{O}_E by definition.

Morphisms have a number of useful properties, one of which is the following:

Theorem 2.1.7. [10, Theorem 8.2.1] *Morphisms of curves are either constant or surjective.*

As noted earlier, the set of \bar{k} -rational points of an elliptic curve form a group, and in typical mathematical fashion we will be interested in maps that preserve the structure of this group. The following result provides us with a necessary condition on morphisms to achieve this.

Theorem 2.1.8. [10, Theorem 9.2.1] *Given E_1 and E_2 as above, a morphism $\varphi : E_1 \rightarrow E_2$ such that $\varphi(\mathcal{O}_{E_1}) = \mathcal{O}_{E_2}$ is a group homomorphism.*

This leads naturally to the concept of isomorphism of elliptic curves.

Definition 2.1.3. *Two curves E_1/k and E_2/k will be called **isomorphic** over k if there exist morphisms $\varphi : E_1 \rightarrow E_2$ and $\psi : E_2 \rightarrow E_1$ over k such that*

$$\varphi \circ \psi = id_{E_2} \text{ and } \psi \circ \varphi = id_{E_1}.$$

There is a convenient computational definition for determining whether two curves are isomorphic based on their Weierstrass equations. A general definition can be found in Chapter 4 of [15], but we will only present the specializations for short Weierstrass equations in the cases when $\text{char}(k) = 2$ and $\text{char}(k) \neq 2, 3$:

Proposition 2.1.9. [15, Theorem 3.15]. *Let $k = \mathbb{F}_q$ be a finite field with $\text{char}(k) \neq 2, 3$. Two elliptic curves $E_1/k : Y^2 = X^3 + aX + b$ and $E_2/k : Y^2 = X^3 + a'X + b'$ are isomorphic over k if and only if there exists $u \in k^*$ such that*

$$u^4 a' = a \quad \text{and} \quad u^6 b' = b.$$

Given that such a u exists, the change of variables $(X, Y) \rightarrow (u^2 X, u^3 Y)$ transforms E_1 to E_2 .

Proposition 2.1.10. [15, Theorem 3.18]. Let $k = \mathbb{F}_{2^m}$ be a field of characteristic 2. Two ordinary elliptic curves $E_1 : Y^2 + XY = X^3 + aX^2 + b$ and $E_2 : Y^2 + XY = X^3 + a'X^2 + b'$ defined over k are isomorphic over k if and only if $b = b'$ and $\text{Tr}(a) = \text{Tr}(a')$.

$\text{Tr}(c)$ denotes the **trace** of an element c using the extension $\mathbb{F}_{2^m}/\mathbb{F}_2$ and can be calculated with the formula

$$\text{Tr}(c) = c + c^2 + c^{2^2} + \cdots + c^{2^{m-1}}.$$

Isomorphisms will be important for us due to the following fact [15, Section 3.1.5]: if two curves E_1 and E_2 defined over k are isomorphic over k , then $E_1(k)$ and $E_2(k)$ are isomorphic as groups. Isomorphism is an equivalence relation, and the equivalence classes are referred to as *isomorphism classes*. For the most part, the properties we are interested in will be preserved by isomorphisms, therefore we will often deal with isomorphism classes rather than the individual Weierstrass equations. Of particular importance is the fact that instances of the ECDLP are equivalent in isomorphic curves.

A necessary condition for elliptic curves to be isomorphic over k makes use of the *j-invariant*.

Definition 2.1.4. Let k be a field of characteristic not equal to 2 or 3, and E/k an elliptic curve with short Weierstrass equation 2.2. The **j-invariant** of E is defined to be

$$j = j(E) = 1728 \frac{4a^3}{4a^3 + 27b^2}.$$

When the characteristic of k is 2 and E is in the form (2.3), then $j(E) = 1/b$.

The condition is encompassed by the theorem below.

Theorem 2.1.11. [10, Theorem 9.3.6] Let k be a field and $E_1/k, E_2/k$ two elliptic curves. There is an isomorphism from E_1 to E_2 over \bar{k} if and only if $j(E_1) = j(E_2)$.

A simple consequence of this is that if two elliptic curves are isomorphic over a field k , then they are isomorphic over any extension of k , and hence have the same *j-invariant*. One might expect that curves which are isomorphic over some extension of k , but not necessarily k itself should be somehow related. Indeed, there is a convenient formula relating their orders. Before we state this theorem though, we introduce the notion of twist curves.

Definition 2.1.5. Let E/k be an elliptic curve. A **twist** of E is an elliptic curve E'/k such that there is an isomorphism $\varphi : E \rightarrow E'$ over \bar{k} (but not necessarily k). One notes that $j(E) = j(E')$.

A twist E' of E will be called *trivial* if it is isomorphic over k to E , and *non-trivial* otherwise.

Let q be an odd prime power and $E : Y^2 = X^3 + aX + b$ an elliptic curve over k . If $d \in k^*$ (the set of non-zero elements of k) is a non-square, then the curve $E^{(d)} : Y^2 = X^3 + d^2aX + d^3b$ is a non-trivial twist of E over k . Furthermore, as long as $j(E) \neq 0$ or 1728, all other non-trivial twists of E are isomorphic to $E^{(d)}$. $E^{(d)}$ is called a *quadratic twist* of E .

A similar notion exists when the characteristic of k is 2. If γ is an element with $\text{Tr}(\gamma) = 1$, then the ordinary curves $E : Y^2 + XY = X^3 + b$ and $E' : Y^2 + XY = X^3 + \gamma X^2 + b$ are non-trivial twists over k . By Proposition 2.1.10, any twist of E must be isomorphic over k to one of these two curves. The relationship between the cardinalities of twist curves is given by the following theorem.

Theorem 2.1.12. [10, Section 9.5] *Let q be a prime power and E/\mathbb{F}_q an elliptic curve with $\#E(\mathbb{F}_q) = q + 1 - t$. Then a quadratic twist over \mathbb{F}_q has cardinality $q + 1 + t$.*

This theorem will be particularly useful in our experimental results since if we know the number of distinct isomorphism classes with a given cardinality $q + 1 - t$, we immediately know the number of isomorphism classes of cardinality $q + 1 + t$.

2.2 Algebraic Number Theory

It seems necessary to now turn our discussion towards some of the essential concepts of algebraic number theory that will help in our understanding of isogenies. A more complete treatment can be found in any algebraic number theory textbook; we found [37] to be quite helpful as well.

2.2.1 Quadratic Number Fields

Let $L = \mathbb{Q}(\sqrt{d})$ for some square-free integer $d \neq 0, 1$. When $d < 0$, L is called an *imaginary quadratic number field* and its *degree* is the degree of L as an extension of \mathbb{Q} (which will be 2). An element $\alpha \in \mathbb{C}$ is said to be *algebraic* if it is the root of some polynomial $A(x) \in \mathbb{Z}[x]$. If it is possible to find a monic such $A(x)$, then we call α an *algebraic integer*. The set \mathbb{A} of all algebraic integers is a ring, and one defines $\mathcal{O}_L = \mathbb{A} \cap L$ to be the *ring of integers* of L .

\mathcal{O}_L is a special case of the more general notion of an *order*. An order of L is defined to be a subring of L that is a finitely generated \mathbb{Z} -submodule of rank $n = \deg(L)$ (recall that intuitively, for a ring A , an A -module is an Abelian group M on which A acts linearly). In a quadratic number field, one can think of an order as a set

$$\mathcal{O} = \{a\omega_1 + b\omega_2 \mid a, b \in \mathbb{Z}\},$$

where ω_1, ω_2 are \mathbb{Q} -linearly independent elements of L such that the set forms a ring. It is possible to show that every order must be contained in the ring of integers of L , and so \mathcal{O}_L is actually the unique maximal order of L . One defines the *discriminant* of L to be

$$D = \text{Disc}(L) = \begin{cases} d & \text{if } d \equiv 1 \pmod{4}, \\ 4d & \text{if } d \not\equiv 1 \pmod{4}. \end{cases} \quad (2.4)$$

It can be shown that $\mathcal{O}_L = \mathbb{Z}[(D + \sqrt{D})/2]$ and that any order $\mathcal{O} = \mathbb{Z}[c(D + \sqrt{D})/2]$ for some $c \in \mathbb{N}$. The value c is known as the *conductor* of \mathcal{O} and is equal to the index $[\mathcal{O}_L : \mathcal{O}]$.

2.2.2 Ideals and Factorization

While we are assuming an elementary background in ring theory, there are a few additional definitions and concepts that we will require pertaining to ideals in imaginary quadratic number fields. In particular, we wish to define the ideal class group and mention a result regarding the factorization of ideals.

Let \mathcal{O} denote the ring of integers of an imaginary quadratic field L , and recall that L is the fraction field of \mathcal{O} . A *fractional ideal* I of \mathcal{O} is an \mathcal{O} -submodule of L such that there exists some $\alpha \in L$ with $\alpha I \subseteq \mathcal{O}$. In this section we will only be referring to fractional ideals of \mathcal{O} , so from now on we will just say fractional ideal. As a quick example, suppose $\mathcal{O} = \mathbb{Z}[\sqrt{-2}] \subseteq \mathbb{Q}(\sqrt{-2})$. Then $\frac{1}{2}\mathcal{O}$ is a fractional \mathcal{O} -ideal in $\mathbb{Q}(\sqrt{-2})$, but not an ideal of \mathcal{O} .

A fractional ideal I is said to be *invertible* if there exists some fractional ideal J that

$$IJ = \{ab : a \in I, b \in J\} = \mathcal{O}.$$

The set of fractional ideals form a group under ideal multiplication, and we define the *ideal class group* of \mathcal{O} to be the set of nonzero fractional ideals modulo the principal fractional ideals (those generated by a single element). The *class number* of \mathcal{O} is the cardinality of the class group.

Now recall that the factorization of \mathcal{O} -ideals into prime ideals is unique. One defines the *norm* of an \mathcal{O} -ideal I to be

$$N(I) = \#(\mathcal{O}/I),$$

and we have the useful fact that the ideal generated by a rational prime l , denoted $l\mathcal{O}$, can factor in three different ways:

1. $l\mathcal{O} = \mathfrak{q}_1\mathfrak{q}_2$, with $\mathfrak{q}_1 \neq \mathfrak{q}_2$ and $N(\mathfrak{q}_1) = N(\mathfrak{q}_2) = l$;
2. $l\mathcal{O} = \mathfrak{q}^2$ and $N(\mathfrak{q}) = l$; or
3. $l\mathcal{O}$ is itself prime.

Note that in each case $N(l\mathcal{O}) = l^2$. In order, we refer to these cases by saying l *splits*, *ramifies* or is *inert*. It turns out that an ideal of norm l exists only in the first two cases, and that there are no other such ideals beside the factors of $l\mathcal{O}$.

An important theorem by Dedekind (see Proposition 5.11 [2]) tells us that if $L = \mathbb{Q}(\alpha)$ and l is a prime that does not divide $[\mathcal{O}_L : \mathbb{Z}[\alpha]]$, then one may determine the behaviour of $l\mathcal{O}$ by looking at the factorization of the minimal polynomial f of α modulo l . In particular, cases 1, 2, 3 above correspond to a factorization of f into 2, 1 or 0 (irreducible) distinct degree one factors. Moreover, if μ and λ are the roots of f modulo l , then the norm- l ideals are of the form $(l, \alpha - \mu)$ and $(l, \alpha - \lambda)$.

For a quadratic field, we can determine how a prime l (that does not divide $[\mathcal{O}_L : \mathbb{Z}[\alpha]]$) behaves by computing the Legendre symbol $\left(\frac{\text{Disc}(L)}{l}\right)$. This can be seen from the quadratic formula mod l since the number of roots corresponds to whether the discriminant is 0, a quadratic residue or a quadratic non-residue mod l (these correspond to the ramify, split and inert cases respectively). We will see how the behaviour of primes provides us with important information about the isogenies of an elliptic curve.

Chapter 3

Isogenies

3.1 Preliminaries

Now that we have sufficient background knowledge, we introduce the key ingredient used in extending the GHS attack.

Definition 3.1.1. *Let E_1/k and E_2/k be elliptic curves. An **isogeny** over k is a morphism $\varphi : E_1 \rightarrow E_2$ over k such that $\varphi(\mathcal{O}_{E_1}) = \mathcal{O}_{E_2}$.*

Note that Theorems 2.1.7 and 2.1.8 tell us that isogenies are group homomorphisms and that every isogeny is surjective except for the one that takes all points to \mathcal{O}_{E_2} . Given an isogeny φ as above, its **kernel** is the set $\ker(\varphi) = \{P \in E_1(\overline{K}) : \varphi(P) = \mathcal{O}_{E_2}\}$.

We denote the set of all isogenies defined over k from E_1 to E_2 by $\text{Hom}_k(E_1, E_2)$. When $E_2 = E_1$, we obtain the **endomorphism ring** of E , which will be denoted by $\text{End}_k(E)$, and consists of all isogenies from E to itself defined over k . We will drop the subscript k in the case where we are considering the algebraic closure of the field of definition of E . One may verify that $\text{End}_k(E)$ is a ring with addition and multiplication given by point-wise addition and composition respectively, and the structure of this ring will be of interest to us in later sections.

To motivate our discussion of isogenies, we state Tate's isogeny theorem, which illustrates an important property of isogenies in relation to cryptography. Essentially instances of the ECDLP on one curve can be mapped to instances of an isogenous one, given that the group in which the ECDLP sits is not annihilated by the kernel of the isogeny. This caveat is typically not a concern in practice since the ECDLP is usually chosen to be in a

subgroup of large prime order and the known methods for efficiently constructing isogenies do not construct ones with kernels that have large prime subgroups.

