Explicit Descent On Elliptic Curves Over Function Fields

by David Roberts, MMath

Thesis submitted to the University of Nottingham for the degree of Doctor of Philosophy, September 2007

Abstract

In this thesis we develop and present algorithms for two-descent and descent via two-isogeny on elliptic curves over function fields of the form $\mathbb{F}_q(T)$. The approach taken requires several sub-algorithms to be developed in the function field case: for solving genus zero curves, local solvability testing of homogeneous spaces and searching for points on everywhere-locally-solvable homogeneous spaces. The latter requires the notion of basis reduction for polynomial lattices, building on similar approaches taken over $\mathbb{Q}$ using LLL basis reduction. The algorithms presented here have all been implemented in a computer algebra package, MAGMA.

Contents

0.1 Introduction

A well known problem in number theory is to develop ways of calculating the rank, Mordell-Weil group or finding points on elliptic curves. Two such methods for doing this are two-descent and descent via two-isogeny. These methods are already very well developed for elliptic curves over the rational numbers and much has also been accomplished for when the base-field is a number field.

But there is another direction in which the descent methods can be extended which has seen much less work than other cases: when the base field of an elliptic curve is a function field. It is the purpose of this thesis to establish a basis for further study in this area. The algorithms for two-descent and descent via two-isogeny require many sub-algorithms during the course of their application, and these must therefore also be developed in this new setting.

There are subtle differences in the theory behind these many algorithms. Some (e.g. deciding local-solvability of homogeneous spaces) require not too much extra work to adapt them. Others however (e.g. lattice basis reduction or conic solving) need new techniques and theory to be developed.

Once all the "smaller" algorithms that form the building blocks of the descent algorithms are established, we can put them all together and use them to find the rank and Mordell-Weil groups of elliptic curves over function fields, and this, as will become clear from the content of this thesis, is what has been achieved with explicit worked examples given in many cases throughout.

The algorithms formulated and presented in this thesis have all been implemented in the computer algebra package MAGMA [Mag] (by the author in most cases) and it is hoped that these programs will be included in future releases of the software.

Chapter 1

Background on function fields

1.1 Definitions

A *rational function field* in one variable (univariate) over a field F is the field of fractions of the polynomial ring $F[X]$. Let K denote this function field. Elements of K are rational-functions in X over the base field F. A typical element $f = f(X) \in K$ has the form $f(X) = P(X)/Q(X)$ where P and Q are polynomials in X with coefficients in F (notation: $K = F(X)$). For the rest of this thesis, unless otherwise specified, "function field" will be taken to mean "rational function field". An element $P(X)$ of $F[X]$ is written uniquely in the form $P(X) = a_0 + a_1X + a_2X^2 + \cdots + a_dX^d$ where $a_i \in F$ and $a_d \neq 0$. The positive integer d is called the *degree* of P and a_d the *leading coefficient* of P. If the leading coefficient of P is equal to 1 (in F) then P is said to be *monic*.

Function fields of more than one variable are the fields of fractions of multivariate polynomial rings (denoted $K = F(X_1, X_2, \ldots, X_n)$), however they may be thought of as univariate function fields in the following way: an element $f \in F(X_1, \ldots, X_n)$ has the form $P(X_1, \ldots, X_n)/Q(X_1, \ldots, X_n)$. The polynomials P and Q can be written in the form $a_0 + a_1X_n + a_2X_n^2 + \cdots + a_kX_n^k$ where $a_i \in F(X_1, \ldots, X_{n-1})$ for $0 \leq i \leq k$. Conversely a rational function in X_n with coefficients in $F(X_1, \ldots, X_{n-1})$ can clearly be written as an element of $F(X_1, X_2, \ldots, X_n)$.

We therefore can conclude the following useful fact:

Let F be a field. Let $X_1, \ldots, X_n$ be all the variables of a polynomial ring over F. Denote by K_i $(1 \leq i \leq n)$ the function field $F(X_1, \ldots, X_i)$. Then if $2 \leq k \leq n$:

$$K_k = K_{k-1}(X_k). \tag{1.1}$$

This is useful from a programming perspective: if there exists an algorithm that provides the solution to problems that are intrinsically linked to fields (e.g. finding points on curves over a field), then to generalize to multivariate function fields the problem reduces to finding a solution in the univariate case that is dependent on there being a solution over the base field. If this is the case then problems over $F(X_1, \ldots, X_n)$ can be reduced to problems over F, for which there are often easy or already well-established solutions. An example of this is the conic solving algorithm (see chapter 5) in which the solution is generalized inductively to the multivariate case after reducing the problem of finding points over $F(X)$ to finding points over F.

1.2 Primes and factorization in $F[T]$

When working with curves etc over univariate function fields, the polynomial ring $F[T]$ plays a similar rôle to the integers in the rational case, and many useful properties of $\mathbb{Z}$ have analogues in $F[T]$ (see [Ros]); one of the most notable is the existence of prime elements and unique factorization. Unique factorization follows from $F[T]$ being a Euclidean domain (and thus a principal ideal domain).

The unit group of $F[T]$ is simply F^*. In $\mathbb{Z}$, the unit group is $\{-1, 1\}$ and its finiteness is important in many algorithms. In order to be able to formulate equivalent or analogous algorithms in the function field setting, we require in many cases that $F[T]$ has a finite unit group. This means that F needs to be a finite field. We therefore will see that a closer analogue to $\mathbb{Z}$ is $\mathbb{F}_q[T]$ where $\mathbb{F}_q$ denotes the finite field of $q = p^n$ (p prime) elements.

Before we can say that $F[T]$ is a unique-factorization-domain we must first define what is meant by a prime element.

Definition. A prime element in $F[T]$ is a monic irreducible polynomial of positive degree.

By irreducible we mean that a polynomial cannot be factorized into two or more polynomials of smaller (positive) degree.

Choosing monic polynomials in the definition is for convenience - a prime is actually unique up to multiplication by any element of the unit group F^* (in the same way as 2 and -2 represent the same prime in $\mathbb{Z}$). With the above definition of a prime element we now conclude the following:

Every non-zero element P of $F[T]$ can be written uniquely in the form

$$P = \alpha P_1^{e_1} P_2^{e_2} \cdots P_k^{e_k} \tag{1.2}$$

where $\alpha \in F^*$, P_i is monic and irreducible for all i, $P_i \neq P_j$ for $i \neq j$ and each e_i is a positive integer and $k \geq 0$. More generally, if $f \in F(T)$, then $f = \alpha P_1^{e_1}, \ldots, P_k^{e_k} \quad (e_i \in \mathbb{Z})$.

1.3 Valuations on $F(T)$

1.3.1 The ord_p function

For every prime p in $F(T)$ we may define a function $\mathrm{ord}_p : F(T)^* \longrightarrow \mathbb{Z}$ in the following way:

We begin by writing a non-zero element f of $F(T)$ in the form

$$f = \alpha \cdot \frac{P_1^{a_1} P_2^{a_2} \cdots P_r^{a_r}}{Q_1^{b_1} Q_2^{b_2} \cdots Q_s^{b_s}} \tag{1.3}$$

where $\alpha \in F^*$ and the P_i and Q_i are all distinct prime elements of $F(T)$. Then:

$$\mathrm{ord}_p(f) = \begin{cases} a_i & \text{if } p = P_i \\ -b_i & \text{if } p = Q_i \\ 0 & \text{otherwise} \end{cases} \tag{1.4}$$

ord_p has the following properties:

1) $\mathrm{ord}_p(f) = 0$ for all $f \in F^*$

2) $\mathrm{ord}_p(fg) = \mathrm{ord}_p(f) \cdot \mathrm{ord}_p(g)$

3) $\mathrm{ord}_p(f+g) \geq \min\{\mathrm{ord}_p(f), \mathrm{ord}_p(g)\}$ with equality if $\mathrm{ord}_p(f) \neq \mathrm{ord}_p(g)$.

The domain of ord_p can be extended to include zero if we introduce the symbol ∞ with the property $x < \infty$ for all $x \in \mathbb{Z}$. Then defining $\mathrm{ord}_p(0) = \infty$, the function now satisfies the above properties over all of $F(T)$.

1.3.2 The ord_∞ function

We can define another important valuation on $F(T)$ which is not associated to any of the primes in $F[T]$. This valuation indicates how a function $f \in F(T)$ behaves "at infinity" (regarding f as a function on $\mathbb{P}^1(f)$) we can define the valuation, which we call ord_∞ as follows:

Definition.

$$\mathrm{ord}_\infty(P(T)/Q(T)) = \deg(Q) - \deg(P) \tag{1.5}$$

where P and Q are polynomials in $F[T]$.

It is a simple exercise to show that ord_∞ satisfies the same conditions 1-3 above that are satisfied also by ord_p.

It can be shown that calculating $\mathrm{ord}_\infty(f(T))$ is the same as looking at how $f(1/T)$ behaves at zero, i.e $\mathrm{ord}_T(f(1/T))$:

Proposition 1. *Let $f(T)$ be an element of the function field $F(T)$. Then*

$$\mathrm{ord}_\infty(f(T)) = \mathrm{ord}_T(f(1/T)).$$

Proof. First we make the observation that for any non-zero polynomial $P(T)$ of degree d, the polynomial $T^d P(1/T)$ has non-zero constant term, since

$$P(T) = a_0 + a_1 T + \cdots + a_d T^d \Rightarrow T^d P(1/T) = a_d + a_{d-1} T + \cdots + a_0 T^d$$

i.e. evaluating P at $1/T$ and multiplying by $T^{\deg(P)}$ simply reverses the coefficients. From this we see that $\mathrm{ord}_T(T^d P(1/T)) = 0$.

Any $f \in F(T)^*$ can be written as $f(T) = P(T)/Q(T)$ where P and Q are non-zero polynomials. It follows that

$$f(1/T) = T^{\deg(Q)-\deg(P)} \cdot \frac{P^*(T)}{Q^*(T)}$$

where P^* and Q^* are the reverse-coefficient polynomials of P and Q respectively. Hence

$$\mathrm{ord}_T(f(1/T)) = (\deg(Q) - \deg(P)) + \mathrm{ord}_T(P^*) - \mathrm{ord}_T(Q^*).$$

Therefore $\mathrm{ord}_T(f(1/T)) = (\deg(Q) - \deg(P)) = \mathrm{ord}_\infty(f)$. □

1.4 Completions of $\mathbb{F}_q(T)$

In subsequent chapters it will become necessary to determine whether or not certain curves over $\mathbb{F}_q(T)$ contain any points. One method of showing that no points exist on a particular curve is to prove that it is not locally solvable - that is, it has no points over the completion of $\mathbb{F}_q(T)$ at a valuation $v = \mathrm{ord}_p$ or $v = \mathrm{ord}_\infty$. Some explanation on the completions of $\mathbb{F}_q(T)$ and associated maps is therefore necessary.

Notation: $K = \mathbb{F}_q(T)$, $R = \mathbb{F}_q[T]$.

Definition. The *residue degree* of a valuation v on K is $\deg(p)$ if $v = \mathrm{ord}_p$ for a monic irreducible p. If $v = \mathrm{ord}_\infty$ then the residue degree of v is equal to 1.

If a valuation has residue degree d, then the corresponding residue field is $\mathbb{F}_{q^d}$.

Proposition 2. *Let v be a valuation on K with residue degree d. The completion of K at v is isomorphic to $L = \mathbb{F}_{q^d}((U))$. That is the formal Laurent series field in one variable, U over $\mathbb{F}_{q^d}$. The embedding $i : K \hookrightarrow L$ is given as follows:*

- *if $x \in \mathbb{F}_q$ then $i(x) = x$ (considering $\mathbb{F}_q$ as a subfield of $\mathbb{F}_{q^d}$).*
- *if $v = \mathrm{ord}_p$ for a monic irreducible $p(T)$ then $i(T) = U + \alpha$ where $\alpha \in \mathbb{F}_{q^d}$ is a root of p.*

- *if* $v = \mathrm{ord}_\infty$ *then* $i(T) = U^{-1}$.

Moreover, when $d > 1$, L *is also the completion of* $\mathbb{F}_{q^d}(T)$ *at the valuation* $\mathrm{ord}_{T-\alpha}$.

Proof. Let L be the completion of K. L then has the following properties:

- $K \subseteq L$ and K is dense in L
- L is complete with respect to a discrete valuation v, extending that on K
- L has valuation ring $\mathcal{O}_L$ and maximal ideal $\mathfrak{m}_L$
- the residue field $\mathcal{O}_L/\mathfrak{m}_L$ is $\mathbb{F}_{q^d}$ (completion does not change the residue field).