Theorem 3.1.1 (Tate’s Isogeny Theorem). *Let E_1 and E_2 be elliptic curves over a finite field k . E_1 is isogenous over k to E_2 if and only if $\#E_1(k) = \#E_2(k)$.*

It should be noted that Tate’s isogeny theorem is technically only the reverse implication of the above, and that originally it appeared in the more general context of abelian varieties; we have simply given its specialization to the elliptic curve setting.

3.2 Standard Form

Isogenies have two particularly important invariants which we will require; those of **degree** and **separability**. These concepts have definitions rooted in algebraic geometry, which the interested reader may refer to [9] or [10] for details. It is not our intention to dive into those aspects here, so we will use a more concrete definition. Our first order of business will be to develop a standard form for a given isogeny as in Section 2.9 of [40]. For notational ease, we will assume that we are working with a curve E given in the form (2.2).

Let $r(X, Y)$ be any rational function with coefficients in k . Since we are interested in points $(X, Y) \in E(\bar{k})$, we can repeatedly use the curve equation to substitute even powers of Y for a polynomial in X . It is thus possible to express r in the form

$$r(X, Y) = \frac{p_1(X) + p_2(X)Y}{p_3(X) + p_4(X)Y}.$$

Further, by multiplying numerator and denominator by $p_3(X) - p_4(X)Y$ and doing a few more substitutions, we can obtain

$$r(X, Y) = \frac{q_1(X) + q_2(X)Y}{q_3(X)}. \tag{3.1}$$

Now suppose $\varphi(P) = \varphi((X, Y)) = (f(X, Y), g(X, Y))$ is an isogeny. Using the fact that φ is a group homomorphism and that the inverse of a point $P = (X, Y)$ is $-P = (X, -Y)$, it follows that

$$\varphi((X, -Y)) = \varphi(-(X, Y)) = -\varphi(X, Y).$$

Hence if we write f in the form (3.1) we can deduce that $q_2(X) = 0$. Writing g in the same way, similar reasoning yields that the corresponding q_1 is the zero polynomial. After eliminating common factors, we finally arrive at

$$\varphi(X, Y) = (f(X), g(X)Y),$$

where f and g are now rational functions only in the variable X . This is the **standard form** of an isogeny.

Definition 3.2.1. Consider an isogeny φ in standard form. Write

$$f(X) = \frac{r_1(X)}{r_2(X)},$$

where the r_i are coprime polynomial functions in X (with coefficients in k). We define the **degree** of φ as

$$\deg(\varphi) = \max\{\deg(r_1), \deg(r_2)\}.$$

Definition 3.2.2. An isogeny is called **separable** if the derivative of r_1/r_2 is non-zero. Otherwise we call the isogeny **inseparable**.

Our primary focus will be on separable isogenies, so we will often drop the word ‘separable’ unless the distinction needs to be made. The following well-known result ties the notions of degree and separability together.

Lemma 3.2.1. [10, Lemma 9.6.4] A nonzero separable isogeny $\varphi : E_1 \rightarrow E_2$ over k of degree d has $\#\ker(\varphi) = d$.

Some examples of isogenies include the *multiplication by n maps*, which send $P \mapsto nP$ for all points $P \in E(\overline{K})$ and the **Frobenius endomorphism**, denoted by π , which given $k = \mathbb{F}_q$ takes

$$(x, y) \mapsto (x^q, y^q).$$

For n relatively prime to the characteristic, $[n]$ is separable and has degree n^2 , whereas π is inseparable of degree q .

Example 1. Consider the elliptic curve E given by the equation $E : Y^2 = X^3 + 12X + 1$ over \mathbb{F}_{13} . If we write $\mathbb{F}_{13^3} \cong \mathbb{F}_{13}[t]/(t^3 + 2t + 11)$, then E has three 2-torsion points, $(2t^2 + 6t + 7, 0)$, $(t^2 + 1t + 10, 0)$, and $(10t^2 + 6t + 10, 0)$. There are thus three separable 2-isogenies over \mathbb{F}_{13^3} whose kernels are \mathcal{O}_E and one of these points. Explicitly the isogenies are (respectively) as follows:

(i) The map which takes (X, Y) to

$$\left(\frac{X^2 + (2t^2 + 6t + 7)X + (4t^2 + 5t + 5)}{X + (11t^2 + 7t + 6)}, \frac{X^2Y + (9t^2 + t + 12)XY + (5t^2 + 3t + 2)Y}{X^2 + (9t^2 + t + 12)X + (4t^2 + 5t + 6)} \right)$$

is a degree-2 isogeny from E to the curve

$$E_1/\mathbb{F}_{13^3} : Y^2 = X^3 + (t^2 + 3t + 10)X^2 + (t^2 + 11t + 12)X,$$

(ii) The map which takes (X, Y) to

$$\left(\frac{X^2 + (t^2 + t + 10)X + (6t^2 + 5t + 12)}{X + (12t^2 + 12t + 3)}, \frac{X^2Y + (11t^2 + 11t + 6)XY + (t^2 + 3t + 1)Y}{X^2 + (11t^2 + 11t + 6)X + (6t^2 + 5t)} \right)$$

is a degree-2 isogeny from E to the curve

$$E_1/\mathbb{F}_{13^3} : Y^2 = X^3 + (7t^2 + 7t + 5)X^2 + (8t^2 + 11t + 4)X,$$

(iii) The map which takes (X, Y) to

$$\left(\frac{X^2 + (10t^2 + 6t + 9)X + (3t^2 + 3t + 8)}{X + (3t^2 + 7t + 4)}, \frac{X^2Y + (6t^2 + 1t + 8)XY + (7t^2 + 7t + 9)Y}{X^2 + (6t^2 + t + 8)X + (3t^2 + 3t + 9)} \right)$$

is a degree-2 isogeny from E to the curve

$$E_1/\mathbb{F}_{13^3} : Y^2 = X^3 + (5t^2 + 3t + 11)X^2 + (4t^2 + 4t + 3)X.$$

Note that the image curves are not in Weierstrass form, but these may be easily converted.

One might wonder (and hope) that if E_1 is isogenous to E_2 whether E_2 is isogenous to E_1 . This question is answered in the affirmative by the existence of the *dual isogeny*.

Theorem 3.2.2. [10, Theorem 9.6.21] *Let E_1 and E_2 be two elliptic curves over k , and $\varphi : E_1 \rightarrow E_2$ be a nonzero isogeny over k of degree m . Then there is a nonzero isogeny $\hat{\varphi} : E_2 \rightarrow E_1$ over k such that*

$$\hat{\varphi} \circ \varphi = [m] : E_1 \rightarrow E_1.$$

Furthermore, $\hat{\varphi}$ is unique and is called the *dual isogeny* of φ .

3.3 Kernels of Isogenies

Now that we have an idea of what isogenies are, a natural question is what sort of isogenies can arise from a particular elliptic curve. In this section we present a series of results that will shed some light on the structure of such maps.

Theorem 3.3.1. [10, Theorem 9.6.18] *Let E_1, E_2, E_3 be elliptic curves over k and $\varphi : E_1 \rightarrow E_2, \psi : E_1 \rightarrow E_3$ be isogenies over k . Suppose that $\ker(\varphi) \subseteq \ker(\psi)$ and that ψ is separable. Then there is a unique isogeny $\lambda : E_2 \rightarrow E_3$ over k such that $\psi = \lambda \circ \varphi$.*

What this theorem says is that by taking a nontrivial subgroup of the kernel of some isogeny, one can decompose it into two isogenies of smaller degree.

Definition 3.3.1. *A set G is said to be **defined** over k if $\sigma(P) \in G$ for all $P \in G$ and $\sigma \in \text{Gal}(\bar{k}/k)$; i.e. the Galois group of \bar{k}/k fixes G .*

Theorem 3.3.2. *[10, Theorem 9.6.19] Let E/k be an elliptic curve and $G \subseteq E(\bar{k})$ a finite group defined over k . Then there is a unique (up to isomorphism over \bar{k}) elliptic curve E_1/k and a separable isogeny $\varphi : E \rightarrow E_1$ over k such that $\ker(\varphi) = G$.*

It is also true that if ψ and φ are two isogenies mapping from E to E_1 with the same kernel, then $\psi = \lambda \circ \varphi$, where λ is an automorphism on the image curve. Further, if ψ instead maps to another curve E_2 , then $\psi = \lambda \circ \varphi$ where $\lambda : E_1 \rightarrow E_2$ is an isomorphism over k .

The main conclusion here is that an isogeny is essentially determined by its kernel. We will say that two separable isogenies are **equivalent** if their kernels are the same. A corollary of the Theorem 3.3.2 is that every isogeny may be decomposed into a chain of prime degree isogenies. We state this formally below.

Theorem 3.3.3. *[10, Theorem 25.1.2] Let E_1 and E_2 be elliptic curves over k and $\varphi : E_1 \rightarrow E_2$ a separable isogeny over k . Then we can write*

$$\varphi = \varphi_1 \circ \varphi_2 \circ \cdots \circ \varphi_k \circ [n],$$

where each φ_i is a separable isogeny over k of prime degree and n is the largest integer such that $E[n] \subseteq \ker(\varphi)$. Furthermore,

$$\deg \varphi = n^2 \prod_{i=1}^k \deg(\varphi_i).$$

This result plays a critical role in the concept of an *isogeny walk* since it tells us that any isogeny can be constructed through a series of “steps” using prime degree isogenies.

3.4 Endomorphism Rings

Let k be a field and recall that $\text{End}(E)$ is the set of all isogenies over \bar{k} from E to itself. We begin by listing a few facts regarding the structure of this ring.

Theorem 3.4.1. [40, Theorem 10.6] Let E/k be an elliptic curve, $\text{char}(k) = p$.

1. If E is ordinary, then $\text{End}(E)$ is an order in an imaginary quadratic field L .
2. If E is supersingular, then $\text{End}(E)$ is a maximal order in a definite quaternion algebra that is ramified at p and ∞ and is split at the other primes.

It can be shown that if E is ordinary and $k = \mathbb{F}_q$, then $L = \mathbb{Q}(\sqrt{t^2 - 4q})$, where t is the trace of Frobenius. To help gain a rough intuition, recall that for any $n \in \mathbb{Z}$ the multiplication-by- n map $[n]$ is an endomorphism. Each such endomorphism is distinct and non-trivial, hence by identifying $[n]$ with its corresponding integer, one can conclude that $\mathbb{Z} \subseteq \text{End}(E)$. With finite fields, the Frobenius endomorphism π also lies in $\text{End}(E)$, thus we always have the following containments:

$$\mathbb{Z}[\pi] \subseteq \text{End}(E) \subseteq \mathcal{O}_L,$$

where \mathcal{O}_L denotes the ring of integers of L . Because of this containment, we will often abuse notation by simply referring to endomorphisms as complex numbers.

In the ordinary case, $\pi \neq [n]$ for any $n \in \mathbb{Z}$, and may be identified with a complex number having minimal polynomial $x^2 - tx + q$. Since L is a quadratic number field, $\{1, \pi\}$ is a \mathbb{Q} -basis for L .

One could write

$$\pi = \frac{-t \pm \sqrt{t^2 - 4q}}{2},$$

but everything not in the root will be absorbed by \mathbb{Q} , yielding

$$L = \mathbb{Q}(\pi) = \mathbb{Q}(\sqrt{t^2 - 4q}).$$

An important note is that if c_π is the conductor of $\mathbb{Z}[\pi]$ in \mathcal{O}_L , the conductor c of any endomorphism ring must divide c_π . Moreover one can show that c_π is the largest integer such that $(t^2 - 4q)/c_\pi^2 \equiv 0, 1 \pmod{4}$.

If we denote $\text{End}(E)$ by \mathcal{O} , whenever E has an endomorphism ring that strictly contains a copy of \mathbb{Z} , we say that E has *complex multiplication* or CM (by \mathcal{O}). It follows that every ordinary elliptic curve defined over a finite field has CM. The theory of complex multiplication is very powerful, and we will make use of a number of useful facts derived from it [37].

1. Given an ordinary elliptic curve E/k with $\text{End}(E) = \mathcal{O}$ an order in an imaginary quadratic field L , suppose that there exists an isogeny $\varphi : E \rightarrow E'$ of prime degree l . Then $\text{End}(E') = \mathcal{O}'$ is also an order in L , and one of the following holds:

- (a) $\mathcal{O} = \mathcal{O}'$,
- (b) $[\mathcal{O} : \mathcal{O}'] = l$,
- (c) $[\mathcal{O}' : \mathcal{O}] = l$.

In the three cases above, the isogeny is referred to as *horizontal*, *descending*, and *ascending* respectively. Horizontal isogenies have distinctive characteristics from the other two cases, and so to distinguish them the latter two cases are called *vertical* isogenies.

2. Every horizontal l -isogeny arises from the action of an invertible \mathcal{O} -ideal of norm l . Hence from our discussion in Section 2.2.2, if one is dealing with a curve whose endomorphism ring is the maximal order \mathcal{O}_L , then one may look at the Legendre symbol $\left(\frac{\text{Disc}(L)}{l}\right)$ to determine exactly how many distinct horizontal l -isogenies arise from a particular curve.

3.5 Isogeny Graphs

To reach our goal of understanding why the isogeny walk strategy is effective, we will need to understand how curves defined over a field k are related through different degree isogenies. It seems natural to consider the graph whose vertices correspond to elliptic curves and whose edges correspond to isogenies. Recall that Theorem 3.3.3 tells us that any separable isogeny may be constructed out of a series of prime degree isogenies. Hence it will be helpful to first look at the structure of such graphs where edges only exist if there is an isogeny of degree l for a fixed prime l . We will see how such graphs combine together for curves of a fixed order N .

To begin, we can consider the vertex set to be the elements of k thanks to the following theorem.

Theorem 3.5.1. [37, Theorem 14.12] *For every j_0 in k , there exists an elliptic curve E/k such that $j(E) = j_0$.*

If we want our edges to correspond to l -isogenies, it would be nice if we had a way to easily determine whether two curves were l -isogenous over k . Fortunately, the l -th modular polynomial allows us to do just that.

For each prime l there exists a degree- $(l + 1)$ polynomial $\Phi_l(X, Y) \in \mathbb{Z}[X, Y]$ with the property that if one computes $\Phi_l(j, Y)$ where j is the j -invariant of a curve E/k , then the

roots of $\Phi_l(j, Y)$ in k are exactly the j -invariants of the curves which are l -isogenous to E over k . It is known that if $l \neq \text{char}(k)$, then the number of l -isogenies over k up to isomorphism is either 0, 1, 2 or $l + 1$. This corresponds to the cases where there are exactly 0, 1, or 2 order- l subgroups of $E(\overline{k})$ that are defined over k , or if $E[l]$ is defined over k . We refer to $\Phi_l(X, Y)$ as the **l -th modular polynomial**, which is used to formally define the l -isogeny graph:

Definition 3.5.1. *The l -isogeny graph $G_l(k)$ is the graph with vertex set k and directed edges (j_1, j_2) occurring with multiplicity equal to the multiplicity of j_2 as a root of $\Phi_l(j_1, Y)$.*

It is worth mentioning that given two isogenous elliptic curves related by a degree- l isogeny, their twists are also related by a degree- l isogeny. Hence there is no loss of generality if we simply treat twist curves as equivalent. If one is actually using the modular polynomial to map to a different curve, the choice of twist is important, but not a difficult task in practice.

In order to represent non-prime isogenies in our graph, we can essentially stack a number of l -isogeny graphs on top of each other. Of course, we could do this for infinitely many primes, but beyond a certain number computations would become quite impractical. Our hope would be that a reasonable selection of primes will be enough to connect all the vertices in an isogeny class.

Note that for each l , $G_l(k)$ is a priori known to be disconnected as a consequence of Tate's isogeny theorem. It thus makes sense to restrict our attention to the j -invariants of curves having the same order N . This leads us to the following definition:

Definition 3.5.2. *[10, 25.3.1] Let E/k be an elliptic curve with $\text{char}(k) = p$. Let $S \subseteq \mathbb{N}$ be a finite set of prime numbers. We define*

$$X_{E,k,S}$$

to be the (directed) graph with vertex set consisting of elements from the k -isogeny class of E . One can think of the vertices as being labeled by j -invariants of isogenous curves. For each $l \in S$, there is an edge $(j(E_1), j(E_2))$ labeled by l for each equivalence class of l -isogenies from E_1 to E_2 defined over k . Note that because every isogeny has a dual isogeny of the same degree going in the opposite direction, we can just think of this as an undirected graph.

Things are slightly more complicated when $j(E) = 0$ or 1728 in that we cannot assume the graph is undirected. However these instances are rare enough that we can ignore them in the interest of simplification. We would also like to point out that if we have a particular order N in mind rather than a curve, we can write $X_{N,k,S}$ without any ambiguity.

3.5.1 The Structure of Isogeny Graphs

Recall that supersingular curves are only isogenous to other supersingular curves. The structure of subgraphs consisting of supersingular curves is distinct from those of ordinary curves, but as usual we will only be discussing the ordinary case.

In 1996, Kohel published the seminal work on the structure of isogeny graphs of elliptic curves over finite fields in [20] and we refer to [36] for an excellent exposition on the subject. The main structural result is that the ordinary components of $G_l(k)$ are a type of graph known as *l-volcanoes*. We will merely state the definition and Kohel's main result.

Definition 3.5.3. *An l -volcano V is a connected, undirected graph whose vertices are partitioned into one or more levels V_0, \dots, V_d such that the following hold:*

- (i) *The subgraph induced by V_0 (called the surface) is a regular graph of degree at most 2.*
- (ii) *For $i > 0$, each vertex in V_i has exactly one neighbour in level V_{i-1} and this accounts for every edge not on the surface.*
- (iii) *For $i < d$, each vertex in V_i has degree $l + 1$.*

The value d is referred to as the depth of the volcano.

Theorem 3.5.2 (Kohel). *Let V be an ordinary component of the l -isogeny graph $G_l(k)$ that does not contain the values 0 or 1728. Then V is an l -volcano with the following properties:*

- (i) *The vertices in level V_i all have the same endomorphism ring \mathcal{O}_i .*
- (ii) *The subgraph on V_0 has degree $1 + (\frac{D_0}{l})$, where $D_0 = \text{disc}(\mathcal{O}_0)$.*
- (iii) *If $(\frac{D_0}{l}) \geq 0$, then $|V_0|$ is the order of \mathfrak{l} in the class group of \mathcal{O}_0 , where \mathfrak{l} is a norm- l ideal; otherwise $|V_0| = 1$.*
- (iv) *The depth of V is equal to $r/2$, where r is the largest integer such that $l^r \mid ((t^2 - 4q)/D_0)$. Here t is the trace of Frobenius for a curve whose j -invariant is in V .*
- (v) *$l \nmid [\mathcal{O}_L : \mathcal{O}_0]$ and $[\mathcal{O}_i : \mathcal{O}_{i+1}] = l$ for $0 \leq i < d$.*

We would like to emphasize a few specific notions to take from this result. Suppose E/k is an elliptic curve with endomorphism ring $\mathcal{O} \neq \mathcal{O}_L$ and conductor c . If φ is an isogeny (over k) between E and a curve E'/k with $\text{End}(E') = \mathcal{O}_L$, then c must divide $\deg(\varphi)$. Secondly, if one is considering an elliptic curve with endomorphism ring equal to \mathcal{O}_L , and l does not divide $[\mathcal{O}_L : \mathbb{Z}[\pi]]$, then the number of l -isogenies originating from this curve is completely determined by the value of $(\frac{D_0}{l})$ (since there are no descending isogenies).

In 2005, Jao et al. [18] published results affirming the random reducibility of discrete logarithms amongst curves of the same order under the assumption of the Generalized Riemann Hypothesis (GRH). One of the primary theorems in this work is the following:

Theorem 3.5.3. [18, Theorem 1.1] *Consider the isogeny graph $X_{N,k,S}$ where S is the set of all primes less than some number m . Assuming the GRH, there exists a polynomial $p(x)$, independent of N and q such that for $m = p(\log q)$ the isogeny graph $X_{N,k,S}$ on each level is an expander graph, in the sense that any random walk on $X_{N,k,S}$ will reach a subset of size h with probability at least $\frac{h}{2^{|X_{N,k,S}|}}$ after $\text{polylog}(q)$ steps (where the implicit polynomial is again independent of N and q).*

We will use this property extensively in our analyses of the random walk in the coming chapters.

Chapter 4

Cryptographic Background

4.1 The Elliptic Curve Discrete Log Problem

Let E/k be an elliptic curve and consider the group $E(k)$. Suppose that $P \in E(k)$ has order N and $Q \in \langle P \rangle$, i.e. $Q = \alpha P$ for some $\alpha \in [0, N - 1]$. The number α is called the *discrete logarithm* of Q with respect to P , and denoted by $\log_P Q$. The *elliptic curve discrete logarithm problem* (ECDLP) is the problem of finding α given only E, P and Q .

In general, the ECDLP is thought to be a hard problem. There are of course instances in which the properties of a particular instance of ECDLP may be exploited to solve it faster than expected, but in cryptographically interesting cases one would assume such techniques are not applicable. Some attacks to consider during selection of a curve/instance are the Pohlig-Hellman attack [29], Weil and Tate pairing attacks [22, 8], and those that affect prime-field anomalous curves [31, 33, 35].

Given that one has chosen an instance appropriately, the problem is believed to be about as hard as in a general group of size N , which means that the best methods of attack are so-called “generic DLP solvers”— methods which may be applied to any curve/ECDLP instance as they do not make use of any special properties. The expected run time of such algorithms is approximately \sqrt{N} . We will begin our discussion with one of the most effective and well-known algorithms, Pollard’s ρ method.

4.1.1 Pollard’s Rho

In a general group G of order N , the method is as follows [40, Chapter 5.2]. Begin by selecting a function $f : G \rightarrow G$ that “looks” random, and initialize with a random element

P_0 . The next step is to compute iterations $P_{i+1} = f(P_i)$. Since G is finite, at some point there will be indices i_0 and j_0 , $i_0 < j_0$, such that $P_{i_0} = P_{j_0}$. A consequence of f being deterministic is that now we have

$$P_{i_0+l} = P_{j_0+l}$$

for all $l \geq 0$. Assuming that the difference $j_0 - i_0$ is the minimal number such that the above collision occurs, the sequence P_i will become periodic after some value i with period $j_0 - i_0$.

We will see that once we have obtained i_0 and j_0 , the DLP is solved. The run time thus boils down to finding a collision in the function f , which assuming f is truly random, is on the order of \sqrt{N} by the birthday paradox. It has been shown that this is essentially optimal (one can refer to [37, Chapter 10] for a nice exposition). Pollard's ρ has extra appeal since it may be modified in such a way as to require very little storage (only two points at any one time) without much more computation. Moreover, it can be parallelized to achieve a speedup that is linear in the number of processors [39].

The typical method that one extracts information about the discrete logarithm from a collision is to choose f in such a way that for each i ,

$$P_i = a_i P + b_i Q,$$

where a_i, b_i are known. Hence a collision will result in a relation

$$a_i P + b_i Q = a_j P + b_j Q.$$

Rearranging will yield the congruence

$$\alpha \equiv \frac{b_i - b_j}{a_j - a_i} \pmod{N/d},$$

where $d = \gcd(N, a_j - a_i)$. One can then guess at values for α until the appropriate value is found, assuming d is small. In cryptographic applications, N is often a prime number so $d = 1$ or N . In the latter case we have a trivial relationship and we simply repeat until $d = 1$.

It is an important observation that the asymptotic running time of Pollard's rho algorithm, $\mathcal{O}(\sqrt{N})$ does not seem to hide any large constants, and is in fact very close to the actual running time [38].

4.2 Index Calculus

An alternate method for solving DLOG that has seen quite a bit of success is *index calculus*. Returning to the setting of a general cyclic group G of size N , suppose g is a generator and that we are trying to find $\log_g h$ for some h in G .

The index calculus method begins by selecting a set $\mathcal{F} \subseteq G$ of elements known as a *factor base*. The elements of \mathcal{F} should be such that a large portion of elements of G may be efficiently expressed as a product of those in \mathcal{F} . The reason for this is that one desires to collect relations of the form

$$g^k = a_1^{b_1} a_2^{b_2} \cdots a_r^{b_r},$$

where $a_1, \dots, a_r \in \mathcal{F}$, $b_1, \dots, b_r \in \mathbb{Z}_{>0}$, and k is a randomly chosen integer. From such a relation, one can obtain a linear equation in the logarithms of the a_i :

$$k \equiv b_1 \log_g a_1 + \cdots + b_r \log_g a_r \pmod{N}.$$

If $\#\mathcal{F} = t$, then it is quite likely that after a little more than t such relations are found, we will have a system of linear equations with rank t , which will allow us to use standard Gaussian elimination to determine $\log_g a_i$ for all i .

Once these values have been obtained, the next step is to generate relations of the form

$$hg^k = a_1^{b_1} a_2^{b_2} \cdots a_r^{b_r},$$

where the notation is the same as above. From here it is straightforward to solve for $\log_g h$. The running times of index calculus algorithms depend heavily on the choice of factor base elements, and often there is no clear choice. One must also be wary of the number of elements chosen to be in \mathcal{F} as there is a trade-off— more elements will make it easier to find a relation, however this also means that more relations will be needed in order to obtain a full-rank matrix.

4.3 Point Decomposition

In the previous section, we saw how the index calculus method could be applied to a general group to solve discrete logarithms. However if one wishes to apply this method to elliptic curves (or hyperelliptic curves) it is somewhat unclear how to do so effectively. In particular, there is no obvious choice of factor base elements and the concept of “factorization” in an elliptic curve group has no real meaning.

A significant breakthrough to these issues came from Semaev [34] in 2004. His idea was that instead of factorization, one could aim to decompose points into the sum of a fixed number of points in a factor base. To this end, he introduced *summation polynomials*, which gave an algebraic expression for the aforementioned problem. We present an outline of this seminal work below.

4.3.1 Summation Polynomials

In order to simplify the notation, suppose that E is an elliptic curve over a field \mathbb{F}_q of characteristic greater than 3 (note that summation polynomials apply in characteristic 2 as well, however this restriction allows us to simplify the notation). We can then express E using the equation $Y^2 = X^3 + aX + b$. The summation polynomials f_n are defined recursively as follows:

- $f_2(X_1, X_2) = X_1 + X_2$,
- $f_3(X_1, X_2, X_3) = (X_1 - X_2)^2 X_3^2 - 2((X_1 + X_2)(X_1 X_2 + a) + 2b)X_3 + ((X_1 - a)^2 - 4b(X_1 + X_2))$,
- and for $n \geq 4$ and $1 \leq r \leq n - 3$,

$$f_n(X_1, \dots, X_n) = \text{Res}_X(f_{n-r}(X_1, \dots, X_{n-r-1}, X), f_{r+2}(X_{n-r}, \dots, X_n, X)).$$

In essence, a solution to the n th summation polynomial satisfying specific conditions is equivalent to a decomposition into the sum of n factor base elements. This is formalized by the following theorem.

Theorem 4.3.1 (Semaev). *Let E be an elliptic curve defined over k , $n \geq 2$ an integer and f_n its n -th summation polynomial. Let X_1, X_2, \dots, X_n be n elements of the algebraic closure \bar{k} of k . Then $f_n(X_1, \dots, X_n) = 0$ if and only if there exists an n -tuple (Y_1, \dots, Y_n) in \bar{k} such that for all i , $P_i = (X_i, Y_i)$ is a point of E and*

$$P_1 + \dots + P_n = \mathcal{O}_E.$$

Furthermore, if $n \geq 3$, the polynomial f_n is symmetric of degree 2^{n-2} in each variable.