Consider the case $d = 1$. If $v = \mathrm{ord}_T$ then the completion of K at v is clearly isomorphic to $\mathbb{F}_q((U))$ with the obvious embedding $T \mapsto U$. If $v = \mathrm{ord}_{T-\alpha}$ for some $\alpha \in \mathbb{F}_q$ then the completion is still $\mathbb{F}_q((U))$, this time with the embedding $T \mapsto U + \alpha$ so $T - \alpha$ maps to U.

If $v = \mathrm{ord}_\infty$ then we reduce to the previous case after applying the map $\mathbb{F}_q(T) \to \mathbb{F}_q(U)$ via $T \mapsto U^{-1}$.

Now consider the case $v = \mathrm{ord}_p$ for some monic irreducible polynomial $p \in \mathbb{F}_q[T]$ of degree d. We can regard K as a subset of L via its embedding in its completion, so $v(T) \geq 0 \Rightarrow T \in \mathcal{O}_L$. T then reduces to an element α of $\mathcal{O}_L/\mathfrak{m}_L = \mathbb{F}_{q^d}$. Hence $p(T)$ maps to $p(\alpha)$, but since $v(p(T)) > 0$, we have $p(\alpha) = 0$ in $\mathbb{F}_{q^d}$. Thus L contains a root α of $p(T)$. Note also that $\mathbb{F}_{q^d} = \mathbb{F}_q(\alpha)$ so:

$$K = \mathbb{F}_q(T) \subseteq \mathbb{F}_{q^d}(T) \subseteq L. \tag{1.6}$$

The restriction of v to $\mathbb{F}_{q^d}(T)$ is a valuation of that field in which $v(T-\alpha) > 0$. This can only be $\mathrm{ord}_{T-\alpha}$. From the definition of completeness we see that L is

also the completion of $\mathbb{F}_{q^d}(T)$ with respect to $\mathrm{ord}_{T-\alpha}$. We are now in the case $d = 1$ and the result is proved.

□

Corollary 3. *Let $\mathcal{C}$ be a curve over $K = \mathbb{F}_q(T)$ and $v = \mathrm{ord}_p$ be a valuation on K of residue degree d. Let $\mathcal{C}'$ be the same curve over $K' = \mathbb{F}_{q^d}(T)$ (after extending the base-field). Let α be a root of p in $\mathbb{F}_{q^d}$. Let v' be the valuation $\mathrm{ord}_{T-\alpha}$ on K'. Then $\mathcal{C}$ is locally solvable at v if and only if $\mathcal{C}'$ is locally solvable at v'.*

Proof. K and K' both have common completion L at v and v' respectively, therefore both statements mean the same thing, namely that $\mathcal{C}(L) \neq \varnothing$. □

1.4.1 Application to local-solvability

The above corollary is useful when implementing an algorithm to determine certain local properties of a curve over a function field, in particular whether it contains any points after a base change to the completion of its base field (local-solvability). Instead of having to work with higher degree primes p we can simply extend the "constant" field $\mathbb{F}_q$ to include a root α of p and work over the residue field $\mathbb{F}_{q^{\deg(p)}}$, with reduction mod p now equivalent to evaluating a polynomial $f(T)$ at $f(\alpha)$ as one does for the degree 1 case. In Magma we therefore avoid having to create more complicated algebraic structures than is necessary.

Chapter 2

Models of curves of genus 1

Throughout this thesis, use is made of several different models for curves of genus one. I will not here give a definition of the genus of a curve as it will not be used at any point, but rather will give a brief description of the various curves and models used.

Throughout the rest of the following chapter we will assume that the characteristic of the field K is not 2.

2.1 Elliptic Curves and Weierstrass equations

Definition. An *elliptic curve* over a field K is a projective curve of genus one that is smooth over K and contains a K-rational point $\mathcal{O}$.

More correctly the elliptic curve is defined by the pair $(E, \mathcal{O})$ where E is the curve and $\mathcal{O}$ the specified point. It is a well known fact that the K-points of an elliptic curve form a finitely generated abelian group with $\mathcal{O}$ as the identity (the *Mordell-Weil group* of E).

The most familiar model of an elliptic curve is its embedding into $P^2(K)$

where it is given by a *Weierstrass equation*. This model takes the following form:

$$E : Y^2 + a_1XY + a_3Y = X^3 + a_2X^2 + a_4X + a_6 \tag{2.1}$$

where $a_1, a_2, a_3, a_4, a_6 \in K$ - the so-called *a-invariants* of E. In fact, since $\text{char}(K) \neq 2$ we can assume that $a_1 = a_3 = 0$. Upon homogenization of the Weierstrass equation with a third variable Z, it is clear that the projective curve E always has a point in $\mathbb{P}^2(K)$, where $X = Z = 0$ and $Y = 1$ regardless of what K is. This point $(0 : 1 : 0)$ is nearly always taken to be the identity $\mathcal{O}$ in the Mordell-Weil group of E. If, in addition, K does not have characteristic 3 then the equation of the curve can be given in the simpler form:

$$E : Y^2 = X^3 + a_4X + a_6. \tag{2.2}$$

The set of K-rational points of E is written as $E(K)$. This notation is also used to represent the Mordell-Weil group of E.

We will also come across curves of genus one that do not have a (known) rational point. These are given by other models.

2.2 $Y^2 = quartic$

Similar in form to the Weierstrass equation of an elliptic curve are equations of the form $Y^2 = f(X)$ where f is a polynomial over a field K with degree 4. Provided f has no repeated roots, these define curves of genus 1 over K. A curve of genus 1 has such a model if and only if it has a K-rational divisor of degree 2. The quartic f can also be written as a polynomial in one variable $[u : v]$ in $\mathbb{P}^1$, i.e. $f(U, V)$ with f homogeneous of degree 4 in U and V. The equation $Y^2 = f(U, V)$ then defines a curve in weighted-projective space with

Y, U and V having weights 2, 1 and 1 respectively.

The weighted version of the curve occurs during descent on elliptic curves when deciding solvability of quadric intersections. This will be discussed further in chapter 6.

2.2.1 I, J and Δ

A quartic of the form $aX^4 + bX^3 + cX^2 + dX + e$ over a field K has associated to it elements of K known as the I and J invariants, and $\Delta = 4I^3 - J^2$, which is associated to the discriminant (see [Cr1, p.89]). We can use the definitions of I and J to extend this to homogeneous quartics and curves of the form $Y^2 = quartic$ in the obvious way:

Definition. Let f be a homogeneous quartic in U and V over K so $f = aU^4 + bU^3V + cU^2V^2 + dUV^3 + eV^4$. Then the I and J invariants of f are defined as follows:

$$\begin{aligned} I = &\quad 12ae - 3bd + c^2, \\ J = &\quad 72ace + 9bcd - 27ad^2 - 27b^2e - 2c^3. \end{aligned} \tag{2.3}$$

2.3 Intersections of quadric surfaces

Another model of a genus 1 curves occurs if and only if a curve has a K-rational divisor of degree 4. If this is the case, the curve can be embedded into 3-dimensional projective space as an intersection of two quadric surfaces.

A quadric surface in $\mathbb{P}^3$ over a field K is a surface defined by a single homogeneous quadratic equation:

$$\sum_{1 \le i,j \le 4} a_{ij} X_i X_j = 0 \qquad (a_{ij} \in K). \tag{2.4}$$

Alternatively this can be written as $\mathbf{X}^T M\mathbf{X} = 0$ where $\mathbf{X}$ is the column-vector of the variables $X_1, \ldots, X_4$ and M is the 4×4 symmetric matrix (a_{ij}) with $a_{ij} = a_{ji}$. The quadric is said to be singular if its defining matrix is singular.

An intersection of two such quadric surfaces (provided the resulting curve is smooth) defines a genus one curve in $\mathbb{P}^3(K)$. A quadric intersection $\mathcal{C}$ given by two symmetric matrices A and B is smooth if and only if $\text{disc}(\det(uA - vB)) \neq 0$.

2.3.1 Pencil of quadrics

Any two quadrics over K, given by matrices A, B defines what is called a pencil of quadrics in $\mathbb{P}^3(K)$, provided $A \neq cB$ for some $c \in K$. In terms of the matrices A and B, what this means is the following: if A and B are the matrices defining a quadric-intersection $\mathcal{C}$, then the pencil of quadrics associated to $\mathcal{C}$ is given by the matrix $M(\lambda, \mu) = \lambda A - \mu B$. Evaluating M at a point (λ, μ) in $\mathbb{P}^1(K)$ gives a quadric surface over K. It is clear that any distinct pair of quadrics in the pencil define the same pencil, and hence the same curve $\mathcal{C}$. We write (A, B) to mean the pencil defined by the matrices A and B.

If $\mathcal{C}$ is smooth and A is non-singular, then the determinant of the matrix defining the pencil has four distinct roots and can be written as follows:

$$\det(uA + vB) = a(u - \epsilon_1 v)(u - \epsilon_2 v)(u - \epsilon_3 v)(u - \epsilon_4 v) \tag{2.5}$$

where $a, \epsilon_i \in K, \quad a \neq 0$. Therefore there are four singular quadrics in the pencil given by the four roots $[\epsilon_i : 1]$ (not all necessarily defined over K) and their matrices $\epsilon_i A + B$ all have rank 3. This then allows us to change variables and write the equations defining the singular quadrics in terms of exactly 3 of the 4 variables. The redundant variable will be different for each of the singular quadrics in the pencil. In the case that A is singular, one of the four factors in

the above equation degenerates to v, corresponding to the root $[1:0]$.

Example. Let C be the quadric intersection over $\mathbb{F}_{11}$ given by the following equations:

$$\begin{aligned} {x_1}^2 + 2{x_2}^2 &= 5{x_4}^2 \\ 2{x_1}^2 + {x_3}^2 &= 2{x_4}^2 \end{aligned} \tag{2.6}$$

C is also represented by the pair of matrices

$$A = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \end{pmatrix} \text{ and } B = \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 9 \end{pmatrix}. \tag{2.7}$$

The pencil of quadrics defined by C is represented by the matrix $M(\lambda, \mu) = \lambda A - \mu B$. This has determinant $\det(M) = 10\lambda\mu(\lambda - 7\mu)(\lambda - 2\mu)$ so the four singular quadrics occur at $(\lambda, \mu) = (1:0), (0:1), (2:1), (7:1)$. The two quadrics we used to define C come from the roots $(1:0)$ and $(0:1)$. The other two are:

$$\begin{aligned} 4{x_2}^2 + 10{x_3}^2 + 3{x_4}^2 &= 0 \\ 5{x_1}^2 + 3{x_2}^2 + 10{x_3}^2 &= 0 \end{aligned} \tag{2.8}$$

2.4 QI - quartic relationships

Quadric intersections and curves of the from $Y^2 = \textit{quartic}$ are closely related. Each quadric intersection has an associated quartic $f(x) = \det(Ax - B)$ where A and B are the defining matrices. The quartic may also be given in homogeneous form, in which case it is $\det(uA - vB)$.

To transform $Y^2 = aX^4 + bX^3 + cX^2 + dX + e$ into a quadric intersection we make the substitution $Z = X^2$ (1). Then the equation becomes $Y^2 =$

$aZ^2 + bXZ + cZ + dX + e$ (2). Homogenizing both (1) and (2) with a fourth variable W, we obtain

$$\begin{cases} X^2 - ZW = 0 \\ aZ^2 + bXZ + cZW + dXW - Y^2 + eW^2 = 0. \end{cases} \tag{2.9}$$

Note that the first quadric above is singular - it does not involve the variable Y and its associated matrix has rank 3 rather than 4. Therefore the associated quartic $\det(uA - vB)$ has a K-rational root $(u : v) = (1 : 0)$.

Conversely if a pencil of quadrics (A, B) has a K-rational root in the associated quartic, then we may change the basis of the pencil to make this root $(1 : 0)$, so $\det(A) = 0$. If in addition the pencil is nonsingular then A must have rank 3, so after a change of variables the equation of the associated quadric only involves 3 of the 4 variables and can be projected down to $\mathbb{P}^2$ as a smooth conic. Furthermore, if this conic has a K-rational point then it may be parametrized, and hence all points on the quadric may be parametrized. Substituting this parametrization into the second quadric and completing the square we recover an equation of the form $Y^2 = f(U, V)$, where f is a homogeneous quartic.