It is easy to see that in order to decompose some point $R = (X', Y')$ into N points, it will suffice to find a solution $(X_1, \dots, X_N) \in \mathbb{F}_q^N$ to the equation

$$f_{N+1}(X_1, \dots, X_N, X') = 0.$$

One will have to do some guesswork regarding the sign of Y_i for each i since the X_i only determine the value of Y_i up to sign, but this takes as most 2^N trials for a fixed (and usually small) N .

Building on Semaev's work, Diem [4] and Gaudry [12] independently proposed methodologies for attacking curves defined over extension fields $\mathbb{F}_{q^n}/\mathbb{F}_q$. The underlying idea of their attacks were similar in spirit to those of index calculus, but are commonly referred to as *decomposition attacks* in order to distinguish them. The main benefit of working in an extension field is that a relatively natural choice for the factor base arises; those points whose X -coordinate lies in the base field \mathbb{F}_q . We provide a rough outline of the strategy used in [12] to illustrate the elegance of their approach.

If one fixes a polynomial representation of \mathbb{F}_{q^n} as an extension of \mathbb{F}_q , say $\mathbb{F}_q[t]/(f(t))$ where $f(t)$ is an irreducible monic polynomial of degree n , then any element of \mathbb{F}_{q^n} may be stored as an n -tuple with entries in \mathbb{F}_q via the bijection

$$(X_0, \dots, X_{n-1}) \longleftrightarrow (X_0 + \dots + X_{n-1}t^{n-1}).$$

The *Weil restriction of scalars* for E/\mathbb{F}_{q^n} is the set of $2n$ -tuples

$$(X_0, \dots, X_{n-1}, Y_0, \dots, Y_{n-1}) \in \mathbb{F}_q^{2n}$$

that correspond to points of E under the natural extension of the above map. It is worth noting that the Weil restriction retains the group law of E . This representation is convenient since we can specify the elements whose X -coordinates lie in the base field by simply setting $X_1 = \dots = X_{n-1} = 0$.

If we invoke summation polynomials to solve a point decomposition now, we must rewrite the coefficients of $f_{N+1}(X_1, \dots, X_N, X') = 0$ as elements of $\mathbb{F}_q[t]/(f(t))$ and equate coefficients of powers of t . This will result in a system of n polynomial equations (one for each power of t) in N unknowns; an analysis shows that taking $N = n$ is the best choice. Since the summation polynomials are symmetric, it follows that these n equations are symmetric as well. The authors of [12] propose that it is beneficial to rewrite these polynomials in terms of a symmetric basis. This system may then be solved using a Gröbner basis algorithm, and then we can find rational roots of the subsequent polynomials to obtain the X -coordinates of the points in the decomposition. It turns out that the Gröbner basis step is by far and large the limiting step, with an expected running time that is polynomial in the value $2^{n(n-1)}$.

One can apply a method of index calculus known as the *double large prime variation* [14] to bring the complexity to an estimated $O(q^{2-2/n})$, where n is fixed and q tends to infinity. Unfortunately the hidden constants grow very large as n is increased, and the

method becomes impractical for $n \geq 4$ or 5. However when n is say 3, the running time is $\mathcal{O}(q^{4/3})$, which is asymptotically faster than the Pollard's rho algorithm's expected $\mathcal{O}(q^{3/2})$ running time.

4.3.2 Decomposition Attacks in Higher Genus

It is worthwhile mentioning that decomposition attacks may also be applied to hyperelliptic curves $\mathcal{C}/\mathbb{F}_{q^n}$ with genus $g > 1$. The caveat is that due to the more complicated group law (the set of points no longer possess a natural group structure), summation polynomials are not applicable in this setting. To overcome this, Nagao [28] devised a method using divisors and Riemann-Roch spaces which (for hyperelliptic curves) reduced relation search to solving a quadratic multivariate polynomial system. Nagao's technique is less efficient than Semaev's in the elliptic curve setting, but is the best known approach in all other cases. Fixing both the extension and genus (n and g respectively), the complexity is estimated to be $\mathcal{O}(q^{2-2/ng})$. Thus for example taking $n = 2$ and $g = 3$ gives a running time of $\mathcal{O}(q^{5/3})$, which is asymptotically faster than Pollard's rho algorithm where running time is $\mathcal{O}(q^3)$ (note that the relevant group in this setting is the *Jacobian* of a genus g hyperelliptic curve over \mathbb{F}_{q^n} , which has size approximately q^{ng}). However for the same reasons as with summation polynomials, this approach is limited in use to cases where n and g are small.

4.4 Weil Descent

As mentioned in Chapter 1, Weil descent is a technique that allows one to map instances of the ECDLP in elliptic curves over extension fields \mathbb{F}_{q^n} into the DLP of a higher genus curve ($g \geq n$) defined over \mathbb{F}_q . The *GHS attack* refines the so-called Weil descent methodology to allow the reduction of ECDLP instances into those of hyperelliptic curves defined over a smaller field. In particular, the problem would be embedded into the Jacobian of a hyperelliptic curve \mathcal{C} of genus g , denoted $J_{\mathcal{C}}$.

In characteristic 2, it is known that the genus of \mathcal{C} will always be either 2^{m-1} or $2^{m-1} - 1$, where m may be determined via the following theorem. The value m is referred to as the *magic number*.

Theorem 4.4.1. [13] *Suppose l and n are positive coprime integers and $q = 2^l$. Let $k = \mathbb{F}_q$ and $K = \mathbb{F}_{q^n}$, and consider the elliptic curve E/K given by the equation*

$$E : Y^2 + XY = X^3 + aX^2 + b.$$

Let $a_i = \sigma^i(a)$ and $b_i = \sigma^i(b)$, where $\sigma : K \rightarrow K$ is the Frobenius automorphism defined by $\alpha \mapsto \alpha^q$. Allow \mathbf{u}_i to denote the vector $(1, a_i, b_i^{1/2})$, and let $U = \text{Span}_{\mathbb{F}_2}\{\mathbf{u}_i\}$ where i ranges from 0 to $n - 1$. Also let $V = \{(0, x^2 + x) : x \in K\}$. Finally we define

$$m = \dim_{\mathbb{F}_2}(U/U \cap V).$$

If $a \in \{0, 1\}$, then this expression simplifies to

$$m(b) = \dim_{\mathbb{F}_2}(\text{Span}_{\mathbb{F}_2}\{\mathbf{w}_i\}), \tag{4.1}$$

where $\mathbf{w}_i = (1, b_i^{1/2})$ for $i = 0, \dots, n - 1$.

A few key things to note are that $\sigma^n = \pi$, which is the identity map for all $\alpha \in K$, hence σ has order n as a function (which is why we let i range from 0 to $n - 1$). Secondly, the simplified expression (4.1) indicates that in these cases m depends only on the value of b .

The reason such an approach could result in a faster DLP algorithm is because in the hyperelliptic curve setting there are known subexponential-time algorithms that may be applied. One can refer to [10, Chapter 15.6] for an overview of available methods. In order to decide whether Weil descent is worth using for a particular ECDLP instance, we compare the expected run time of such algorithms with that of Pollard's rho algorithm. Broadly speaking, this approach can be effective for some families of elliptic curves, but fails in general if the genus g of the covering curve \mathcal{C} is too large. One expects the size of the Jacobian of \mathcal{C} to be $\approx q^g$, and the running time of known subexponential algorithms are affected greatly by this quantity. Hence a large g renders these algorithms infeasible.

4.5 Cover and Decomposition Attack

Joux and Vitse showed how one might combine Weil descent (or its variants, commonly referred to as *cover attacks*) with Nagao's decomposition to create an effective attack on the ECDLP for curves over composite extension fields. We give an overview of their method here, but refer to [19] for details.

Suppose that $\mathbb{F}_{q^n}/\mathbb{F}_p$ is an extension of finite fields with q being a power of p , and E an elliptic curve defined over \mathbb{F}_{q^n} . Note that p itself is a large prime in most applications, but more generally p can be any prime power. Although a decomposition attack can be applied in this setting, if n is bigger than 5 the approach is quite impractical. Instead, one can utilize the tower of extensions $\mathbb{F}_{q^n} - \mathbb{F}_q - \mathbb{F}_p$ and combine both the cover and decomposition

attacks. The general idea is to first employ (GHS) Weil descent or some variant on the extension $\mathbb{F}_{q^n}/\mathbb{F}_q$ to obtain a (hyperelliptic) covering curve defined over \mathbb{F}_q , with a small enough genus (relative to directly applying Weil descent on the extension $\mathbb{F}_{q^n}/\mathbb{F}_p$). From here we can use a decomposition attack on this new curve with the extension $\mathbb{F}_q/\mathbb{F}_p$.

Joux and Vitse also proposed a variant of Nagao’s decomposition technique that is reminiscent of the number field and function field sieves. It was shown that while the complexity of this variant is asymptotically higher, a smaller hidden constant implies that it is faster in practice for smaller values of q and p .

To showcase the applicability of their attack, analyses were done in fields of the form \mathbb{F}_{p^6} and \mathbb{F}_{p^4} . Joux and Vitse looked at the two different towers $\mathbb{F}_{p^6} - \mathbb{F}_{p^3} - \mathbb{F}_p$ and $\mathbb{F}_{p^6} - \mathbb{F}_{p^2} - \mathbb{F}_p$ in the former case, and the tower $\mathbb{F}_{p^4} - \mathbb{F}_{p^2} - \mathbb{F}_p$ in the latter.

For \mathbb{F}_{p^6} , the asymptotic complexity of their attack was estimated to be $\Theta(p^{5/3})$ using Nagao’s decomposition for either tower, and is expected to directly affect $\Theta(p^4)$ curves as compared to the expected complexity $\Theta(p^3)$ of Pollard’s rho method which affects all curves. However, it was seen that using a genus 3 cover (that is, the smallest attainable, non-trivial genus using the tower $\mathbb{F}_{p^6} - \mathbb{F}_{p^2} - \mathbb{F}_p$) reduced the problem to a smaller polynomial system than using a genus 2 cover (which is the smallest attainable, non-trivial genus in the $\mathbb{F}_{p^6} - \mathbb{F}_{p^3} - \mathbb{F}_p$ case). Hence the former approach is better in practice. The use of their variant yielded an asymptotic complexity of $\Theta(p^{12/7})$. Analysis of the case \mathbb{F}_{p^4} showed that the attack was effective, although the improvements were less significant.

Joux and Vitse provided an example of their attack in \mathbb{F}_{p^6} with $p = 33554467$. The size of the elliptic curve group in their example was a 149-bit prime, and they successfully computed discrete logarithms using their attack which otherwise would not have been practical using any previously known algorithm.

4.6 Isogeny walks

In 2002, Galbraith, Hess and Smart [11] proposed a novel way to increase the number of curves for which the Weil descent attack could be applied. The principal idea of their work was that one could use isogenies to map the discrete logarithm from a seemingly safe curve into one for which the GHS attack is effective. We present the ideas for their algorithm here.

Note that we will be considering elliptic curves whose endomorphism ring is the maximal order \mathcal{O}_L of the corresponding quadratic number field L . Typically one should be able to use Kohel’s algorithm to find a chain of isogenies from a particular curve E to one in the maximal order and carry out the isogeny walk from there, but this step can be problematic

if $c = [\mathcal{O}_L : \text{End}(E)]$ is divisible by a large prime (the largest prime dividing this value is known as the *conductor gap*). In this case Kohel's algorithm is inefficient, but we are choosing to not worry about this since such a case rarely occurs in practice for randomly selected elliptic curves.

As usual, let $k = \mathbb{F}_q$ where q is some prime power. In order to simplify notation, we will suppose that $\text{char}(k) \neq 2, 3$. Let E/k be an elliptic curve in the form (2.2) and l be a prime not equal to $\text{char}(k)$. We wish to address how one finds an l -isogeny to some elliptic curve E_1/k (which at this point is unknown).

Suppose that $\text{End}_k(E) \simeq \mathcal{O} \subseteq L = \mathbb{Q}(\pi)$, where π denotes the Frobenius endomorphism. Letting t denote the trace of Frobenius, recall that

$$\pi^2 - t\pi + q = 0.$$

From our discussion in Section 2.2, we know how the behaviour of $(l) = l\mathcal{O}$ as an ideal of \mathcal{O} reflects the number of norm- l ideals, and hence the number of horizontal l -isogenies over \mathbb{F}_q . We assume l is a prime that splits in \mathcal{O} , noting that the ramifying case is easier and in the inert case there's nothing to do. Since l splits, we can write (l) as the product of two prime ideals of norm l , say $(l) = \mathfrak{l}_1\mathfrak{l}_2$. It follows that the characteristic polynomial of Frobenius factorizes as

$$X^2 - tX + q = (X - \mu)(X - \lambda) \pmod{l}$$

for some $\mu, \lambda \in \mathbb{Z}/l\mathbb{Z}$. Thus we can write

$$l = (l, \pi - \mu)(l, \pi - \lambda).$$

Assume that we start with an elliptic curve E/\mathbb{F}_q having endomorphism ring \mathcal{O} . We will keep track of the walk by storing only the j -invariant of the current curve along with an ideal that is updated with each step. The ideal will contain information about an isogeny that maps from the starting curve to the current one. Note that even though twist curves have the same j -invariant, our starting curve is unambiguous so there will be no confusion as the algorithm progresses. We select a set \mathcal{F} of small primes such that \mathcal{F} satisfies two properties: First, the set of prime ideals \mathfrak{l} corresponding to the primes in \mathcal{F} should generate the ideal class group of \mathcal{O} (the idea here being that we want our isogeny graph to be connected). Second, we need enough primes so that our walk 'looks' random.