Remark. Note that if K is a number field or function field and the quadric intersection has points everywhere locally then so must the conic obtained from the singular quadric. Then, by the Hasse local-global principal it must have a K-rational point an can therefore be parametrized. So every curve of the form $Y^2 = f(X)$ (f a quartic) can be converted into a quadric intersection (as shown above) but only those QIs with: a) points everywhere locally, and b) A K-rational root of the associated quartic can be "converted back".

2.5 The Hasse bound

For genus 1 curves over finite fields, the following is a well known result:

Theorem 4. *Let $\mathcal{C}$ be a projective curve of genus 1 over the finite field $\mathbb{F}_q$. Then $\#\mathcal{C}(\mathbb{F}_q)$ satisfies the following inequality:*

$$|\#\mathcal{C}(\mathbb{F}_q) - (q+1)| \leq 2\sqrt{q}.$$

This is known as Hasse's theorem and is a standard result (see [Sil1, Ch.V]). A corollary of the above is that a genus 1 curve over a finite field always has at least one point.

Chapter 3

Full 2-descent on elliptic curves

3.1 Introduction

In this chapter we describe an algorithm for performing two descent on an elliptic curve E over $K = \mathbb{F}_q(T)$ when the curve has full 2-torsion. The algorithm will be used to compute the Mordell-Weil group of E (as an abstract group $E(K)_{tors} \oplus \mathbb{Z}^r$). The algorithm requires the following sub-algorithms: conic solving over function fields (see chapter 5), local solvability of quartics over $\mathbb{F}_q(T)$ (chapter 6) and an algorithm for searching for points on quadric-intersections over $\mathbb{F}_q(T)$ (chapter 8).

The input for the algorithm will be an elliptic curve with full two-torsion (defined either by its Weierstrass equation or the x-coordinates of its two-torsion points), and a non-negative integer H which sets a height-bound on the point

search. The algorithm at the very least returns lower and upper bounds on the rank r. Increasing H may raise the returned lower bound on r, however there are curves where the upper bound really is greater than the rank. This occurs when some of the homogeneous spaces, which we describe in the next section, are everywhere locally solvable but do not have a global point (this is the same as saying that E has a non trivial element of order 2 in its Tate-Shafarevich group).

So while the program described here may indeed succeed in finding the rank, for some curves it cannot do so *in principle*. Higher descents would then be needed, which we do not concern ourselves with here.

3.2 Preliminaries

Let E be an elliptic curve over $\mathbb{F}_q(T)$ defined by the following equation:

$$E : Y^2 = (X - e_1)(X - e_2)(X - e_3) \tag{3.1}$$

with $e_1, e_2, e_3 \in \mathbb{F}_q[T]$.

3.2.1 The group $K(S, 2)$

Let K be a field. We let K^*/K^{*2} denote the multiplicative group "K mod squares". Thus if $K = \mathbb{Q}$, elements of $\mathbb{Q}^*/\mathbb{Q}^{*2}$ are represented by sets of distinct prime numbers, together with -1, with multiplication modulo squares as the group law; e.g.

$$2 \cdot 5 \cdot 11 \times 7 \cdot 11 \cdot 31 = 2 \cdot 5 \cdot 7 \cdot 31$$

The identity element is 1 which could be thought of as an "empty" set of primes.

Things are similar when $K = \mathbb{F}_q(T)$ - the primes are now monic irreducible polynomials, and -1 is replaced by a primitive element of $\mathbb{F}_q$. Only one such element α is needed since $\{x^2 : x \in \mathbb{F}_q\} \cup \{\alpha x^2 : x \in \mathbb{F}_q\} = \mathbb{F}_q$. An example of the group law in K^*/K^{*2} when $K = \mathbb{F}_5(T)$ is:

$$2 \cdot T \cdot (T+3) \cdot (T^2+3) \times 2 \cdot T \cdot (T+4) = (T+3) \cdot (T+4) \cdot (T^2+3).$$

We now define an important subgroup of the above group:

Definition. Let K be a field and S a finite set of valuations on K. We denote by $K(S,2)$ the following group:

$$K(S,2) = \left\{ x \in K^*/K^{*2} : v(x) \equiv 0 \mod 2 \quad \forall v \notin S \right\}. \tag{3.2}$$

It is important to note that when $K = \mathbb{F}_q(T)$ the group K^*/K^{*2} also contains a non-square unit α and therefore $K(S,2)$ also contains α. This is only a small problem over $\mathbb{F}_q(T)$ since $\mathbb{F}_q^*/\mathbb{F}_q^{*2}$ is simply $\{1, \alpha\}$ (this is analogous to including -1 in $\mathbb{Q}(S,2)$). Over $\mathbb{Q}(T)$ however, every element of $\mathbb{Q}^*/\mathbb{Q}^{*2}$ is in $\mathbb{Q}(T)(S,2)$ - an infinite group. The following approach to 2 descent relies on the finiteness of $K(S,2)$. Another point to note is that when K is a function field there is the valuation ord_∞. We do *not* require in $K(S,2)$ that $\mathrm{ord}_\infty(x)$ is even, so the definition of $K(S,2)$ should more correctly read:

$$K(S,2) = \left\{ x \in K^*/K^{*2} : v(x) \equiv 0 \mod 2 \quad \forall v \notin S \cup \infty \right\}. \tag{3.3}$$

As an abstract group $K(S,2)$ is isomorphic to $(\mathbb{Z}/2\mathbb{Z})^n$ where $n = 1 + \#S$.

3.2.2 Homogeneous spaces

The existence of points on E depends on the existence of points on a number of curves associated to E. Suppose there exists a point $(X, Y) \in E(K)$; it is

a well known property of elliptic curves over global fields that $X = x/z^2$ and $Y = y/z^3$ where $x, y, z \in \mathbb{F}_q[T]$ and $\gcd(x, z) = \gcd(y, z) = 1$. Replacing X and Y with these values in the equation of E and multiplying by z^6 gives:

$$y^2 = (x - e_1 z^2)(x - e_2 z^2)(x - e_3 z^2). \tag{3.4}$$

Let p be a prime in $\mathbb{F}_q[T]$ (a monic irreducible). Suppose p divides $x - e_i z^2$ and $x - e_j z^2$ $(1 \leq i, j \leq 3,\ i \neq j)$. Then $(x - e_j z^2) - (x - e_i z^2) = (e_i - e_j)z^2$ is divisible by p and so is $e_i(x - e_j z^2) - e_j(x - e_i z^2) = (e_i - e_j)x$. Since $\gcd(x, z) = 1$ it follows that p divides $e_i - e_j$. Therefore we see that if p does not divide $(e_1 - e_2)(e_2 - e_3)(e_3 - e_1)$ then p divides *at most one* of $x - e_i z^2$, so $\mathrm{ord}_p(x - e_i z^2)$ is even (it is either zero or equal to $\mathrm{ord}_p(y^2)$).

We now see that for each i, $x - e_i z^2$ can be written as $d_i z_i^2$ where $z_i, d_i \in \mathbb{F}_q[T]$ and d_i is square-free and divides $(e_1 - e_2)(e_2 - e_3)(e_3 - e_1)$. So we have $x - e_1 z^2 = d_1 z_1^2$, $x - e_2 z^2 = d_2 z_2^2$ and $x - e_3 z^2 = d_3 z_3^2$. Substituting the first two into the equation of E gives $y^2 = d_1 d_2 (z_1 z_2)^2 (x - e_3 z^2)$, and we see that $x - e_3 z^2$ must in fact be equal to $d_1 d_2 z_3^2$ for some $z_3 \in \mathbb{F}_q(T)$. We therefore have the following (after re-naming some of the variables):

$$\begin{cases} (1) & x - e_1 z_4^2 = d_1 z_1^2 \\ (2) & x - e_2 z_4^2 = d_2 z_2^2 \\ (3) & x - e_3 z_4^2 = d_1 d_2 z_3^2. \end{cases} \tag{3.5}$$

Taking the differences (1)-(2) and (1)-(3) of the above, we arrive at:

$$\begin{cases} d_1 {Z_1}^2 & - & d_2 {Z_2}^2 & + & (e_1 - e_2){Z_4}^2 & = & 0 \\ d_1 {Z_1}^2 & - & d_1 d_2 {Z_3}^2 & + & (e_1 - e_3){Z_4}^2 & = & 0. \end{cases} \tag{3.6}$$

These equations define a curve $\mathcal{C}_{d_1,d_2}$ in $\mathbb{P}^3(K)$ for each pair d_1, d_2 square-free dividing $(e_1 - e_2)(e_2 - e_3)(e_3 - e_1)$. This condition on d_1, d_2 is the same as

saying $d_1, d_2 \in K(S,2)$ where S is the set of bad places of E.

A point in $E(K)$ must come from a point on $\mathcal{C}_{d_1,d_2}$ for some d_1, d_2 pair. In fact these curves are even more closely related to E than this as we see in the following two theorems:

Theorem 5. *Let E be an elliptic curve over a global field K of characteristic not 2. Let $E(K)$ denote the Mordell-Weil group of E. If S is the set of places where E has bad reduction (including the place at infinity if K is a function field) then there is an embedding $\phi : E(K)/2E(K) \hookrightarrow K(S,2) \times K(S,2)$ given by*

$$\phi(x,y) = \begin{cases} (x-e_1, x-e_2) & \text{if } x \notin \{e_1, e_2\} \\ ((e_1-e_2)(e_1-e_3), e_1-e_2) & \text{if } x = e_1 \\ (e_2-e_1, (e_2-e_3)(e_2-e_1)) & \text{if } x = e_2 \\ (1,1) & \text{if } x = \infty. \end{cases} \tag{3.7}$$

S is easily computed. It is simply the set of places given by the prime factors of $(e_1-e_2)(e_2-e_3)(e_3-e_1)$. In the function field case, if we did not include ord_∞ in S then we would be further restricted to places coming from primes of even degree.

Theorem 6. *A pair $(d_1, d_2) \in K(S,2) \times K(S,2)$ is an image of an element in $E(K)/2E(K)$ under the above embedding if and only if there is a solution $(z_1, z_2, z_3) \in K$ to the following set of equations:*

$$\begin{cases} d_1 {Z_1}^2 - d_2 {Z_2}^2 &= e_2 - e_1 \\ d_1 {Z_1}^2 - d_1 d_2 {Z_3}^2 &= e_3 - e_1. \end{cases} \tag{3.8}$$

Moreover if a solution (z_1, z_2, z_3) exists then (d_1, d_2) is the image of the point

$$(d_1 {z_1}^2 + e_1, d_1 d_2 z_1 z_2 z_3) \in E(K).$$

If (z_1, z_2, z_3) satisfy (3.8) then $P = \psi(z_1, z_2, z_3) = (d_1 {z_1}^2 + e_1, d_1 d_2 z_1 z_2 z_3)$ lies in $E(K)$ and satisfies $\phi(P) = (d_1, d_2)$.

Proof of these theorems are given in [Sil1, Ch.X, §1] in the case when K is a number field. However there are no assumptions made about K that prevent the proof from working analogously over $\mathbb{F}_q(T)$.

The above pair of equations define a quadric intersection over K (see chapter 2). For each quadric intersection $\mathcal{C}$ arising from a pair (d_1, d_2) we must either find a point on $\mathcal{C}$ or show that no points exist. Such curves have genus one and may be everywhere-locally-solvable yet not globally solvable (the Hasse local-global principle does not apply to genus one curves). It is therefore possible that bad cases will be encountered, in which it cannot be shown that no points exist on $\mathcal{C}$ yet a point may not exist. In such cases we will not be able to compute the rank of E exactly without doing a further descent.

The quadric intersections described above are known as *homogeneous spaces* for E. The next part of the algorithm is to decide whether or not each of them can have points. We use two methods of showing that there are no points: conic solving and local-solvability tests.

Remark

An improvement to the algorithm may be made by observing that not all pairs $(d_1, d_2) \in K(S, 2) \times K(S, 2)$ need to be checked. Assume as in 3.2.2 that $x - e_1 z^2 = d_1 z_1^2$, then if a prime p divides d_1, $\mathrm{ord}_p(x - e_1 z^2)$ is odd. So p must divide at least one of $(x - e_2 z^2), (x - e_3 z^2)$ which implies (using the same argument as in 3.2.2) that p divides $(e_1 - e_2)(e_1 - e_3)$. Similarly if p divides d_2 then p also divides $(e_1 - e_2)(e_2 - e_3)$. Therefore we need only check pairs $(d_1, d_2) \in K(S_1, 2) \times K(S_2, 2)$ where S_1 is the set of places coming from the prime factors

of $(e_1 - e_2)(e_1 - e_3)$. S_2 is similar with the prime factors of $(e_1 - e_2)(e_2 - e_3)$. This immediately reduces the amount of quadric intersections to analyse in subsequent steps. Further refinement may be made by observing that if p divides both d_1 and d_2, then it follows from the above that p must divide $e_1 - e_2$. Also if p divides d_1 but not d_2 (resp. d_2 but not d_1) then p divides $e_1 - e_3$ (resp. $e_2 - e_3$). If, therefore, $d_1 = ab$, $d_2 = ac$ with b, c coprime then $a \mid e_1 - e_2$, $b \mid e_1 - e_3$ and $c \mid e_2 - e_3$.

3.3 The upper bound on the rank

3.3.1 Elimination via conic solving

From now on we work over the field $K = \mathbb{F}_q(T)$. This section requires an algorithm for solving conics over function fields which we give in chapter 5.

Recall that in 3.2.2 we construct for each pair $(d_1, d_2) \in K(S,2) \times K(S,2)$, a quadric intersection over K. Homogenizing the defining equations we obtain

$$\begin{cases} d_1 {Z_1}^2 & - & d_2 {Z_2}^2 & + & (e_1 - e_2){Z_4}^2 & = & 0 \\ d_1 {Z_1}^2 & - & d_1 d_2 {Z_3}^2 & + & (e_1 - e_3){Z_4}^2 & = & 0. \end{cases} \tag{3.9}$$

The above two equations, considered individually involve only three out of the four variables Z_i and are therefore conics in $\mathbb{P}^2(K)$. Denote by $\mathcal{C}_1$ the conic in Z_1, Z_2, Z_4 and $\mathcal{C}_2$ the conic in Z_1, Z_3, Z_4. If the quadric intersection is solvable then so must be both $\mathcal{C}_1$ and C_2. It is therefore a logical first step to check solvability of these using the algorithm described in chapter 5.

In fact we can do a lot better than this. The above equations also define a pencil of quadrics (see 2.3.1) in which the associated quartic has all four roots

in K. In this case the quartic is

$$f(\lambda,\mu) = \lambda \cdot \mu \cdot (\lambda - \mu) \cdot \left(\lambda - \left(\frac{e_1 - e_3}{e_1 - e_2}\right)\mu\right) \cdot g \tag{3.10}$$

where the constant g depends on d_1, d_2 and the e_i. So there are two more singular quadrics in the pencil of which we can check the conic solvability, corresponding to $(\lambda : \mu) = (1 : 1)$ and $(\lambda : \mu) = (e_1 - e_3 : e_1 - e_2)$. These conics are:

$$\begin{aligned} \mathcal{C}_3 : &\quad d_2 {Z_2}^2 - d_1 d_2 {Z_3}^2 + (e_2 - e_3){Z_4}^2 = 0 \\ \mathcal{C}_4 : &\quad (e_2 - e_3) d_1 {Z_1}^2 + (e_3 - e_1) d_2 {Z_2}^2 + (e_1 - e_2) d_1 d_2 {Z_3}^2. \end{aligned} \tag{3.11}$$

If any one of the four conics fails the solvability test we can eliminate the pair (d_1, d_2) from consideration. Once this has been done for every pair $(d_1, d_2) \in K(S,2) \times K(S,2)$ we test the local solvability of the remaining quadric-intersections.

3.3.2 Elimination via local solvability

At this stage in our algorithm we have a set of pairs $(d_1, d_2) \in K(S,2) \times K(S,2)$, each defining a quadric intersection $\mathcal{C}_{d_1,d_2}$ (and hence a pencil of quadrics) over K. All the singular quadrics in each pencil are solvable as conics. Instead of working directly with the quadric intersections we convert them to equivalent curves of the form $Y^2 = f(U,V)$ where $f(U,V)$ is a homogeneous quartic. This is because it is a relatively simple matter to adapt an existing algorithm (see [MSS]) for curves given by $Y^2 = f(X)$ over $\mathbb{Q}$ to work over $\mathbb{F}_q(T)$. This will be discussed further in chapter 6. Since all conics in the pencil are solvable and therefore parametrizable we can solve any one of the four and substitute its parametrization into another. There is no obvious advantage to be gained in choosing one of the twelve possible ways of doing this over another. The

MAGMA implementation of this algorithm parametrizes $\mathcal{C}_1$ and substitutes into $\mathcal{C}_2$.

3.3.3 The upper bound

Up to this point we have calculated the group $K(S,2)$, and for each element $(d_1, d_2) \in K(S,2) \times K(S,2)$ defined a quadric intersection over K. After eliminating some of these via the methods described in this section we are left with a subgroup of pairs (d_1, d_2) whose associated quadric intersections are everywhere locally solvable. This set actually forms a subgroup B of $K(S,2) \times K(S,2)$ containing the image under ϕ of $E(K)/2E(K)$. Hence $\#E(K)/2E(K) \leq \#B$, therefore $r = \text{rank}(E) \leq s$ where $\#B = 2^{s+2}$. This is our upper bound on r.

3.4 The lower bound on the rank

We now have an upper bound s on r. Our next aim is to find a lower bound. This will be accomplished by explicitly finding points on the homogeneous spaces $\mathcal{C}_{d_1,d_2}$. Often we do not have to search all 2^{s+2} curves for points: since B is itself a group if $\mathcal{C}_{a,b}$ and $\mathcal{C}_{c,d}$ have points then so does $\mathcal{C}_{ab,cd}$ $(a,b,c,d \in K(S,2))$. So we may be able to establish the existence of points on some curves even when the search algorithm used does not find any.

3.4.1 Height bound

To search for points on quadric intersections over K we first choose a bound H on the degree (or "height") of such points. This immediately means that we may not find any points on a curve either because none exist or because they

lie outside the bound. Increasing H increases the likelihood of the algorithm succeeding in finding a point but at the cost of run-time.

The search algorithm is discussed more fully in chapter 8 but the basic idea is as follows:

a) Choose a height bound H on the points being searched for;

b) Construct a set of lattices over $\mathbb{F}_q[T]$ (see chapter 7) such that each point on $\mathcal{C}$ of degree $\leq H$ is contained in exactly one of these lattices;

c) Use polynomial lattice-basis reduction (chapter 7) to find short vectors in the lattices.

A lattice is constructed for each point on $\mathcal{C}$ mod p for some prime $p \in \mathbb{F}_q[T]$ where $\deg(p) = 1 + \lfloor \frac{2H}{3} \rfloor$ therefore the number of lattices is approximately $q^{\frac{2H}{3}}$ by the Hasse Bound. The run-time therefore increases exponentially with H.

3.4.2 The lower bound

Once we have chosen H and the search algorithm has completed its run we have a subgroup A of $K(S,2) \times K(S,2)$ consisting of (d_1, d_2) whose associated quadric intersections each have a global point (though we may not have found a point on all of them). The lower bound on the rank of E is t, where $\#A = 2^{t+2}$.

3.5 The algorithm

This is a summary of the algorithm for performing 2 descent on an elliptic curve E over $K = \mathbb{F}_q(T)$ when E has full 2-torsion over K. We define an operation $\star$ on quadric-intersections as follows:

$$\mathcal{C}_{d_1,d_2} \star \mathcal{C}_{d_3,d_4} = \mathcal{C}_{d_1 d_3, d_2 d_4}. \tag{3.12}$$

ALGORITHM 3.5.1: **Full 2 descent on elliptic curves over $\mathbb{F}_q(T)$**

INPUT: $e_1, e_2, e_3 \in \mathbb{F}_q[T]$ all different, H: a non-negative integer

OUTPUT: A lower and upper bound on $\mathrm{rank}(E)$ where E is the elliptic curve over $\mathbb{F}_q(T)$ given by $Y^2 = (X - e_1)(X - e_2)(X - e_3)$

1. S_1:=$\{\infty\} \cup \{$primes dividing $(e_1 - e_2)(e_1 - e_3)\}$
   S_2:=$\{\infty\} \cup \{$primes dividing $(e_1 - e_2)(e_2 - e_3)\}$
   K:=$\mathbb{F}_q(T)$
   SC:=$\{\mathcal{C}_{1,1}\}$ (the set of ELS homogeneous spaces)
   XC:={ } (the set of not ELS homogeneous spaces)
   PC:={ } (the set of homogeneous spaces with a global point)
   SE:={ } (the set of found points on E);
2. FOR $(d_1, d_2) \in K(S_1, 2) \times K(S_2, 2)$ DO;
3. $\mathcal{C}$:=$\mathcal{C}_{d_1,d_2}$;
4. IF $\mathcal{C} \in XC$ THEN CONTINUE;
5. END IF;
6. $\mathcal{C}_1, \ldots, \mathcal{C}_4$:= 4 singular quadrics in pencil associated to $\mathcal{C}$;
7. IF all $\mathcal{C}_i$ are solvable THEN;
8. parametrize $\mathcal{C}_1$;
9. f:=quartic formed by substituting parametrization of $\mathcal{C}_1$ into $\mathcal{C}_2$;
10. ELSE XC:=$XC \cup \{\mathcal{C}\}$;
11. XC:=$XC \cup \{\mathcal{C} \star \mathcal{D} : \mathcal{D} \in SC\}$;
12. CONTINUE;
13. END IF;
14. IF f is not everywhere locally solvable THEN;

15. XC:=$XC \cup \{\mathcal{C} \star \mathcal{D} : \mathcal{D} \in SC\}$;

16. CONTINUE;

17. ELSE SC:=$SC \cup \{\mathcal{C}\}$;

18. XC:=$XC \cup \{\mathcal{C} \star \mathcal{D} : \mathcal{D} \in XC\}$;

19. END IF;

20. search for points of degree $\leq H$ on $\mathcal{C}$;

21. IF $\mathcal{C}$ has a point Q THEN P:=$\psi(Q) \in E(K)$;

22. SE:=$SE \cup \{P\}$;

23. PC:=$PC \cup \{\mathcal{C}\} \cup \{\mathcal{C} \star \mathcal{D} : D \in PC\}$;

24. END IF;

25. END FOR;

26. r_u:=$\log_2(\#SC) - 2$;

27. G:=subgroup of $E(K)$ generated by SE;

28. r_1:=$\#\text{basis}(G)$;

29. RETURN r_l, r_u;

3.6 Example

We give here a worked example of the 2-descent algorithm applied to an elliptic curve over $\mathbb{F}_{59}(T)$.

3.6.1 Constructing the homogeneous spaces

Let K be the function field $\mathbb{F}_{59}(T)$. Let E be the following elliptic curve over K:

$$Y^2 = (X - e_1)(X - e_2)(X - e_3) \tag{3.13}$$

where

$$\begin{aligned} e_1 &= 36T + 49 \\ e_2 &= 2T + 57 \\ e_3 &= 11T^2 + 26T + 4. \end{aligned} \tag{3.14}$$

The bad places of E are ∞ and the primes dividing $(e_1-e_2)(e_2-e_3)(e_3-e_1)$. These are:

$$\{T+17, T+25, T+31, T+32, T+36\}. \tag{3.15}$$

We also require a non-square element α of $\mathbb{F}_{59}$ in order to construct the sets from which we derive the equations of the homogeneous spaces. In this case $\alpha = 2$ will suffice. We construct quadric intersections with equations:

$$\begin{cases} d_1 Z_1^2 - d_2 Z_2^2 + (e_1 - e_2) Z_4^2 = 0 \\ d_1 Z_1^2 - d_1 d_2 Z_3^2 + (e_1 - e_3) Z_4^2 = 0. \end{cases} \tag{3.16}$$

As demonstrated in the remark at the end of 3.2.2, we do not need to construct one of these for each pair $(d_1, d_2) \in K(S,2) \times K(S,2)$ where S is the set of bad places of E. It suffices to choose $d_1 \in K(S_1, 2)$ and $d_2 \in K(S_2, 2)$. In this example, the groups $K(S_1, 2)$ and $K(S_2, 2)$ are generated by:

$$\begin{aligned} K(S_1, 2): &\quad \{2, T+17, T+25, T+31\} \\ K(S_2, 2): &\quad \{2, T+31, T+32, T+56\}. \end{aligned} \tag{3.17}$$

Each of the above groups has sixteen elements, so there are 256 quadric intersections to analyse. We now give worked examples of the analysis of two of these.