Our primes should be chosen so that most split in \mathcal{O} , though some ramified primes may be used as well. According to [1], one chooses an upper bound

$$L = 6(\log(|t^2 - 4q|))^2$$

for the possible primes and takes all suitable primes less than L to be in \mathcal{F} . The authors of [11] suggest that in practice 16 distinct split primes is generally sufficient for the walk to succeed. In addition to \mathcal{F} , we need a deterministic, but near uniform function $f : k \rightarrow \mathcal{F} \times \{0, 1\}$, which essentially allows us to take a j -invariant and deterministically obtain a prime and the resolution of a coin-flip. The walk now proceeds as follows.

The algorithm initializes with the pair $(j_1 = j, \mathfrak{a}_1 = \mathcal{O})$. At the i th stage, suppose we have a pair (j_i, \mathfrak{a}_i) where j_i is a j -invariant for some curve and \mathfrak{a}_i is an ideal. One applies f to j_i to obtain a pair (l, b) and compute $\Phi_l(X, j_i)$. Factoring $\Phi_l(X, j_i)$ in $\mathbb{F}_q[X]$ returns two new j -invariants. The bit b is then used to decide (in some deterministic fashion) which of these two j -invariants becomes j_{i+1} , and the corresponding isogeny is stored as an ideal \mathfrak{l} . In the event that two j -invariants arise, we have that $l\mathcal{O}$ splits into two ideals, one of which corresponds to the relevant isogeny. [11] outlines a method to determine the correct association, but we are content with the fact that this may be done unambiguously. Once we have this, let $\mathfrak{a}_{i+1} = \mathfrak{a}_i \cdot \mathfrak{l}$ and update our pair to be $(j_{i+1}, \mathfrak{a}_{i+1})$.

Note that one should actually take $\mathfrak{a}_{i+1} = \text{Reduce}(\mathfrak{a}_i \cdot \mathfrak{l})$, where the function `Reduce` returns the element of least norm in the ideal class group of $\mathfrak{a}_i \cdot \mathfrak{l}$. The reason for this is that if one is actually trying to compute a discrete log, reducing helps us keep the ideal from getting too large (and hence the corresponding isogeny from becoming unwieldy). In fact once the walk terminates, one should smooth the ideal output into a product of small sized ideals to facilitate storage and computation. We only mention this as an aside because our focus is on the reach of the isogeny walk strategy rather than actually computing logarithms. These details may be found in [11]. At each step we check whether or not the new curve (which corresponds to the stored j -invariant) is vulnerable, and halt if such a curve is found.

4.6.1 Modeling the Walk

From the above description it is entirely possible that this walk never ends, so it seems sensible to set a limit on the number of steps to take before calling it quits. Typically, one assumes a roughly uniform random distribution of weak curves so that combined with Theorem 3.5.3 we can easily model the probability of encountering a weak curve during the walk. Phrased as a problem one might find in an introductory statistics textbook, consider an isogeny class as a bin and the isomorphism classes as coloured balls. We will say the weak ones are red and the others are blue, and that there are N balls in total. Our problem is now the following: If we are drawing balls at random from the bin (with replacement), how many draws are needed until we can expect to have found a red ball?

If the total number of red balls is n , one expects to draw a red ball after $\lceil N/n \rceil$

steps. This is of course an oversimplification since whenever the function f gives a split prime, we are presented with two j -invariants and it makes sense to check if either of the corresponding curves is weak before selecting one at random. This would be analogous to drawing two balls at once and checking if either one is red. If one includes any ramified primes in \mathcal{F} , the scenario is slightly different and one must take into account the probability of selecting such a prime. Let p denote the probability of selecting a ramified prime. Then the probability of success after one draw is

$$P(\text{red}) = p\left(\frac{n}{N}\right) + (1 - p)\left(1 - \left(\frac{N - n}{N}\right)\left(\frac{N - n - 1}{N - 1}\right)\right).$$

One may then take the reciprocal of this value to determine the expected number of steps.

We performed a number of experiments in some small fields (\mathbb{F}_{218} , \mathbb{F}_{224} , \mathbb{F}_{230}) and checked that this seemed to provide a reasonable approximation for the number of steps taken in the random walk. Note that in the event that the value n is not known, we can find an estimate on the expected number of steps needed to reach a weak curve and solve for n to get an approximation (we will use this strategy later).

To obtain an estimate on the number of steps needed for a certain success probability, one can model the steps as a random variable and apply Chebyshev's Inequality. Alternatively, a well-studied statistics problem known as the coupon collector's problem [26] gives us the probability of not having reached all nodes after some number of steps. In particular, we have

$$P(T > \beta N \log N) < N^{-\beta+1},$$

where T is the time taken to reach all nodes (steps), N is the number of nodes, and β is some constant. A quick calculation gives that one needs $T \geq (\log N + t \log(10))$ to make the probability less than 10^{-t} .

Chapter 5

Distribution of Curves

In this chapter we will take a closer look at some of the common assumptions that are made in analyses of the isogeny walk strategy. In particular, we would like to discuss the distribution of isomorphism classes of curves amongst isogeny classes, as well as the distribution of curves which are directly vulnerable to the GHS Weil descent attack (those which yield a low genus cover). Since one of our main goals will be to experiment with the effectiveness of Joux-Vitse's cover and decomposition attack on large fields of characteristic 2, we will focus on the case of elliptic curves defined over fields of the form \mathbb{F}_{q^6} , where $q = 2^l$ for some positive integer l . We will perform exhaustive calculations on a few small fields ($\mathbb{F}_{2^{18}}$, $\mathbb{F}_{2^{24}}$ and $\mathbb{F}_{2^{30}}$), as well as limited experiments on some medium-sized fields ($\mathbb{F}_{2^{42}}$, $\mathbb{F}_{2^{48}}$ and $\mathbb{F}_{2^{54}}$) in order to gain confidence in these assumptions.

5.1 Distribution of Isomorphism Classes

Suppose we have fixed a base field \mathbb{F}_q . Adopting the notation of [32], we let $I(t)$ denote the isogeny class of elliptic curves over \mathbb{F}_q having exactly $q + 1 - t$ points and $N(t)$ denote the number of \mathbb{F}_q isomorphism classes in $I(t)$. We have the following theorem paraphrased from [32] which gives us a precise value for the size of an isogeny class of an ordinary elliptic curve.

Theorem 5.1.1. *Let $t \in \mathbb{Z}$ with $t^2 < 4q$ and $p \nmid t$. Then $N(t) = H(t^2 - 4q)$, where $H(\Delta)$ denotes the Kronecker class number.*

One notes that this theorem applies precisely to the ordinary elliptic curves due to Hasse's Theorem and Definition 2.1.2.

The Kronecker class number counts the number of reduced binary quadratic forms of discriminant $\Delta < 0$, which is equal to the class number of a quadratic number field with the same discriminant. For our purposes, we will simply acknowledge that this value is computable via the following theorem (though faster methods exist and are available in Magma).

Theorem 5.1.2. *Let $\Delta \in \mathbb{Z}, \Delta < 0$. Then*

$$H(\Delta) = \{(a, b, c) \in \mathbb{Z}^3 : a > 0, b^2 - 4ac = \Delta, |b| \leq a \leq c, \\ \text{and } b \geq 0 \text{ whenever } a = |b| \text{ or } a = c\}.$$

Now that we have a precise formula for the size of a particular isogeny class, the next question to answer is how these values compare across all isogeny classes. A result of Lenstra's [21] gives us some insight to this question by bounding the potential sizes above and below. Before we can state this result we will first need some notation.

Restricting to curves over a field \mathbb{F}_p where $p > 3$ is prime, there are approximately $p^2 - p$ possible nonsingular Weierstrass equations. Letting $\text{Aut}(E)$ denote the set of all isomorphisms from a curve E/\mathbb{F}_p to itself, the number of curves isomorphic to a particular curve E is $\#\mathbb{F}_p^*/\#\text{Aut}(E) = (p - 1)/\#\text{Aut}(E)$. If we sum over a set of representatives for the isomorphism classes and divide by $(p - 1)$, we obtain the equation

$$\sum_E \frac{1}{\#\text{Aut}(E)} = p.$$

This is expressed by writing

$$\#\{E : E \text{ is an elliptic curve over } \mathbb{F}_p\} / \cong_{\mathbb{F}_p} = p,$$

where $\#'$ denotes the weighted cardinality in which the isomorphism class of a curve E is counted with weight $(\#\text{Aut}E)^{-1}$.

Proposition 5.1.3. [21, Proposition 1.8] *There exist efficiently computable positive constants c_1, c_2 such that for each prime number $p > 3$ the following two assertions hold:*

1. *If S is a set of integers s with $|s - (p + 1)| \leq 2\sqrt{p}$ then*

$$\#\{E : E \text{ elliptic curve over } \mathbb{F}_p, \#E(\mathbb{F}_p) \in S\} / \cong_{\mathbb{F}_p} \\ \leq c_1 \cdot \#S \cdot \sqrt{p} \cdot (\log p)(\log \log p)^2.$$

2. If S is a set of integers s with $|s - (p + 1)| \leq \sqrt{p}$ then

$$\begin{aligned} \#\{E : E \text{ elliptic curve over } \mathbb{F}_p, \#E(\mathbb{F}_p) \in S\} / \cong_{\mathbb{F}_p} \\ \geq c_2 \cdot (\#S - 2) \cdot \sqrt{p} / (\log p). \end{aligned}$$

Since most elliptic curves possess $\#Aut(E) = 2$, we can heuristically replace the weighted cardinality with the usual cardinality and absorb the factor of 2 into the constants. While the version of this theorem that appears in [21] requires the elliptic curves be defined over \mathbb{F}_p for primes $p > 3$, the result can be adjusted for those defined on \mathbb{F}_q where q is any prime power.

Essentially what this tells us is that given an interval centered at $q + 1$ (the middle of the Hasse interval), the sizes of the isogeny classes should be roughly bounded. Admittedly this is a bit vague, so in order to gain a better understanding of what these bounds actually look like, we used the computational algebra package Magma to exhaustively compute the isogeny class sizes for fields of size 2^{18} , 2^{24} and 2^{30} . Before presenting these findings, we first state a lemma that helps to simplify some of our computations.

Lemma 5.1.4. *Let $q = p^m$ and E_1/\mathbb{F}_{q^n} be an elliptic curve with general Weierstrass equation*

$$E_1 : Y^2 + a_1XY + a_3Y = X^3 + a_2X^2 + a_4X + a_6, \quad (5.1)$$

with coefficients in \mathbb{F}_{q^n} , and let σ denote the q -power Frobenius map, which takes rational points $P = (X, Y) \mapsto (X^q, Y^q)$ and $\mathcal{O}_{E_1} \mapsto \mathcal{O}_{E_2}$. Then σ is an isogeny of degree q . Furthermore, if $\varphi : E_1 \rightarrow E_2$ is an isogeny over \mathbb{F}_{q^n} of degree d , then there exists a degree- d isogeny $\varphi_q : E_1^q \rightarrow E_2^q$, where E_i^q is the curve obtained by raising all the coefficients of E_i to the q th power.

Proof. Observe that σ is a rational function by definition and takes the point at infinity from one curve to the other, hence it is an isogeny. The fact that the degree of σ is q follows from X^q having degree q (note that although we did not provide σ in standard form, the transformation will only affect the function's second coordinate).

Observe now that E_i^q is the image curve of E_i under σ . One can see this by considering the general equation

$$Y^2 + a_1XY + a_3Y = X^3 + a_2X^2 + a_4X + a_6$$

and raising both sides to the q^{th} power. Since q is a power of the characteristic of the field, all cross terms will turn to zero, leaving us with the equation

$$Y^{2q} + a_1^q X^q Y^q + a_3^q Y^q = X^{3q} + a_2^q X^{2q} + a_4^q X^q + a_6^q.$$

Thus (X^q, Y^q) will satisfy

$$Y^2 + a_1^q XY + a_3^q Y = X^3 + a_2^q X^2 + a_4^q X + a_6^q,$$

whenever $P = (X, Y) \in E_i(\overline{\mathbb{F}}_{q^n})$. Supposing we have some degree- d isogeny φ in standard form, one can verify using a similar argument that by raising all the coefficients of the isogeny φ to the q^{th} power, one obtains an isogeny between the curves E_1^q and E_2^q . Moreover, changing the coefficients does not change the degree of the map, so this new isogeny is also of degree d . \square

5.1.1 Small Fields

The graphical representations of our findings are displayed in Figures 5.1, 5.2 and 5.3. Note that due to Lemma 5.1.4 we can treat curves that lie in the same orbit under action by σ as equivalent; we refer to these colloquially as *orbits* or *orbit classes*. Recalling that σ has order at most n , this allows us to scale our graphs by the extension degree of the field (note that although there may be some orbits with size strictly dividing n , this number is small enough that we can ignore them).




Figure 5.1: Plot displaying the size of isogeny classes over $\mathbb{F}_{2^{18}}$ as a function of the trace of Frobenius. The isogeny classes have been shifted so that those corresponding to the low end of the Hasse interval appear on the left side of the plot.




Figure 5.2: Plot displaying the size of isogeny classes over $\mathbb{F}_{2^{24}}$ as a function of the trace of Frobenius. The isogeny classes have been shifted so that those corresponding to the low end of the Hasse interval appear on the left side of the plot.




Figure 5.3: Plot displaying the size of isogeny classes over $\mathbb{F}_{2^{30}}$ as a function of the trace of Frobenius. The isogeny classes have been shifted so that those corresponding to the low end of the Hasse interval appear on the left side of the plot.