3.6.2 Analysis of the homogeneous spaces

The curve with $(d_1, d_2) = (2, 1)$

The quadric intersection $\mathcal{C}_{d_1,d_2}$ in this case has the following equations:

$$\begin{cases} (1) \quad 2Z_1^2 - Z_2^2 + (34T + 51)Z_4^2 = 0 \\ (2) \quad 2Z_1^2 - 2Z_3^2 + (48T^2 + 10T + 45)Z_4^2 = 0. \end{cases} \tag{3.18}$$

Applying the conic solving algorithm (chapter 5) to equation (1) above shows that there are no points on the conic and therefore none on the curve $\mathcal{C}_{d_1,d_2}$.

The curve with $(d_1, d_2) = ((T + 17)(T + 35), 1)$

This curve has equations:

$$\begin{cases} (1) \quad (T + 17)(T + 35)Z_1^2 - Z_2^2 + (34T + 51)Z_4^2 = 0 \\ (2) \quad (T + 17)(T + 35)Z_1^2 - (T + 17)(T + 35)Z_3^2 + (48T^2 + 10T + 45)Z_4^2 = 0. \end{cases} \tag{3.19}$$

Applying the conic solving algorithm we find that the conic defined by (1) has a solution $(56 : 56T+10 : 1)$ and the conic defined by (2) has a solution $(55 : 8 : 1)$. We then parametrize the conic (1) to get $Z_1 = q_1(U, V)$ and $Z_4 = q_3(U, V)$ (we do not need Z_2) where:

$$\begin{aligned} q_1(U, V) &= 20U^2 + (16T + 24)V^2 \\ q_3(U, V) &= 13U^2 + (2T + 13)UV + (25T + 8)V^2. \end{aligned} \tag{3.20}$$

Substituting this into (2) and multiplying by $d_1 d_2$ gives

$$d_1^2 d_2^2 Z_3^2 = d_1 d_2 (d_1 q_1^2 + (e_1 - e_3)q_3^2) \tag{3.21}$$

or, replacing $d_1d_2Z_3^2$ by Y^2 we have $Y^2 = f(U,V)$ where

$$\begin{aligned} f(U,V) = \quad & (16T^4+46T^3+52T^2+21T+3)U^4 \\ & +(18T^5+36T^4+4T^3+13T^2+50T+33)U^3V \\ & +(15T^6+19T^5+42T^4+14T^3+8T^2+2T+33)U^2V^2 \\ & +(21T^6+44T^5+48T^4+32T^3+27T^2+8T+43)UV^3 \\ & +(48T^6+46T^5+29T^4+45T^3+18T^2+36T+35)V^4. \end{aligned} \tag{3.22}$$

This quartic can be shown to be everywhere locally solvable (ELS) using the local solvability algorithm described in chapter 6.

3.6.3 Computing the rank

After analysing all 256 homogeneous spaces $\mathcal{C}_{d_1,d_2}$ for solvability we are left with 16, corresponding to the subgroup B of $K(S_1,2)\times K(S_2,2)$ given by the table below. The fourth column of the table represents elements of $K(S_1,2)\times K(S_2,2)$ in the group ${\mathbb{F}_2}^4\times{\mathbb{F}_2}^4$ with the ordering of the generators $T+17, T+25, T+31, 2$ for $K(S_1,2)$ and $T+31, T+32, T+56, 2$ for $K(S_2,2)$.

We now need to search for points on these curves, using the algorithm of chapter 8. Searching with height bound 0 yields points on 5 of them: #5,6,9,10 and 14. We map these back to points on E and label them P_5, P_6, P_9, P_{10} and P_{14} respectively. P_{14} is the 2-torsion point $(11T^2+26T+4, 0)$ whereas the others are all points of infinite order. They are:

$$\begin{aligned} P_5 &= (18T+38, 10T^2+18T+12) \\ P_6 &= (48T+54, 51T^2+16T+31) \\ P_9 &= (49T+34, 8T^2+53T+5) \\ P_{10} &= (20T+3, 10T^2+43T+17). \end{aligned} \tag{3.23}$$

We know of three other 2-torsion points on E (using (3.7) we see that

these are P_1, P_4 and P_{15}) so we have 8 points. This gives a lower bound on rank(E) of 1 since $\lceil \log_2(8) - 2 \rceil = 1$ but we can get a higher bound by observing from the table that P_5 and P_6 are independent. In fact the subgroup B of $K(S_1, 2) \times K(S_2, 2)$ is isomorphic to the subgroup of ${\mathbb{F}_2}^4 \times {\mathbb{F}_2}^4$ generated by $(1000, 0100), (0100, 0100), (1101, 0111)$ and $(0010, 1110)$, corresponding

Figure 3.1: Table showing ELS homogeneous spaces.

#	d_1	d_2	${\mathbb{F}_2}^4 \times {\mathbb{F}_2}^4$
1	1	1	$(0000, 0000)$
2	$(T+17)(T+25)$	1	$(1100, 0000)$
3	$2(T+31)$	$2(T+31)$	$(0011, 1001)$
4	$2(T+17)(T+25)(T+31)$	$2(T+31)$	$(1111, 1001)$
5	$T+17$	$T+32$	$(1000, 0100)$
6	$T+25$	$T+32$	$(0100, 0100)$
7	$2(T+17)(T+31)$	$2(T+31)(T+32)$	$(1011, 1101)$
8	$2(T+25)(T+31)$	$2(T+31)(T+32)$	$(0111, 1101)$
9	$2(T+17)$	$2(T+56)$	$(1001, 0011)$
10	$2(T+25)$	$2(T+56)$	$(0101, 0011)$
11	$(T+17)(T+31)$	$(T+31)(T+56)$	$(1010, 1010)$
12	$(T+25)(T+31)$	$(T+31)(T+56)$	$(0110, 1010)$
13	2	$2(T+32)(T+56)$	$(0001, 0111)$
14	$2(T+17)(T+25)$	$2(T+32)(T+56)$	$(1101, 0111)$
15	$T+31$	$(T+31)(T+32)(T+56)$	$(0010, 1110)$
16	$(T+17)(T+25)(T+31)$	$(T+31)(T+32)(T+56)$	$(1110, 1110)$

to P_5, P_6, P_{14} and P_{15} respectively. This raises the lower bound to 2. The upper bound is $\log_2(s) - 2$ where s is the size of the subgroup of ELS homogeneous spaces. So our upper bound is $\log_2(16) - 2 = 2$. Since the upper and lower bounds match we conclude that E has rank 2, with P_5 and P_6 generating a subgroup of finite index.

Chapter 4

Descent via 2-isogeny

4.1 Introduction

The full 2-descent algorithm described in chapter 3 only works for elliptic curves with full 2-torsion, i.e. those that can be written as
$Y^2 = (X - e_1)(X - e_2)(X - e_3)$. This section describes an alternative method that may be employed when E has at least one non-trivial 2-torsion point.

The algorithm "descent via 2-isogeny" uses two elliptic curves: E itself and an isogenous curve E' with isogenies $\phi : E \to E'$ and $\phi' : E' \to E$ between them. The isogenies ϕ and ϕ' are dual and of degree 2 (see [Sil1, pp. 84-90]) hence $\phi' \circ \phi : E \to E$ is the multiplication-by-two map on E and $\phi \circ \phi' : E' \to E'$ is the multiplication-by-two map on E'. The aim of the algorithm is, like full 2-descent, to find the rank or rank-bounds by computing $E(K)/2E(K)$ where $K = \mathbb{F}_q(T)$. This is done by first finding $E(K)/\phi'(E'(K))$ and $E'(K)/\phi(E(K))$. From this we find rank bounds and possibly the rank of $E(K)$.

4.2 The isogenous curves

Let E be an elliptic curve over $K = \mathbb{F}_q(T)$, with at least one non-trivial two-torsion point. As before we assume that $\text{char}(K) \neq 2$ so we can write E as a Weierstrass equation:

$$E : Y^2 = (X - e)(X^2 + aX + b) \tag{4.1}$$

with $e, a, b \in \mathbb{F}_q[T]$. By changing variables through $(x, y) \mapsto (x - e, y)$ we can write E as the elliptic curve

$$Y^2 = X(X^2 + cX + d) \tag{4.2}$$

with two-torsion point $(0, 0)$. The isogenous curve E', also defined over K, is given by the equation

$$Y^2 = X(X^2 + c'X + d') \tag{4.3}$$

where $c, c', d, d' \in \mathbb{F}_q[T]$ and $c' = -2c$, $d' = c^2 - 4d$ and $dd' \neq 0$. To calculate $E(K)/\phi'(E'(K))$ we use the following theorem:

Theorem 7. *Let E be an elliptic curve over $K = \mathbb{F}_q(T)$ given by an equation of the form:*

$$E : Y^2 = X(X^2 + cX + d).$$

Then:

a) There is an isogenous curve:

$$E' : Y^2 = X(X^2 + c'X + d')$$

where $c' = -2c$ and $d' - c^2 - 4d$; and dual 2-isogenies:

$$\phi : E \to E'$$

$$\phi' : E' \to E$$

given by the formulae:

$$\begin{array}{l} \phi(x,y) = \left(\frac{x^2+cx+d}{x}, \frac{x^2y-dy}{x^2}\right) \\ \phi'(x,y) = \left(\frac{x^2+c'x+d'}{4x}, \frac{-x^2y+d'y}{8x^2}\right) \end{array} . \tag{4.4}$$

b) There is an embedding ψ of $E(K)/\phi'(E'(K))$ into the (finite) subgroup of K^*/K^{*2} generated by the divisors of d and a non-square element of $\mathbb{F}_q$.

The map ψ is given by:

$$\psi(x,y) = \begin{cases} x \ mod \ K^{*2} & if \ x \neq 0 \\ d \ mod \ K^{*2} & if \ x = 0. \end{cases}$$

Proof. See [Cr3] and [Cr1, chapter 3]. □

The second part of the above theorem applies to E' as well as E just by swapping E' and E, ϕ and ϕ', and d and d'.

To calculate $E(K)/\phi'(E'(K))$ and $E'(K)/\phi(E(K))$ we search for points on homogeneous spaces - one for each factorization of d in the former case, and one for each factorization of d' in the latter.

4.3 The first descent

Let E be an elliptic curve over $K = \mathbb{F}_q(T)$ of the form $Y^2 = X(X^2 + cX + d)$. Similar to the full 2-descent, each factorization $d = d_1d_2$ with d_1 square-free gives rise to a homogeneous space $\mathcal{C}_{d_1}$. We must therefore decide for each d_1 whether C_{d_1} is everywhere locally solvable, and if so we must search for points on $\mathcal{C}_{d_1}$.

4.3.1 Homogeneous spaces

Let G_d be the subgroup of K^*/K^{*2} generated by the divisors of d and a non-square $\alpha \in \mathbb{F}_q$. For example if $K = \mathbb{F}_5(T)$ and $d = T(T+3)^2$ then

$$G_d = \{1, 2, T, 2T, T+3, 2T+1, T(T+3), 2T(T+3)\}.$$

For each element d_1 of G_d we define the homogeneous space $\mathcal{C}_{d_1}$ as follows:

$$\mathcal{C}_{d_1} : U^2 = d_1 V^4 + cV^2 + d_2 \tag{4.5}$$

where $d_2 = d/d_1$ and c is the same c from the equation of E. Since this equation is of the form $Y^2 = \textit{quartic}$ (see 2.2) we can homogenize the RHS and apply the local-solvability algorithm described in chapter 6. If $\mathcal{C}_{d_1}$ turns out to be everywhere locally solvable, in order to search for points we can convert $\mathcal{C}_{d_1}$ into a quadric intersection according to 2.4 and apply the algorithm of Chapter 8.

If we find a point on $\mathcal{C}_{d_1}$ then we may map back to E via $(u, v) \mapsto (d_1 u^2, d_1 uv)$. Note though that as for full 2-descent, these homogeneous spaces may be everywhere locally solvable and yet not posses a global point. A second descent (see below) may be able to show the non-existence of global points in this case.

Eliminating homogeneous spaces via local solvability testing yields an upper bound on the rank of $E(K)/\phi'(E'(K))$. In fact the ELS curves form a subgroup H_u of G_d. The same is carried out for E' and we end up with a subgroup H'_u of $G_{d'}$, of ELS homogeneous spaces for E. From this the upper bound on the rank of E can be calculated since if $\#H_u = 2^s$ and $\#H'_u = 2^{s'}$ then $\text{rank}(E) \leq s + s' - 2$.