Observe that these experiments indicate that the distribution of isomorphism classes take on the shape of a curved band with highest density in a horizontal bar near the middle and dropping off near the ends. This is consistent with Theorem 5.1.3 which bounds the sizes of the isogeny classes above on the Hasse interval, and below near the middle. Our reasoning for selecting these particular fields is two-fold: their extension degrees are divisible by 6, which means that one may apply a cover and decomposition attack to curves defined over these fields in the same fashion as [19] (see Section 4.5), and these were the largest such fields that we could exhaustively obtain data for vulnerable curves (which we will see in the following section).

5.1.2 Medium Fields

The fields selected in Section 5.1.1 are hardly comparable to those typically used for cryptographic applications in terms of size, so in an attempt to provide more convincing evidence, experiments in slightly larger fields were also performed. The fields selected were of size 2^{48} and 2^{54} . As it was not feasible to do exhaustive computations, we sampled a random subset of curves over these fields. The sample data (Figures 5.4 and 5.5) appear to be consistent with Theorem 5.1.3.

5.2 Distribution of Weak Curves

One of the key assumptions needed for the evaluation of the effectiveness of the random walk is that the curves vulnerable to the GHS attack (i.e. with the smallest non-trivial ‘ m ’ value) are distributed roughly uniformly amongst the isogeny classes. We sketch the heuristic analysis from [11] for completeness.

Suppose we have a fixed field $K = \mathbb{F}_{q^n}$, where $q = 2^r$. By Hasse’s theorem there are about $4q^{n/2}$ potential isogeny classes, and Theorem 2.1.3 gives that amongst these roughly $2q^{n/2}$ are ordinary isogeny classes. By Theorem 2.1.10, the number of distinct isomorphism classes is approximately $2q^n$.

In order to compute the number of weak classes, let t be the multiplicative order of 2 mod n , and $s = (n - 1)/t$. Then the number of weak isomorphism classes may be bounded above by $2sq^{t+1}$. The probability of a random curve being vulnerable is thus estimated to be

$$\frac{2sq^{t+1}}{2q^n} = sq^{t+1-n}.$$




Figure 5.4: Plot displaying the size of isogeny classes for roughly 100 000 random curves over $\mathbb{F}_{2^{48}}$ as a function of the trace of Frobenius. The isogeny classes have been shifted so that those corresponding to the low end of the Hasse interval appear on the left side of the plot.

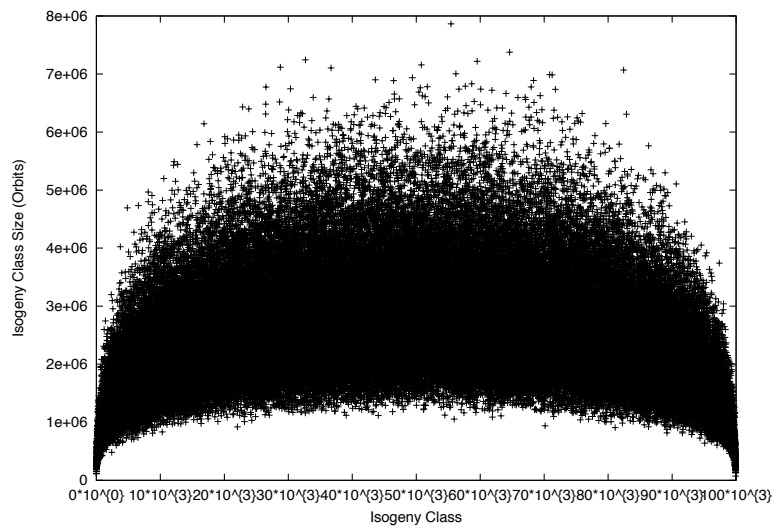


Figure 5.5: Plot displaying the size of isogeny classes for roughly 100 000 random curves over $\mathbb{F}_{2^{54}}$ as a function of the trace of Frobenius. The isogeny classes have been shifted so that those corresponding to the low end of the Hasse interval appear on the left side of the plot.

One can also verify that the 2-power Frobenius map preserves the value of m , and it is assumed this is the only map with such a property. In most cases, the orbit of the 2-power Frobenius map is nr . What this means is that for the most part, vulnerable curves distribute across isogeny classes in sets of size nr . Hence if we treat curves as equivalent if they are related by 2-power Frobenius action (recall these were referred to as *orbit classes*), the number of weak curves becomes

$$\frac{2sq^{t+1}}{nr}.$$

Under the assumption of a roughly uniform distribution amongst the isogeny classes, the expected number of weak orbit classes in an arbitrary isogeny class is expected to be close to

$$\frac{sq^{t+1}}{nrq^{n/2}} = \frac{sq^{t+1-n/2}}{nr}.$$

In the special case of $n = 3$ we have $t = 2$ and $s = 1$, so our analysis is slightly more concrete. The number of isomorphism classes is precisely $2(q^3 - 1)$, and the number of weak curves can be enumerated simply by counting the number of elements b such that $Tr_{\mathbb{F}_{q^3}/\mathbb{F}_q}(b) = 0$ and doubling, which yields a value of $2(q^2 - 1)$ (note that this is somewhat smaller than our above upper bound). Our probabilities become roughly $1/q$ for a random curve being weak, and we would expect

$$\sqrt{q}/nr$$

weak orbit classes per isogeny class. This ratio is less than 1 for small fields (for example, in $\mathbb{F}_{2^{24}}$, the ratio is $16/24$), but as r increases this number grows proportionally.

Note that ahead of time we know that around half of the ordinary isogeny classes will not contain any weak curves. This follows from the fact that $Tr_{\mathbb{F}_{q^3}/\mathbb{F}_q}(b) = 0$ implies that $Tr_{\mathbb{F}_{q^3}/\mathbb{F}_2}(b) = 0$, and the following theorem:

Theorem 5.2.1. [25, Lemma 7] *Let $E : y^2 + xy = x^3 + b$ be an elliptic curve over $\mathbb{F}_{2^N} = \mathbb{F}_{q^N}$ where $N \geq 3$. Then $Tr_{\mathbb{F}_{q^3}/\mathbb{F}_2}(b) = 0$ if and only if $\#E(\mathbb{F}_{2^N}) \equiv 0 \pmod{8}$.*

Recall that the twist E' of weak curve E is also weak. Further, by Theorem 2.1.12, if $\#E(\mathbb{F}_{2^N}) = 2^N + 1 - t$, then $\#E'(\mathbb{F}_{2^N}) = 2^N + 1 + t$. Thus if $\#E(\mathbb{F}_{2^N}) \equiv 0 \pmod{8}$, it is easily verified that $\#E'(\mathbb{F}_{2^N}) \equiv 2 \pmod{8}$. It now follows that we would only expect half of the ordinary isogeny classes (those which contain curves with cardinality congruent to 0 or 2 $\pmod{8}$) to contain any weak curves. It is conjectured that all isogeny classes that satisfy this criterion contain a weak curve, and are thus vulnerable to the cover and decomposition attack [19].

5.2.1 Small Fields

In order to gain confidence in the assumption on the distribution of weak curves, we exhaustively computed the number of weak curves per isogeny class for elliptic curves defined over $\mathbb{F}_{2^{18}}$, $\mathbb{F}_{2^{24}}$ and $\mathbb{F}_{2^{30}}$. It should be noted that we shall be treating these as fields of the form \mathbb{F}_{q^n} , where $n = 3$ is fixed. Thus we directly make use of the above analysis. We make note that the number of weak orbit classes was obtained by first determining the total number of weak isomorphism classes, say x , and then computing $\lfloor x/nr \rfloor$, where nr is the degree of the field of definition over \mathbb{F}_2 . As noted earlier this is a heuristic, but the number of orbits that deviate from this size is practically negligible.

We then compared our data to an experimental model in which one assumes that each isogeny class is roughly the same size. This assumption is justified by Theorem 5.1.3 and the experiments performed in Section 5.1. We frame the model as a simple statistics problem: Suppose that we have N bins and we are throwing n balls into these bins at random. We would like to determine the expected number of bins containing exactly r balls for all $r < n$.

Consider the probability that a particular bin contains exactly r balls. This probability is given by

$$p = \binom{n}{r} \left(\frac{1}{N}\right)^r \left(1 - \frac{1}{N}\right)^{n-r}.$$

Letting X_i denote the Bernoulli random variable

$$X_i = \begin{cases} 0 & \text{if bin } i \text{ contains exactly } r \text{ balls,} \\ 1 & \text{otherwise,} \end{cases} \quad (5.2)$$

the expected value of X_i is $E(X_i) = p$. By the linearity of expectation, one anticipates the number of bins with r balls to be

$$E(X_1 + \cdots + X_N) = N \binom{n}{r} \left(\frac{1}{N}\right)^r \left(1 - \frac{1}{N}\right)^{n-r}.$$

Figures 5.6, 5.7, and 5.8 compare the theoretical model (top) with the data for our three chosen fields \mathbb{F}_{18} , \mathbb{F}_{24} and \mathbb{F}_{30} . One notes the discrepancy between the data and the predicted number of isogeny classes with no weak curves for \mathbb{F}_{24} and \mathbb{F}_{30} . This difference could be due simply to statistical variation, but given how similar the distribution is to the theoretical model otherwise, it is possible that there is some other condition not being considered that prevents some isogeny classes from having any weak curves. Further experimentation would need to be done to strengthen such a claim.





Figure 5.6: Number of isogeny classes defined over \mathbb{F}_{218} possessing a given number of weak orbits classes. Above is the theoretical expectation under the proposed model; below is the data collected.





Figure 5.7: Number of isogeny classes defined over $\mathbb{F}_{2^{24}}$ possessing a given number of weak orbits classes. Above is the theoretical expectation under the proposed model; below is the data collected.





Figure 5.8: Number of isogeny classes defined over \mathbb{F}_{230} possessing a given number of weak orbits classes. Above is the theoretical expectation under the proposed model; below is the data collected.

5.2.2 Medium Fields

As mentioned earlier, in order to gain a more accurate idea of the weak curve distribution we attempted to gather data for curves defined on the slightly larger fields $\mathbb{F}_{2^{42}}$, $\mathbb{F}_{2^{48}}$ and $\mathbb{F}_{2^{54}}$, where exhaustive computation was no longer possible. Our experiments were done by randomly selecting curves E such that $t^2 - 4q$ is a squarefree integer (which occurs asymptotically with probability $6/\pi^2$), and such that the congruence conditions discussed at the end of Section 5.2 were satisfied. The first condition ensures that the endomorphism ring of E is the maximal order of its associated quadratic number field. For the isogeny walk, we selected our prime bound to be 59. From Section 4.6, the expander property is guaranteed if all primes less than $L = 6(\log(|t^2 - 4q|))^2$ are selected, which is maximized when $t = 0$. Hence for our largest field we would need to have selected all primes less than around 9040, which is much too large for practical purposes. Upon implementation, we found (similar to Galbraith et al. [11] who argued for 16 split primes) that much lower numbers sufficed. Incidentally, the default Magma database contains only the (classical) modular polynomials up to $N < 60$, so we thought this was a reasonable bound to choose.

Using Magma, we performed a random walk on the orbit classes in the isogeny graph. The representative for a particular orbit class was determined by computing the ‘ b ’ coefficient of each elliptic curve in the class (done by successive squaring) and reinterpreting them as integers (recall that an element of \mathbb{F}_{2^n} may be represented as a polynomial with coefficients in \mathbb{F}_2 of degree less than n , and one can obtain an integer in binary by reading off the coefficients). The element with the smallest such integer representation was taken as the orbit’s representative. Such a walk was performed until we either encountered a weak curve, or took enough ‘steps’ so that the probability of there being a weak curve and us not having found it was sufficiently small (less than 1%, see Section 4.6.1).

For curves which the walk yielded a weak curve, we repeated the experiment and averaged the number of steps taken, which allowed us to estimate the number of weak orbit classes in the isogeny class. Note that in order to obtain a better estimate, it would be necessary to repeat these walks many times, as well as begin the walk at different curves in the same isogeny class. Due to computational constraints, we opted to perform walks in more isogeny classes rather than improving the estimates, as our primary concern was to gather evidence that most isogeny classes do indeed contain a weak curve in cryptographically sized fields.

Due to the above drawbacks in our experimental method, we are hesitant to draw any conclusions about the number of weak curves in the isogeny classes, however we include the results for posterity (Figures 5.9, 5.10 and 5.11). One notes that the experimental data is much less concentrated than the theoretical models, but this can be explained by the fact that starting curves of our walk experiments could vary greatly in terms of distance to

weak curves. This could either shorten or lengthen a walk and lead to skewed estimates on the number of weak curves. Since our experiments were quite few, these anomalies would have had a significant impact on the data.

What can be reasonably concluded from these graphs is that the proportion of curves that do not contain any weak curves seem to be decreasing as the fields size grows, which is what the theory suggests.





Figure 5.9: Estimated number of isogeny classes containing a certain number of weak orbits; data is for a sample of roughly 5000 isogeny classes of curves defined over $\mathbb{F}_{2^{42}}$. Above is the theoretical expectation under the proposed model; below is the data collected. The expected number of weak orbits per isogeny class (under a uniform distribution) is roughly 3.





Figure 5.10: Estimated number of isogeny classes containing a certain number of weak orbits; data is for a sample of roughly 5000 isogeny classes of curves defined over \mathbb{F}_{248} . Above is the theoretical expectation under the proposed model; below is the data collected. The expected number of weak orbits per isogeny class (under a uniform distribution) is roughly 5.





Figure 5.11: Estimated number of isogeny classes containing a certain number of weak orbits; data is for a sample of roughly 4000 isogeny classes of curves defined over $\mathbb{F}_{2^{54}}$. Above is the theoretical expectation under the proposed model; below is the data collected. The expected number of weak orbits per isogeny class (under a uniform distribution) is roughly 9.

Chapter 6

Experiments in Large Fields

In order to determine whether the isogeny walk is a viable strategy for attacking the ECDLP in fields of cryptographic size, we performed some experiments in the field $\mathbb{F}_{2^{150}}$ and $\mathbb{F}_{2^{210}}$. We chose the field $\mathbb{F}_{2^{150}}$ because it is a sextic extension of $\mathbb{F}_{2^{25}}$ and is comparable in size to the field used in the example found in the work of Joux and Vitse (see Section 4.5). The field $\mathbb{F}_{2^{210}}$ was chosen for similar reasons, but additionally this field is of special interest due to [25], which identified it as a potential candidate for being a bad field. We will begin with a brief overview of the known methods for computing discrete logarithms as they apply to curves defined over fields of the form \mathbb{F}_{q^3} where $q = 2^{2^m}$, followed by concrete performance estimates in the fields mentioned above.

6.1 Discrete Logarithms in \mathbb{F}_{q^3}

Let $K = \mathbb{F}_{q^3}$ where $q = 2^{2^m}$ for a positive integer m . The performance of any algorithm for solving the ECDLP in a curve E/K is compared to that of Pollard's rho, which can be applied to any curve defined over K . The expected running time of Pollard's rho is $\mathcal{O}(q^{3/2})$ and possesses essentially no memory requirements. An algorithm is deemed effective if its running time is better asymptotically than Pollard's rho, though in practice these asymptotic estimates may not tell the whole story. We describe the known such algorithms below.

One possible approach to solving discrete logarithms in K is to use Gaudry's index calculus algorithm. Roughly described in Section 4.3, this method works on all elliptic curves and requires $\mathcal{O}(q)$ memory. The running is expected to be $\mathcal{O}(q^{4/3})$, however there

are large hidden constants due to the use of a Gröbner basis computation to solve a large multivariate polynomial system.

Alternatively, one might be able apply the GHS reduction followed by the Enge-Gaudry index calculus algorithm. Using the extension $\mathbb{F}_{q^3}/\mathbb{F}_q$, the GHS attack directly targets roughly $2q^2$ curves of the total $2q^3$, or around $1/q$ curves. On its own, the reduction is quite fast, and the index calculus is the bottleneck in both memory and running time. The memory requirements are $\mathcal{O}(q)$ and achieves a running time of $\mathcal{O}(q^{4/3})$. The range of this approach can be extended to near q^3 curves using an isogeny walk, and which itself should take time $\mathcal{O}(q)$. Note that one may use the GHS with different extensions, though the options will depend on the specific value of m .

It is also theoretically possible to use the generalised version of the GHS attack [16], in which a genus-3 non-hyperelliptic curve defined over \mathbb{F}_q is returned. Using this method, it was observed in in [27] that around $q^{3/4}$ curves are directly affected, and this can be extended to all curves which possess $\text{Tr}(a) = 1$ and $\text{Tr}(b) = 0$. One can then make use of an algorithm by Diem and Thomé [5], which expects to run in time $\mathcal{O}(q)$ and uses memory $\mathcal{O}(q^{1/2})$. The main drawback of this method is that the gGHS has never been implemented, and operations in the Jacobian of the non-hyperelliptic curve are likely quite expensive.

Finally, we consider Joux-Vitse’s attack, which applies GHS on the extension $\mathbb{F}_{q^3}/\mathbb{F}_q$ followed by Nagao’s decomposition (or a variant) on the extension $\mathbb{F}_{2^{2m}}/\mathbb{F}_{2^m}$. This approach targets the same curves as the GHS attack, and is similarly extendable via isogeny walks in time $\mathcal{O}(q)$. Since Nagao’s method (or its variant) has running time $\mathcal{O}(q^{5/6})$ (respectively $\mathcal{O}(q^{6/7})$), the walk becomes a dominating factor in the running time analysis. Memory requirements in this setting are only $\mathcal{O}(q^{1/2})$.

6.2 The Field $\mathbb{F}_{2^{150}}$

Previous attempts at establishing the weakness of $\mathbb{F}_{2^{150}}$ made use of the GHS attack on the extension $\mathbb{F}_{2^{25 \cdot 30}}/\mathbb{F}_{2^{30}}$, where the smallest nontrivial magic number is $m = 5$, and predominantly yields a genus 16 hyperelliptic curve over the base field (see [25], Section 3.2 for a more precise statement). This attack is applicable to most curves (essentially all in fact) over this field. The time needed for the actual GHS reduction is considered negligible, so the dominating step is the Enge-Gaudry index-calculus algorithm for solving the resulting HCDLP instance. Suppose T_1 and T_2 denote the running times for the two phases of Enge-Gaudry, relation generation and linear-algebra, respectively. If we write $\mathbb{F}_{2^{150}} = \mathbb{F}_{2^{n \cdot l}}$ with $n = 5$ and $l = 30$, it was estimated in [25] that

$$T_1 \approx 2^{l+43} = 2^{73}$$

where the unit of operation is in essence one Jacobian addition and a smoothness test, and

$$T_2 \approx 2^{2l+2} = 2^{62}$$

where the operation is multiplication modulo the group size. It is safe to assume that the former operation is more time-consuming, so the algorithm is dominated by the relation generation phase. Comparatively, Pollard's rho algorithm targets all elliptic curves, and is expected to take $\sqrt{\pi 2^{150}}/2 \approx 2^{75}$ steps or elliptic curve additions. Though the GHS attack seems to do better by a factor of 4, it is evident that the more cumbersome Jacobian operation and smoothness test is more time-consuming than an elliptic curve addition, resulting in no improvement.

It should be noted that the relation generation phase of Enge-Gaudry and Pollard's rho are both effectively parallelizable, and essentially obtain a linear speedup in the number of processors. This is not known to be possible with the linear-algebra phase of Enge-Gaudry, so in practice the linear-algebra may be the algorithm's rate-limiting step. If one also takes into account memory requirements, Pollard's rho uses almost no memory, whereas the linear-algebra phase requires storing a matrix of size roughly $2^{30} \times 2^{30}$ (2^{30} is the typical size of the factor base in this setting). Thus any improvement that might be offered by the GHS method is likely not significant enough to warrant its use over Pollard's rho.

One can also implement the GHS method using the extension $\mathbb{F}_{2^{150}}/\mathbb{F}_{2^{50}}$. Here the GHS reduction yields a genus-3 hyperelliptic curve defined over $\mathbb{F}_{2^{50}}$ for approximately 2^{101} isomorphism classes over $\mathbb{F}_{2^{150}}$. The benefit of this approach is the *double-large prime variant* of index calculus (Section 4.3.1) offers the best improvement with this extension. If the GHS is used with the extension $\mathbb{F}_{q^n}/\mathbb{F}_q$ and provides a genus g hyperelliptic curve defined over \mathbb{F}_q , then the total running time of the attack is $\mathcal{O}(q^{2-\frac{2}{g}})$, though one must be wary of hidden constants. Nonetheless our case ends up with running time on the order of $2^{50(4/3)} \approx 2^{66}$ (the operation here is again Jacobian addition and smoothness testing). Roughly half of all curves, in particular those with $\text{Tr}(b) \equiv 0, 2 \pmod{8}$, may be reached with this attack using an isogeny walk, though a significant drawback is that it is not parallelizable since the main idea of the double-large prime variant is to balance the time-cost of relation-generation with the linear algebra. Thus even if we disregard the cost of the isogeny walk, with a moderate number of processors (around 2^9) this too is less effective than Pollard's rho.

We now consider Joux-Vitse's method, which is able to target half of all curves defined over $\mathbb{F}_{2^{150}}$ (in fact these are the same curves as the GHS attack on $\mathbb{F}_{2^{150}}/\mathbb{F}_{2^{50}}$). One expects to solve an ECDLP instance on a directly weak curve (those which yield a genus 3 hyperelliptic cover via the GHS attack, since we are using the tower $\mathbb{F}_{2^{150}} - \mathbb{F}_{2^{50}} - \mathbb{F}_{2^{25}}$) in

time $2^{25(5/3)} \approx 2^{42}$ using Nagao decomposition, or $2^{25(12/7)} \approx 2^{43}$ using the sieving variant. The unit of operation here is multiplication in $\mathbb{F}_{2^{25}}$, although these estimates are a bit misleading due to potentially large suppressed constants.

Supposing we are in an isogeny class that may contain a weak curve, i.e. one which contains elliptic curves with order congruent to 0 or 2 modulo 8, our heuristic assumptions suggest there are around 2^{76} isomorphism classes in an isogeny class, and roughly 2^{26} weak curves amongst them. These values come from the fact that there are roughly 2^{151} isomorphism classes of curves, 2^{75} isogeny classes, and 2^{101} weak curves. Therefore we expect to take about 2^{50} steps to find a weak curve. It seems reasonable to assume that a walk step is not 2^8 times faster than the unit of operation in the cover and decomposition, so we expect the isogeny walk to dominate the complexity of this attack.

Our purpose is to now compare the time needed to perform a random walk with that of Pollard's rho. Scaling, this reduces to comparing the time of a single walk step against 2^{25} Pollard rho steps or elliptic curve additions. If c_{150} denotes the cost of a single field multiplication in $\mathbb{F}_{2^{150}}$ and c_E denotes the cost of a single elliptic curve addition, we have that $c_E \approx 8c_{150}$ (with the use of mixed affine-projective coordinates). In our implementation of the random walk, the most time consuming step is the computation of $\mathbb{F}_{2^{150}}$ -roots for the polynomial $f = \Phi_l(X, j)$, where j is the j -invariant of an elliptic curve and l is a prime. This is done by computing $g = X^{2^{150}} - X \pmod f$, determining $\gcd(f, g)$, and then factoring the result (which is of degree at most 2). The time to find the g.c.d. and factor a degree 2 polynomial is overshadowed by the cost of computing g .

Breaking this up further, one computes $g = X^{2^{150}} - X \pmod f$ by successively squaring X modulo f (this is repeated 150 times). Since f is of degree $l + 1$, if we precompute the powers $X^2, X^4, \dots, X^{2^l} \pmod f$, the square of a polynomial

$$a_0 + a_1X + \dots + a_lX^l \in \mathbb{F}_{2^{150}}[X]$$

modulo f is equal to

$$a_0^2 + a_1^2X^2 + \dots + a_l^2X^{2^l} \pmod f.$$

Since our precomputed values will be polynomials of degree at most l , ignoring addition gives that the total work comes from squaring l distinct coefficients, and multiplying each with a degree- l polynomial. The cost of one polynomial-squaring can be estimated to be l^2c_{150} , so the cost to fully compute g is $150l^2c_{150}$. If we assume equal probability of a prime being split versus inert and aim for the 16 split prime mark as in [11], we could set our prime bound to be 131, which means on average, $l \approx 2^6$.

Thus, we find that a single walk step takes time roughly $150(2^6)^2c_{150} \approx 2^{19}c_{150}$, versus $2^{25}c_E \approx 2^{28}c_{150}$. Thus a random walk should be faster than Pollard's rho algorithm by

about a factor of 2^9 . We note that both the random walk and Pollard’s rho are parallelizable in a way that they are both essentially linear in the number of processors, and both have low storage requirements, so neither has advantage in these areas.

We implemented the random walk using an AMD Opteron 6168 processor at 2.3 GHz. As mentioned earlier, we selected our prime bound to be much lower than recommended (59), so on average $l \approx 2^{4.5}$. Using the above analysis, we would expect the walk to be close to $2^{28}/2^{16} = 2^{12}$ times faster than Pollard’s rho. Averaging over 1000 walks consisting of 100 steps each, we estimated the time needed to perform a single walk step to be roughly 0.0914 seconds. The estimated time for a single step of Pollard’s rho algorithm was 4.06×10^{-5} seconds. The heuristic analysis predicts that one would require 2^{50} random walk steps before expecting to find a weak curve, which amounts to a total time of 102907251485415.8336 seconds ($\approx 2.859 \times 10^{10}$ hours or over 3.2 million years). On the other hand, Pollard’s rho is expected to take 2^{75} steps, which totals to 1533824633636060765.4085 seconds ($\approx 4.261 \times 10^{14}$ hours or over 48 billion years). The ratio of these two values is $14903.812 \approx 2^{14}$, which seems reasonably close to the predicted ratio. The length of these times are not within the feasible range, so we cannot conclude that $\mathbb{F}_{2^{150}}$ is partially bad under this attack. However, our implementation of the random walk was admittedly very sub-optimal, and could likely be sped up significantly.

Hence we have given reasonable evidence that roughly half of isogeny classes are susceptible to an attack that performs significantly faster than Pollard’s rho, which suggests that $\mathbb{F}_{2^{150}}$ is a partially weak field. However to be absolutely certain, one would have to actually compute the isogeny from a successful walk and transfer an instance of the discrete logarithm problem from the initial elliptic curve to the weak one. We are assuming the process of computing and reducing the ideal at each step (see Section 4.6) as well as smoothing out the final ideal obtained from the random walk is fast compared to the walk itself (asymptotically, this has running time on the order of 2^{35} , but should be implemented for concreteness). To obtain the stronger result that $\mathbb{F}_{2^{150}}$ is partially bad, one would need to actually implement Joux-Vitse’s attack and use an isogeny walk to solve a discrete logarithm in an elliptic curve defined over this field.

6.3 The Field $\mathbb{F}_{2^{210}}$

Turning our attention now to the field $\mathbb{F}_{2^{210}}$, Menezes, Teske and Weng [25] proposed this as a prime candidate for being a bad field. We recall the details of their analyses, which also takes into account the relative time-consumption of the various operations. We will then investigate how Joux-Vitse’s cover and decomposition attack influences the argument. In this field, one expects Pollard’s rho to have a running time near 2^{110} (note that this

value has been scaled with respect to multiplication in $\mathbb{F}_{2^{42}}$); we will use this value as a benchmark to gauge the effectiveness of the methods described below.