4.3.2 Searching for points on the homogeneous spaces

To search for points on these homogeneous spaces we use the algorithm described in chapter 8. This algorithm is written for curves that are the intersection of two quadrics whereas our homogeneous spaces $\mathcal{C}_{d_1}$ have the form $U^2 = d_1V^4 + cV^2 + d_2$. However if we make the substitution $W = V^2$ into this equation, and homogenize with a fourth variable X we obtain a quadric intersection:

$$\begin{cases} U^2 = d_1W^2 + cV^2 + d_2X^2 \\ V^2 = WX. \end{cases} \tag{4.6}$$

The elements of G_d whose associated homogeneous spaces have a K-rational point form a subgroup H_l of G_d. Therefore if we have found a point on $\mathcal{C}_a$ and $\mathcal{C}_b$ for $a, b \in G_d$ then the curve $\mathcal{C}_{ab}$ also has a K-rational point, and we do not need to apply the search algorithm. Doing the same for E' we get a subgroup H'_l of $G_{d'}$. This then gives us a lower bound on the rank of E: if $\#H_l = 2^t$ and $\#H'_l = 2^{t'}$ then $\text{rank}(E) \geq t + t' - 2$.

4.4 The second descent

During the course of the first descent, we may encounter homogeneous spaces that are everywhere locally solvable but have no points within the height bound of the search algorithm. In such cases we can perform a second descent on these curves, examining the descendants of each in order to a) find points on the first descent curves, or b) prove that no points exist.

In this section we let $\mathcal{C}$ be a first descent homogeneous space for an elliptic curve E defined over $K = \mathbb{F}_q(T)$.

4.4.1 Mapping $\mathcal{C}$ to a conic

A homogeneous space for E arising in the first descent is a curve $\mathcal{C}$ with equation:

$$U^2 = d_1V^4 + cV^2 + d_2 \tag{4.7}$$

with $c, d_1, d_2 \in \mathbb{F}_q[T]$, $d_1d_2 \neq 0$, $c^2 - 4d_1d_2 \neq 0$. At a glance, it can be seen that this equation is actually a conic in U and V^2, so we make the substitutions $Y = U$, $X = V^2$ and homogenize with a third variable Z to obtain the conic:

$$Y^2 = d_1X^2 + cXZ + d_2Z^2. \tag{4.8}$$

This is a semi-diagonal conic (see 5.4.1) and we can find a point on it using the conic-solving algorithm of chapter 5. In fact the curve (4.8) always has a K-rational point by the Hasse principle, since $\mathcal{C}$ is everywhere locally solvable. This does not immediately lead to a point on $\mathcal{C}$ however - this is the case if and only if X/Z is a square. The map back to $\mathcal{C}$ given a solution to (4.8) is

$$(x : y : z) = (a^2 : y : b^2) \mapsto (y, a/b) \in \mathcal{C}. \tag{4.9}$$

4.4.2 Constructing the descendants

Since we always have a K-rational solution to (4.8) we can parametrize all solutions:

$$(\lambda : \mu) \mapsto (q_1(\lambda, \mu) : q_2(\lambda, \mu) : q_3(\lambda, \mu)) \tag{4.10}$$

where q_1, q_2, q_3 are homogeneous quadratic polynomials over $\mathbb{F}_q[T]$. Solutions $(x : y : z)$ to (4.8) only correspond to points on $\mathcal{C}$ when x/z is a square, so we require that $\frac{q_1(\lambda,\mu)}{q_3(\lambda,\mu)}$ is a square. This corresponds to a solution $\lambda, \mu, s, t \in \mathbb{F}_q[T]$ to the equations:

$$\begin{cases} q_1(\lambda, \mu) = d_3s^2 \\ q_3(\lambda, \mu) = d_3t^2 \end{cases} \tag{4.11}$$

where d_3 is a square-free divisor of the resultant of $q_1(\lambda, 1)$ and $q_2(\lambda, 1)$.

It can be seen immediately that the above equations define a quadric intersection D_{d_3} over K - one for each square-free divisor of $\mathrm{Res}(q_1, q_3)$ - and we can deal with these in almost exactly the same way as the quadric intersections arising in the full 2-descent. Firstly, by checking each one is solvable as a conic (this time semi-diagonal) and if so, parametrizing one and substituting into the other to obtain a curve $Y^2 = \mathit{quartic}$ in order to check the local solvability using the algorithm in chapter 6. To those that are everywhere locally solvable we apply the point search algorithm of chapter 8. If all the descendants $\mathcal{D}_{d_3}$ are shown to have no points then the curve $\mathcal{C}$ does not, and is therefore an example of an everywhere locally solvable curve with no global point. If this is not the case, and no points have been found then no new information about $\mathcal{C}$ has been found by performing the second descent.

4.5 The algorithm

Descent via 2-isogeny can be summarised as follows: Firstly, the elliptic curve E has a non-trivial 2-torsion point which we move to $(0, 0)$ via a change of variables. We then construct another elliptic curve E' isogenous to E via dual isogenies ϕ and ϕ' of degree 2. Next, for each square-free divisor of the coefficient d in the equation of E, we construct a homogeneous space $\mathcal{C}$ of the form $Y^2 = \mathit{quartic}$ and check its local solvability. If it is everywhere locally solvable then we may wish to perform a second descent. Otherwise we search for points on $\mathcal{C}$, and hence obtain lower and upper bounds on $\#E(K)/\phi'(E'(K))$. Repeating this for E' we obtain bounds on $\#E'(K)/\phi(E(K))$ and hence bounds on $\mathrm{rank}(E)$.

If the second descent is applied, for each homogeneous space $\mathcal{C}$ that is every-

where locally solvable, we construct a set of descendants $\mathcal{D}_{d_3}$ and check their local solvability and search for points, sharpening the lower and upper bounds obtained through the first descent.

When using the second descent it must be decided whether to apply it a) after a search for points on $\mathcal{C}$ has yielded nothing, or b) straight away once $\mathcal{C}$ is established to be everywhere locally solvable. One is more likely to find points on the second descent because if they exist they will have smaller height than on $\mathcal{C}$, however this approach risks wasting computer time on the second descent if $\mathcal{C}$ has an easy to find point. In the case (a) above the search algorithm (chapter 8) is invoked twice so a height bound must be chosen twice. The default setting in the MAGMA implementation is to set one bound used for both, but the user can choose them independently if desired.

Here is the algorithm for performing descent via 2-isogeny on an elliptic curve E over $K = \mathbb{F}_q(T)$ when E has non-trivial 2-torsion over K. Similar to the full 2-descent, we define an operation $\star$ on the homogeneous spaces of E as follows:

$$\mathcal{C}_a \star \mathcal{C}_b = \mathcal{C}_{ab} \tag{4.12}$$

where $\mathcal{C}_x$ is the curve $U^2 = xV^4 + cV^2 + d/x$.

ALGORITHM 4.5.1: **Descent via 2-isogeny on elliptic curves over $\mathbb{F}_q(T)$**

`INPUT:` E: `an elliptic curve over` $\mathbb{F}_q(T)$ `with a non-trivial 2-torsion point,` H: `a non-negative integer`

`OUTPUT: A lower and upper bound on the rank of` E

1. `Transform` E `to an elliptic curve given by` $Y^2 = X(X^2 + cX + d)$;
2. c'`:=` $-2c$, d'`:=` $c^2 - 4d$

E':= the elliptic curve $Y^2 = X(X^2 + c'X + d')$

ϕ:= the 2-isogeny from E to E'

ϕ':= the dual isogeny from E' to E;

3. G_d:= the subgroup of K^*/K^{*2} consisting of

the square-free divisors of d and a non-square $\alpha \in \mathbb{F}_q^*$

SC_1:= { } - the set of ELS homogeneous spaces of E

PC_1:= { } - the set of homogeneous spaces of E with a found point

PE_1:= { } - the set of points found on $E(K)/\phi'(E'(K))$

XC_1:= { } - the set of not ELS homogeneous spaces of E;

4. $G_{d'}$:= the subgroup of K^*/K^{*2} consisting of

α and the square-free divisors of d'

SC_2:= { } - the set of ELS homogeneous spaces of E'

PC_2:= { } - the set of homogeneous spaces of E' with a found point

PE_2:= { } - the set of points found on $E'(K)/\phi(E(K))$;

XC_2:= { } - the set of not ELS homogeneous spaces of E';

5. FOR $d_1 \in G_d$ DO;

6. $\mathcal{C}$:= the curve $U^2 = d_1V^4 + cV^2 + d_2$ where $d_2 = d/d_1$;

7. IF $\mathcal{C} \in XC_1$ OR $\mathcal{C} \in SC_1$ THEN CONTINUE; END IF;

8. IF $\mathcal{C}$ is everywhere locally solvable THEN;

9. $SC_1 := SC_1 \cup \{\mathcal{C}\} \cup \{\mathcal{C} \star \mathcal{D} : \mathcal{D} \in SC_1\}$;

10. $XC_1 := XC_1 \cup \{\mathcal{C} \star \mathcal{D} : \mathcal{D} \in XC_1\}$;

11. $\mathcal{C}_0$:= the quadric intersection $\begin{cases} U^2 = d_1W^2 + cV^2 + d_2X^2 \\ V^2 = WX \end{cases}$;

12. search for points on $\mathcal{C}_0$ up to height bound H;

13. IF a point on C_0 is found THEN;

14. map point to a point P in $E(K)/\phi'(E'(K))$;

15. PC_1:= $PC_1 \cup \{\mathcal{C}\}$, PE_1:= $PE_1 \cup \{P\}$;

16. END IF;

17. ELSE XC_1:=$XC_1 \cup \{\mathcal{C}\} \cup \{\mathcal{C} \star \mathcal{D} : \mathcal{D} \in SC_1\}$,

18. END IF;

19. END FOR;

20. FOR $d_1 \in G_{d'}$ DO;

21. $\mathcal{C}$:= the curve $U^2 = d_1V^4 + cV^2 + d_2$ where $d_2 = d/d_1$;

22. IF $\mathcal{C} \in XC_2$ THEN CONTINUE; END IF;

23. IF $\mathcal{C}$ is everywhere locally solvable THEN;

24. SC_2:= $SC_2 \cup \{\mathcal{C}\} \cup \{\mathcal{C} \star \mathcal{D} : \mathcal{D} \in SC_2\}$;

25. XC_2:= $XC_2 \cup \{\mathcal{C} \star \mathcal{D} : \mathcal{D} \in XC_2\}$;

26. $\mathcal{C}_0$:= the quadric intersection $\begin{cases} U^2 = d_1W^2 + cV^2 + d_2X^2 \\ V^2 = WX \end{cases}$;

27. search for points on $\mathcal{C}_0$ up to height bound H;

28. IF a point on C_0 is found THEN;

29. map point to a point P in $E'(K)/\phi(E(K))$;

30. PC_2:= $PC_2 \cup \{\mathcal{C}\}$, PE_2:= $PE_2 \cup \{P\}$;

31. END IF;

32. ELSE XC_2:=$XC_2 \cup \{\mathcal{C}\} \cup \{\mathcal{C} \star \mathcal{D} : \mathcal{D} \in SC_2\}$;

33. END IF;

34. END FOR;

35. s_l:= $\lceil \log_2(\#PC_1) \rceil$, s'_l:= $\lceil \log_2(\#PC_2) \rceil$,
s_u:= $\log_2(\#SC_1)$, s'_u:= $\log_2(\#SC_2)$;

36. PE:= $PE_1 \cup \{\phi'(P) : P \in PE_2\}$;

37. r_l:= #{basis of group generated by PE};

38. r_u:= $s_u + s'_u - 2$;

39. RETURN r_l, r_u;

ALGORITHM 4.5.2: **Subroutine: second descent**

INPUT: An ELS curve $\mathcal{C}$ over K with equation $U^2 = d_1V^4 + cV^2 + d_2$,
H a non-negative integer

OUTPUT: $i \in -1, 0, 1$ with $i = -1$ if $\mathcal{C}$ has no points,
$i = 0$ if no points are found in the search,
and $i = 1$ if a point is found; a point P on $\mathcal{C}$ if $i = 1$

1. $\mathcal{C}_0$:= the conic $Y^2 = d_1X^2 + cXZ + d_2Z^2$;
2. $(X : Y : Z) = q_1(\lambda, \mu) : q_2(\lambda, \mu) : q_3(\lambda, \mu)$:= parametrization of $\mathcal{C}_0$;
3. R:= $\mathrm{Res}(q_1(\lambda, 1), q_3(\lambda, 1))$;
4. S:= {square-free divisors of R};
5. j:=0 (counter);
6. FOR $d_3 \in S$ DO;
7. $\mathcal{D}$:= the quadric intersection $\begin{cases} q_1(\lambda, \mu) := d_3 s^2 & (1) \\ q_3(\lambda, \mu) := d_3 t^2 & (2) \end{cases}$;
8. IF conics (1) and (2) are both solvable THEN;
9. parametrize (1) and substitute into (2);
10. $\mathcal{D}_0$:= the $Y^2 = quartic$ curve resulting from this;
11. IF $\mathcal{D}_0$ is NOT ELS THEN;
12. $j := j + 1$;

```
13.             CONTINUE; END IF;
14.         search for points on 𝒟 with height bound H;
15.         IF 𝒟 has a point then map to a point P on 𝒞;
16.             RETURN 1,P; END IF;
17.         ELSE j := j + 1;
18.         CONTINUE; END IF;
19.     END FOR;
20. IF j = #S THEN;
21.     RETURN -1; END IF;
22. RETURN 0;
```

The second descent sub-routine can be called either when $\mathcal{C}$ is shown to be ELS or after additionally searching for points on $\mathcal{C}$ has returned nothing. If the second descent returns 1 then $\mathcal{C}$ is added to PC, the point mapped to a point on E and added to PE. If -1 is returned then $\mathcal{C}$ is removed from SC. No changes are made if the second descent returns 0.

4.6 Example

Here we give an example of descent via 2-isogeny over $\mathbb{F}_{31}(T)$.

4.6.1 Construction of the homogeneous spaces

Let E be the elliptic curve over $K = \mathbb{F}_{31}(T)$ with equation:

$$Y^2 = X(X^2 + (6T + 1)X + (3T + 6)). \tag{4.13}$$

E has 2-torsion point $(0, 0)$ and is in the required form so we do not need to change variables here. Using the notation of the rest of this chapter, $c = 6T + 1$

and $d = 3T + 6$. The isogenous curve E' is given by the equation

$$Y^2 = X(X^2 + (19T + 29)X + (5T^2 + 8)). \tag{4.14}$$

Let $c' = -2c = 19T + 29$ and $d' = c^2 - 4d = 5T^2 + 8$. We also need a square-free element α of $\mathbb{F}_{31}$ - in this case we use $\alpha = 3$. Since $d = 3T + 6 = 3(T + 2)$ we construct a homogeneous space $\mathcal{C}_{d_1}$ for each element d_1 of the subgroup G_d of K^*/K^{*2} generated by the divisors of d. In this case $G_d = \{1, T+2, 3, 3(T+2)\}$. For $d_1 \in G_d$ we have a homogeneous space given by the equation

$$U^2 = d_1V^4 + cV^2 + d_2 \tag{4.15}$$

where $d_1d_2 = d$. There are four such curves in this case.

Next we do the same for E'. $d' = 6^2(T^2 + 14)$ so the group $G_{d'}$ consists of four elements: $G_{d'} = \{1, 3, T^2 + 14, 3(T^2 + 14)\}$. We construct homogeneous spaces $\mathcal{C}_{d'_1}$ for each $d'_1 \in G_{d'}$. These have equation

$$U^2 = d'_1V^4 + c'V^2 + d'_2 \tag{4.16}$$

where $d'_1d'_2 = d'$. We now have 8 curves to analyse.

4.6.2 Analysis and rank computation

First consider the four curves $\mathcal{C}_{d_1}$. Applying the local-solvability algorithm (chapter 6) to each shows that all are ELS. Next we convert these into quadric intersections (see 2.4) in order to search for points using the algorithm of chapter 8. A search with height bound 0 yields points on all four curves. Mapping back to E, these points are $(0, 0)$, $(15, 30)$, $(25T + 19, 19T + 7)$ and the point at infinity. Doing the same for E' we find that two of its homogeneous spaces are ELS and both have points. Mapping these back to E' then to E via the dual

isogeny from E' to E yields no new points - these are just $(0,0)$ and the point at infinity.

From the size of the subgroups of ELS curves we get an upper bound on the rank of E. In this case the group sizes are 2^2 and 2^1 respectively so the upper bound is $2+1-2=1$. Since we have found points on all ELS homogeneous spaces we can conclude that the rank of E is in fact 1. The second descent was not needed. A non-torsion generator of $E \pmod{2E}$ is $(15,30)$.

Chapter 5

Solving conics over function fields

5.1 Introduction

The following chapter gives an overview of an algorithm for solving conics over function fields. It is used frequently in the main descent algorithms to decide solvability of certain homogeneous spaces. The algorithm, which is similar to one of the classical algorithms for solving conics over $\mathbb{Q}$ was simultaneously and independently being developed in 2004 by myself and Cremona, and van Hoeij. A full description has been published as VAN HOEIJ, CREMONA - *Solving conics over function fields* ([vHCr]), to which we refer for further details, giving here only the main outline. The algorithm has been implemented in Maple by van Hoeij and in Magma by Cremona and myself.

Throughout this chapter we let F denote a field of characteristic $\neq 2$.

5.1.1 Definition

A conic over a field F is a smooth, projective curve $\mathcal{C}$ in $\mathbb{P}^2(F)$ defined by a quadratic equation of the form:

$$aX^2 + bY^2 + cZ^2 + dXY + eYZ + fXZ = 0 \quad (a, b, c, d, e, f \in F). \tag{5.1}$$

Such a curve can be represented by a 3×3 symmetric matrix M with entries in F satisfying $\mathbf{X}^T M \mathbf{X} = 0$ where $\mathbf{X}$ is a column vector of the variables X, Y and Z. Using the notation of the above equation, M looks like:

$$M = \begin{pmatrix} 2a & d & f \\ d & 2b & e \\ f & e & 2c \end{pmatrix}. \tag{5.2}$$

M is unique up to multiplication by a nonzero element of F. If the matrix M is diagonal then the conic defined by M is called a *diagonal conic*. The condition for $\mathcal{C}$ to be smooth is $\det(M) \neq 0$.

5.2 Solution and parametrization: base cases

If a conic $\mathcal{C}$ has a point $(x : y : z) \in \mathbb{P}^2(F)$ then every point in $\mathcal{C}(F)$ can be easily found via parametrization. This is because a line (defined over F) drawn through a point P in $\mathcal{C}(F)$ intersects the curve at precisely one other point (P itself if the line is a tangent) in $\mathcal{C}(F)$. Intersecting $\mathcal{C}$ with a general line through P (represented by an element of $\mathbb{P}^1(F)$ giving its slope) gives a parametrization:

$$\begin{array}{rcl} \theta : \mathbb{P}^1 & \to & \mathbb{P}^2 \\ (u : v) & \mapsto & (q_1(u, v) : q_2(u, v) : q_3(u, v)). \end{array} \tag{5.3}$$

The q_i are homogeneous quadratics in u and v with no common factor. For efficiency when using a conic parametrization in other algorithms, one must

choose the q_i carefully (see [Sim1]).

5.2.1 Base cases: $\mathbb{Q}$ and $\mathbb{F}_q$

Over the rational numbers, the problem of conic solving was solved by Legendre, and a number of algorithms exist for finding solutions when they exist (see [CrRu]). Over a finite field $\mathbb{F}_q$ ($2 \nmid q$), a simple pigeon-hole argument shows that a solution always exists and it is easy to find a solution: by choosing random elements of $\mathbb{F}_q$ for X and Y the quantity $\frac{1}{c}(aX^2 + bY^2)$ is a square in $\mathbb{F}_q$ with probability $\frac{1}{2}$.

5.3 Reduced conics

In all that follows all conics will be diagonal conics unless otherwise stated.

Let $\mathcal{C}$ be a conic over a function field $F(T)$ given by the equation

$$aX^2 + bY^2 + cZ^2 = 0 \tag{5.4}$$

where $a, b, c \in F(T)$. $\mathcal{C}$ is said to be *reduced* if the following conditions hold:

(1) $a, b, c \in F[T]$;

(2) $\gcd(a, b) = \gcd(b, c) = \gcd(c, a) = 1$;

(3) $a, b,$ and c are square-free (ignoring squares in F).

Condition (1) is achieved by multiplying a, b and c by the LCM of their denominators. For (2) we note that the conic with coefficients a, b, c is equivalent to the conic $\mathcal{C}'$ with coefficients $a' = a/g, b' = b/g, c' = cg$ if g divides both a and b (since if $(x : y : z)$ is a solution to $\mathcal{C}$ then $(x : y : z/g)$ is a solution to $\mathcal{C}'$). Also, since $\deg(abc) > \deg(a'b'c')$ when $\deg(g) > 0$, we can repeat this step a finite

number of times with $g = \gcd(a, b)$, $g = \gcd(b, c)$ and $g = \gcd(c, a)$, until (2) is satisfied. Finally, for (3) if say $a = gh^2, \quad h \notin F$ then replacing a with g gives an equivalent conic $\mathcal{C}'$ with $(x : y : z) \in \mathcal{C} \Leftrightarrow (hx : y : z) \in \mathcal{C}'$.

Cremona and Van Hoeij [vHCr] give an algorithm "ReduceConic" which takes a diagonal conic $\mathcal{C}$ as input and returns a new conic $\mathcal{C}'$ that is reduced, together with $\lambda, \mu, \nu \in K^*$ such that $(x : y : z) \in \mathcal{C}' \Leftrightarrow (\lambda x : \mu y : \nu z) \in \mathcal{C}$.

5.4 The algorithm of Cremona and Van Hoeij

It is shown in [vHCr] that the existence of a point on a diagonal conic depends on the existence of a so-called "solubility certificate". A solubility certificate is a set of conditions on the coefficients a, b and c that must be satisfied in order for a solution to be permitted to exist. It is shown furthermore that the converse is also true, and the existence of a solubility certificate implies that the conic is solvable.

The algorithm first computes a solubility certificate of a reduced conic and then, if a certificate exists runs a sub-algorithm "FindPoint" to produce a solution from it. As a preliminary the following lemma is needed:

Lemma 8. *If a solution* $(x : y : z) \in \mathbb{P}^2(K)$ *of equation* (5.4) *exists, then*

$$l_a X^2 + l_b Y^2 + l_c Z^2 = 0 \tag{5.5}$$

where l_a, l_b *and* l_c *are the leading coefficients of* a, b *and* c *respectively, has a solution in* $\mathbb{P}^2(F)$.

Proof. See [vHCr] □

We will refer to the above equation/conic as the *leading coefficient equation/conic* associated to the conic defined by a, b and c.

Before we define the solubility certificate we establish some more notation:

$$\begin{aligned} L_p &:= \text{the residue field } F[T]/(p) \text{ for a monic irreducible } p \in F[T]; \\ \operatorname{supp}(f) &:= \text{the set of all monic irreducibles dividing } f \in F[T]. \end{aligned}$$

Let f_a, f_b, and f_c be the following polynomials over $F[T]$ in a new variable U:

$$f_a := bU^2 + c, \quad f_b := cU^2 + a, \quad f_c := aU^2 + b. \tag{5.6}$$

We can now define a solubility certificate:

Definition. Let $\mathcal{C}$ be a reduced conic defined by $aX^2 + bY^2 + cZ^2 = 0$. A solubility certificate for $\mathcal{C}$ is a list containing the following:

(1) For every $p \in \operatorname{supp}(a)$, a root α_p of f_a mod p in L_p;

(2) For every $p \in \operatorname{supp}(b)$, a root α_p of f_b mod p in L_p;

(3) For every $p \in \operatorname{supp}(c)$, a root α_p of f_c mod p in L_p;

(4) If $\deg(a) \equiv \deg(b) \equiv \deg(c) \equiv 0 \pmod 2$, either a solution or a solubility certificate for the leading-coefficient equation associated to $\mathcal{C}$.

Given a solubility certificate, it is shown in [vHCr] that a solution exists satisfying the following conditions[1]: $x, y, z \in F[T]$ and

(1) (a) $y \equiv \alpha_p z \pmod p$ for all $p \in \operatorname{supp}(a)$;

(b) $z \equiv \alpha_p x \pmod p$ for all $p \in \operatorname{supp}(b)$;

(c) $x \equiv \alpha_p y \pmod p$ for all $p \in \operatorname{supp}(c)$;

(2) (a) $\deg(x) \leq \left\lceil \frac{1}{2} \deg(bc) \right\rceil - 1$;

[1] for simplicity we only treat here the case where $\deg(a)$, $\deg(b)$, $\deg(c)$ do not all have the same parity. The other case is similar, using also the solution to the leading-coefficient equation: see [vHCr] for details.

(b) $\deg(y) \leq \left\lceil \frac{1}{2} \deg(ac) \right\rceil - 1$;

(c) $\deg(z) \leq \left\lceil \frac{1}{2} \deg(ab) \right\rceil - 1$;

Finding x, y, and z is now a problem in linear algebra over F: the number of equations is $\deg(abc)$ and the number of unknown coefficients is $\deg(abc) + 1$ so there is always a non trivial solution.