Due to the high number of factors of 210, the GHS attack can be applied in a number of different ways. Writing $2^{210} = 2^{n \cdot l}$ and applying the GHS on the extension $\mathbb{F}_{2^{210}}/\mathbb{F}_{2^l}$, [25] analyzed the cases where $n = 5$ and 6 . In both cases, the smallest nontrivial m value is 5 . For the extension $\mathbb{F}_{2^{210}}/\mathbb{F}_{2^{42}}$, the vast majority of curves are susceptible to the GHS attack, and the running time of Enge-Gaudry (with notation as in Section 6.2) is estimated to be around

$$T_1 \approx 2^{97}, \quad T_2 \approx 2^{89}$$

(again, we will not be discussing the operations since they were already accounted for). Using $\mathbb{F}_{2^{210}}/\mathbb{F}_{2^{35}}$, we expect Enge-Gaudry to perform in time

$$T_1 \approx 2^{90}, \quad T_2 \approx 2^{72}.$$

For this extension, approximately 2^{175} isomorphism classes of curves are susceptible in this case, though one can extend the range via isogenies to roughly one quarter of all curves (those with $\text{Tr}(a) = \text{Tr}(b) = 0$). The isogeny walk is expected to be faster than the above estimate, so the overall running time is dominated by Enge-Gaudry. It is clear that for both of these extensions the running time is better than Pollard's rho, however as in our discussion for $\mathbb{F}_{2^{150}}$, the linear-algebra phase is non-parallelizable and has huge memory requirements, which may prove to be highly restrictive in practice.

In [24], the effects of the generalized version of GHS for this field were also analyzed. Introduced in 2004 by Hess [16], the generalized GHS (gGHS) attack differs from the original version in that the covering curve obtained is not necessarily hyperelliptic. Under some conservative assumptions regarding the time-cost of operating in the Jacobian of a non-hyperelliptic curve, the following results were obtained.

When $n = 7$, the gGHS is effective for around 2^{149} curves with $\text{Tr}(a) = 0$. The running time estimates for index-calculus were

$$T_1 \approx 2^{84}, \quad T_2 \approx 2^{61}.$$

One can also use isogeny walks to extend this to half of all non-subfield curves (those with $\text{Tr}(a) = 0$). The isogeny walk is dominated by the above estimates, so it does not affect the overall running time. When $n = 6$, the estimates are

$$T_1 \approx 2^{82}, \quad T_2 \approx 2^{72}$$

for curves which yield a genus 14 cover (which ends up being almost all curves). Alternatively, a small portion (2^{176} isomorphism classes) yield a genus 12 cover. In this instance, the estimates are

$$T_1 \approx 2^{72}, \quad T_2 \approx 2^{72}.$$

This can be extended to half of all curves via isogeny walks to all non-subfield curves E such that $\#E(\mathbb{F}_{2^{210}}) \equiv 0$ or $2 \pmod{8}$. Again, the isogeny walk is dominated by the index-calculus costs, so it does not impact the running time. It should be mentioned that as before, the relation generation time (T_1) is parallelizable, but no effective methods are known for the linear algebra (T_2). Also, while these times are indeed very good, the generalized GHS has yet to be implemented, and there is some uncertainty regarding the actual cost of operating in the Jacobian of a non-hyperelliptic curve.

Finally, we consider the speed of Joux-Vitse’s algorithm using the tower $\mathbb{F}_{2^{210}} - \mathbb{F}_{2^{70}} - \mathbb{F}_{2^{35}}$. Their attack targets roughly 2^{141} elliptic curves over $\mathbb{F}_{2^{210}}$, and with their sieving variant of Nagao decomposition, is estimated to take time $2^{35(12/7)} \approx 2^{60}$, though hidden constants must be kept in mind. Note that we could also consider applying the double-large prime variation here instead, but the running time estimate is 2^{93} operations, which is worse than the GHS employed with the extensions discussed above. Using an argument similar to that of Section 6.2, the isogeny walk will require roughly 2^{70} steps, and allows the attack to reach close to half of all curves (those with cardinality congruent to 0 or 2 modulo 8). Hence this approach seems to outperform even the gGHS attacks (as the walk may be parallelized), and has been implemented before, which gives us a good understanding of its performance in practical settings.

We performed timings in this field as well, finding that a single step of Pollard’s rho takes roughly 5.25×10^{-5} seconds, and a single walk step takes time roughly 0.137 seconds. The ratio of these times is $2610 \approx 2^{11}$. Given that Pollard’s rho takes 2^{105} steps and the random walk takes 2^{70} steps, the latter approach is asymptotically faster. We conclude that the combination of Joux-Vitse’s attack and isogeny walks adds further evidence for the potential badness of $\mathbb{F}_{2^{210}}$ for elliptic curve cryptography.

Chapter 7

Concluding Remarks

Based on our experimental results, we have evidence pointing towards the validity of the assumption that elliptic curves distribute roughly uniformly amongst isogeny classes, and that those susceptible to Weil descent are indeed distributed in a roughly uniform fashion. In particular, this tells us that as fields $\mathbb{F}_{2^{6t}}$ reach cryptographic size, we can expect that a given ordinary isogeny class will contain a vulnerable curve (assuming that the orders of the curves within are congruent to either 0 or 2 modulo 8). An aspect that remains unexplored in this work is the distribution of weak curves when there are multiple layers in the isogeny volcano. An assumption made in [24] is that the ratio of weak curves in an entire isogeny class is the same if one restricts to curves which have the same endomorphism ring, and it would be worthwhile to gather experimental data for this.

While we have shown the partial weakness of the field $\mathbb{F}_{2^{150}}$, it remains to actually implement the cover and decomposition attack on a cryptographically interesting elliptic curve over this field and to compute discrete logarithms. It would be reasonable to expect that using the tower $\mathbb{F}_{2^{150}} - \mathbb{F}_{2^{50}} - \mathbb{F}_{2^{25}}$ the running time would be similar to the 149-bit example done by Joux and Vitse, namely around 110 000 CPU hours. Furthermore, a more refined implementation of the random walk might give the stronger result of partial badness. In addition, an implementation of the generalised GHS attack for these fields would also be worthwhile, so that a more concrete comparison could be drawn between it and the Joux-Vitse attack.

Another potential avenue for future work lies in the case of odd characteristic fields. The work done in [19] takes place in non-binary fields, and according to the authors of [19], “it may be possible to transfer the DLP from [an elliptic curve] E to a more vulnerable isogenous curve E' .” Presumably one would perform an isogeny walk on elliptic curves expressed in Weierstrass form. However in the degree 6 extension setting, the curves for

which the GHS attack succeeds are those with the form

$$y^2 = h(x)(x - \alpha)(x - \sigma(\alpha)),$$

where $h(x)$ is a degree 1 or 2 polynomial in $\mathbb{F}_q[x]$, $\alpha \in \mathbb{F}_{q^3} \setminus \mathbb{F}_q$, and σ is the Frobenius automorphism of $\mathbb{F}_{q^3}/\mathbb{F}_q$ [19]. It seems unclear how one can efficiently check whether a particular curve in Weierstrass form can be transformed into the above, since the relationship between the roots is not preserved by affine transformations in \mathbb{F}_{q^3} . Thus an effective method for discerning whether a curve is weak will be needed in order to utilize the isogeny walk in a meaningful way.

References

- [1] E. Bach. Explicit bounds for primality testing and related problems. *Math. Comp*, 55:355–380, 1990.
- [2] D. Cox. *Primes of the Form $x^2 + ny^2$* . John Wiley & Sons Inc, 2013.
- [3] C. Diem. The GHS attack in odd characteristic. *J. Ramanujan Math. Soc.*, 18:1–32, 2003.
- [4] C. Diem. On the discrete logarithm problem in elliptic curves. *Compos. Math.*, 147(1):75–104, 2011.
- [5] C. Diem and E. Thomè. Index calculus in class groups of non-hyperelliptic curves of genus three. *J. Cryptology*, 21:593–611, 2008.
- [6] A. Enge and P. Gaudry. A general framework for subexponential discrete logarithm algorithms. *Acta Arithmetica*, 102:83–103, 2002.
- [7] G. Frey. Applications of arithmetical geometry to cryptographic constructions. *Proceedings of the Fifth International Conference on Finite Fields and Applications*, pages 128–161, 2001.
- [8] G. Frey and H. Rück. A remark concerning m -divisibility and the discrete logarithm in the divisor class group of curves. *Mathematics of Computation*, 62:865–874, 1994.
- [9] W. Fulton. *Algebraic Curves – An Introduction to Algebraic Geometry*. 2008.
- [10] S. Galbraith. *Mathematics of Public Key Cryptography*. Cambridge University Press, 2012.
- [11] S. Galbraith, F. Hess, and N.P. Smart. Extending the GHS Weil descent attack. *Advances in Cryptology – EUROCRYPT 2002, LNCS*, 1807:29–44, 2002.

- [12] P. Gaudry. Index calculus for abelian varieties of small dimension and the elliptic curve discrete logarithm problem. *J. Symbolic Comput.*, 44(12):1690–1702, 2008.
- [13] P. Gaudry, F. Hess, and N.P. Smart. Constructive and destructive facets of Weil descent on elliptic curves. *J. Cryptology*, 15(1):19–46, 2002.
- [14] P. Gaudry, E. Thomé, N. Thériault, and C. Diem. A double large prime variation for small genus hyperelliptic index calculus. *Math. Comp.*, 76(257):475–492, 2007.
- [15] D. Hankerson, A. Menezes, and S. Vanstone. *Guide to Elliptic Curve Cryptography*. Springer-Verlag, New York, 2003.
- [16] F. Hess. Generalising the GHS attack on the elliptic curve discrete logarithm problem. *LMS Journal of Computation and Mathematics*, 7:167–192, 2004.
- [17] M. Jacobson, A. Menezes, and A. Stein. Solving elliptic curve discrete logarithm problems using Weil descent. *J. Ramanujan Math. Soc.*, 16(3):231–260, 2001.
- [18] D. Jao, S. Miller, and R. Venkatesan. Do all elliptic curves of the same order have the same difficulty of discrete log? *Advances in Cryptology – Asiacrypt 2005, LNCS*, 3788:21–40, 2005.
- [19] A. Joux and V. Vitse. Cover and decomposition index calculus on elliptic curves made practical – Application to a previously unreachable curve over \mathbb{F}_{p^6} . *In: Advances in Cryptology – EUROCRYPT 2012, LNCS*, 7237:9–26, 2012.
- [20] D. Kohel. *Endomorphism rings of elliptic curves over finite fields*. PhD thesis, University of California at Berkeley, 1996.
- [21] H.W. Lenstra. Factoring integers with elliptic curves. *Annals of Mathematics*, 126:649–673, 1987.
- [22] A. Menezes, T. Okamoto, and S. Vanstone. Reducing elliptic curve logarithms to logarithms in a finite field. *IEEE Transactions on Information Theory*, 39:1639–1646, 1993.
- [23] A. Menezes and M. Qu. Analysis of the Weil descent attack of Gaudry, Hess and Smart. *Topics in Cryptology: CT-RSA, LNCS*, 2020:308–318, 2001.
- [24] A. Menezes and E. Teske. Cryptographic implications of Hess’ generalized GHS attack. *Applicable Algebra in Engineering, Communication and Computing*, 15(6):439–460, 2006.

- [25] A. Menezes, E. Teske, and A. Weng. Weak fields for ECC. *Topics in Cryptology: CT-RSA 2004, LNCS*, 2964:366–386, 2004.
- [26] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, Cambridge, 1995.
- [27] M. Musson. *Another Look at the Gaudry Hess and Smart Attack on the Elliptic Curve Discrete Logarithm Problem*. PhD thesis, University of Calgary, 2011.
- [28] K. Nagao. Decomposition attack for the Jacobian of a hyperelliptic curve over an extension field. *ANTS-IX: Algorithmic Number Theory, LNCS*, 6197:285–300, 2010.
- [29] S. Pohlig and M. Hellman. An improved algorithm for computing logarithms over $\text{GF}(p)$ and its cryptographic significance. *IEEE Transactions on Information Theory*, 24:106–110, 1978.
- [30] J. Pollard. Monte Carlo methods for index computation mod p . *Mathematics of Computation*, 32:918–924, 1978.
- [31] T. Satoh and K. Araki. Fermat quotients and the polynomial time discrete log algorithm for anomalous elliptic curves. *Commentarii Mathematici Universitatis Sancti Pauli*, 47:81–92, 1998.
- [32] R. Schoof. Nonsingular plane cubic curves over finite fields. *J. Combinatorial Theory*, 46:183–211, 1987.
- [33] I. Semaev. Evaluation of discrete logarithms in a group of p -torsion points of an elliptic curve in characteristic p . *Mathematics of Computation*, 67:353–356, 1998.
- [34] I. Semaev. Summation polynomials and the discrete logarithm problem on elliptic curves. Cryptology ePrint Archive, Report 2004/31, 2004.
- [35] N. Smart. The discrete logarithm problem on elliptic curves of trace one. *Journal of Cryptology*, 12:193–196, 1999.
- [36] A. Sutherland. Isogeny volcanoes. *ANTS X: Proceedings of the Tenth Algorithmic Number Theory Symposium*, 12:507–530, 2012.
- [37] A. Sutherland. Lecture notes for 18.783: Elliptic Curves. <http://ocw.mit.edu/courses/mathematics/18-783-elliptic-curves-spring-2015/lecture-notes/>, Spring 2015.
- [38] E. Teske. On random walks for pollard’s rho method. *Math. Comp.*, 70:809–825, 2000.

- [39] P. van Oorschot and M. Wiener. Parallel collision search with cryptanalytic applications. *Journal of Cryptology*, 12:1–28, 1999.
- [40] L. Washington. *Elliptic Curves – Number Theory and Cryptography*. CRC Press, Boca Raton, Florida, 2008.