It can be seen from the above definition of a solubility certificate is that there is sometimes a "bad" case when a solution of a conic over F is needed. If $F(T)$ is say, $\mathbb{Q}(T_1, \ldots, T_{n-1}, T)$ then in the worst case we would need to solve a conic over $\mathbb{Q}(T_1, \ldots, T_{n-1})$ then one over $\mathbb{Q}(T_1, \ldots, T_{n-2})$ and so on until the base case over $\mathbb{Q}$ is reached.

5.4.1 Non-diagonal conics

The algorithm in [vHCr] only considers diagonal conics. However a non-diagonal conic may be transformed to an equivalent diagonal conic by completing the square. As far as the 2-descent algorithm for elliptic curves is concerned, the most general form of conic needed to be solved is the "semi-diagonal" one. A semi-diagonal conic has equation $Y^2 = aX^2 + bXZ + cZ^2$. This may be solved easily by completing the square (since $\mathrm{char}(K) \neq 2$) to obtain

$$Y^2 = a\left(X + \frac{b}{2a}Z\right)^2 + \left(\frac{c}{a} - \frac{b^2}{4a^2}\right)Z^2 \tag{5.7}$$

- a diagonal conic with coefficients a, -1 and $\frac{c}{a} - \frac{b^2}{4a^2}$ in the variables $X' = X + \frac{b}{2a}Z$, $Y' = Y$ and $Z' = Z$ respectively.

5.4.2 Implementations in Magma

The programs for solving diagonal conics, semi-diagonal conics and the general case have now all been implemented in Magma by myself and John Cremona

[Cr2, conicsFF.m]. The more general cases reduce first to the diagonal case by completing the square. These programs work for multivariate function fields $F(T_1, \ldots, T_n)$ using recursion, with base cases $F = \mathbb{Q}$ or $F = \mathbb{F}_q$. A second program ([Cr2, ternquadFT.m]) is an implementation of an alternative algorithm for solving the non-diagonal case, using minimization and reduction of the associated 3×3 matrix.

Chapter 6

Local-solvability of quartics over $\mathbb{F}_q(T)$

6.1 Background and motivation

During the descent process it is usually necessary to decide whether or not two quadratic equations in four variables have a common solution. In other words we seek a way of proving that a quadric-intersection (QI) either has or does not have any points. We do this by considering local solvability: if a QI has a point over its base field (i.e a global solution) then it clearly has a solution mod p for any prime p and by Hensel's Lemma this lifts to a p-adic solution. If this existence of a p-adic solution at every p occurs in a curve, it is said to be *everywhere locally solvable*. The converse of the above is more useful to us however: if a curve is not solvable locally at a prime p, then it has no global solution. If the curve has arisen in the process of 2-descent, it can then be

removed from consideration, along with the other curves in the same coset.

While proving that a quadric-intersection has no points is (in context) fairly easy, showing that there is a point is somewhat harder. This is down to the failure of local-global principle for curves of genus 1: A QI may be everywhere locally solvable (ELS) but yet still not have a global solution. The only option left for us currently is to show a QI has a point by explicitly finding one. This is the subject of chapter 8 which presents an algorithm for searching for points on QIs over function fields.

Although the problem is the local-solvability of quadric intersections, it can be transformed fairly easily into the problem of deciding local-solvability of quartics. This step is used because existing methods for quartics over $\mathbb{Q}$ can be adapted to work for quartics over function fields, and the extra time taken in converting a QI to a quartic is negligible.

Definition. A homogeneous quartic of the form

$$f(U,V) = a_0U^4 + a_1U^3V + a_2U^2V^2 + a_3UV^3 + a_4V^4$$

is said to be everywhere locally solvable if the corresponding curve defined by $Y^2 = f(U,V)$ is non-singular and ELS.

Definition. A homogeneous quartic $f(U,V)$ over a field F, in the form given above is said to be solvable if there exists a point $(u:v)$ in $\mathbb{P}^1$ such that $f(u,v)$ is a square in F.

Remark. Local solvability is automatic at places of good reduction. This is because the reduced curve is a genus 1 curve over a finite field and therefore has a point by Hasse's Theorem. This point lifts to a local point by Hensel's Lemma. We only need therefore check local solvability at places of bad reduction, i.e. primes dividing the discriminant of f.

6.1.1 From QIs to quartics

To begin with, it must be noted that in the broader view of the main descent algorithm, the conversion from QI to quartic only needs to be carried out once the QI has passed the conic solving stage. Therefore all QIs we deal with now are solvable as pairs of conics and can be parametrized.

Let $\mathcal{C}$ be a quadric-intersection in $\mathbb{P}^3$ given by the following pair of equations:

$$C: \quad \begin{cases} f_1(W,X,Z) &= 0 \\ f_2(X,Y,Z) &= 0 \end{cases} \tag{6.1}$$

where f_1 and f_2 are homogeneous quadrics in the variables indicated. Since each only involves three of the four variables, they can be thought of as conics, albeit conics in different ambient spaces.

Let $\mathcal{C}_1$ be the conic in $\mathbb{P}^2$ defined by f_1 and $\mathcal{C}_2$ that defined by f_2.

Step 1

Solve and parametrize $\mathcal{C}_1$ to obtain the following equations:

$$\begin{aligned} W &= q_1(U,V) \\ X &= q_2(U,V) \\ Z &= q_3(U,V) \end{aligned} \tag{6.2}$$

where q_1, q_2 and q_3 are homogeneous quadratics in the variables U and V.

Step 2

We now substitute q_2 and q_3 for X and Z respectively in f_2:

$$f_2\big(q_2(U,V), Y, q_3(U,V)\big) = 0. \tag{6.3}$$

Because of the nature of the quadric intersections we are interested in, the only term involving Y in f_2 is cY^2 where c is a non-zero constant. Therefore

subtracting cY^2 from both sides of the equation and dividing by $-c$ we end up with

$$Y^2 = f(U, V) \tag{6.4}$$

where f is a homogeneous quartic in U and V. We can now examine the local-solvability of f.

6.2 The Local solvability algorithm

The algorithm given in this section is a modification of one presented by Merriman, Siksek and Smart [MSS] that tests local solvability of quartics over $\mathbb{Q}$. The basic steps of the algorithm may be applied to quartics over function fields using the same arguments given in [MSS]. There are however a couple of differences. The quartics in [MSS] are in one variable, whereas we use homogeneous two-variable quartics $f(U, V)$ (so the curve $Y^2 = f(U, V)$ is in fact in a weighted projective space). Because of this first change, we need a different method to stop the algorithm from recursing infinitely to that used in [MSS]. This is done in section 6.2.1.

Once we have an equation of the form $Y^2 = f(U, V)$ where f is a homogeneous quartic over a function field $\mathbb{F}_q(T)$, we can test it for local solvability at a prime p using the algorithm described in this section. As in chapter 1, we denote by $k_\mathfrak{p}$ the field $\mathbb{F}_q[T]/(p)$. Recall that $k_\mathfrak{p}$ is isomorphic to $\mathbb{F}_{q^{\deg(p)}}$.

The first step is to clear denominators of f. This is done by multiplying by the *square* of the LCM of the denominators of the coefficients of f. This square term is then absorbed into the Y^2 term on the left-hand-side of the equation via the change of variables $Y' = g^2 Y$. Once denominators have been cleared from f, any remaining common square factors of the coefficients of f can be similarly

absorbed into the Y^2 term. We assume from now on that these steps have been performed and f now looks like

$$f(U,V) = a_0U^4 + a_1U^3V + a_2U^2V^2 + a_3UV^3 + a_4V^4 \tag{6.5}$$

where $a_i \in \mathbb{F}_q[T]$ for all i and HCF$(\{a_i : 0 \leq i \leq 4\})$ is square-free. Let $\bar{f}(U,V)$ denote the reduced quartic f mod p. The remainder of the algorithm runs as follows:

Step 1

Test whether f is divisible by p. If this is the case then proceed to **step 2**. Otherwise go to **step 3**.

Step 2

We now have f divisible by p but not by p^2. Let $f_1 = \frac{1}{p}f$. We now examine the roots of the polynomial $\bar{f}_1$ (the reduction of f_1 mod p):

If f is locally solvable at p then so is the equation $pY^2 = f_1(U,V)$. Reducing this mod p we see that $\bar{f}_1(U,V)$ must have a root. So if $\bar{f}_1$ has no roots then f is not locally solvable and we **stop**. Otherwise:

Step 2a

If $\bar{f}_1$ has a non-repeated root then f is locally solvable and we **stop**. Otherwise $\bar{f}_1$ has either 1 or 2 repeated roots $(u_i : v_i)$ $(i = 1, 2)$. It is then a case of checking whether the following is locally solvable:

$$f_2 = \begin{cases} \frac{1}{p^2}f(pU + \epsilon_i V, V) & \text{if } v_i \neq 0 \\ \frac{1}{p^2}f(U, pV) & \text{if } v_i = 0 \end{cases} \tag{6.6}$$

where $\epsilon_i = u_i/v_i$. This needs to be checked for each i if there are two roots. It is necessary and sufficient that one of these new quartics be locally solvable. To check this we recurse to **step 1** with the new function f_2 instead of f.

Step 3

At this stage, f is not divisible by p. If the reduced polynomial $\bar{f}$ is a square then we have local solvability and **stop**.
If $\bar{f}$ *cannot* be written in the form $\bar{f} = \bar{a}\bar{g}^2$ where $\bar{g} \in k_{\mathfrak{p}}[U, V]$ and $\bar{a} \in k_{\mathfrak{p}}$ then f is locally solvable and we **stop**. Otherwise proceed to **step 4**.

Step 4

Now $\bar{f} = \bar{a}\bar{g}^2$ with $\bar{a} \neq 0$ not a square. Therefore any local solution $(y, [u : v])$ to $Y^2 = f(U, V)$ must satisfy $y \equiv 0 \pmod{p}$ and $\bar{g}(u, v) \equiv 0 \pmod{p}$. If $\bar{g}$ has no roots in $k_{\mathfrak{p}}$ then f not locally solvable at p and we **stop**. Otherwise proceed to **step 5**.

Step 5

Now $\bar{g}$ has two roots $(u_1 : v_1)$ and $(u_2 : v_2)$. Write $f = ag^2 + ph$ where a and g reduce to $\bar{a}$ and $\bar{g}$ respectively. If neither of $(u_i : v_i)$, $i = 1, 2$ is a root of $\bar{h}$ then f is not locally solvable and we **stop**. Otherwise proceed to **step 6**.

Step 6

It now remains to check whether either of the following two quartics f_i is locally solvable.

$$f_i = \begin{cases} \frac{1}{p^2} f(pU + \epsilon_i V) & \text{if } v_i \neq 0 \\ \frac{1}{p^2} f(U, pV) & \text{if } v_i = 0 \end{cases} \tag{6.7}$$

where $\epsilon_i = u_i/v_i$. If one of these is locally solvable then so is f and we **stop**. Otherwise f is not locally solvable and we **stop**. Therefore we recursively repeat **step 1** for each of the two f_i.

6.2.1 Some computational remarks

The local-solvability algorithm contains several steps where the whole program must be recursively called from the top. In a number of these cases it is possible (depending on how certain polynomials factorize) that a recursive call must be made on two input quartics to determine the result. Each such call of the function may in turn split further as the algorithm is repeated worked through (a split recursion). From a programming perspective we handle this as follows:

Within the program there is a function that runs through *one* call of the algorithm until one of the following happens: 1) Local solvability is decided for *the particular quartic being tested at that point*, or 2) A recursion occurs (split or not). If a recursion occurs, the transformation matrix changing the variables U and V is stored. This function, which we will here call R1 (for "one recursion") is then applied to the original quartic q being tested. The object returned by R1 is a 3-tuple of the form (i, M, f) where $i = 1$ if q is locally solvable, $i = -1$ if not and $i = 0$ if we are undecided and therefore need a recursion. If $i = 0$ then there may be a split recursion, in which case we return a sequence of two such 3-tuples. M is the cumulative transformation matrix, the matrix needed to

transform q into f, the quartic to be tested for the next recursion (if necessary). Note that if $i \neq 0$ we do not need M or f. The function R1 is then repeatedly called on the sequence of 3-tuples. If $i = 1$ for any element of the sequence then we have local solvability and stop. If $i = -1$ then that element remains unchanged for every call of R1. If $i = 0$ the element is replaced by either one or two "recursion" elements as described above. This is repeated until we have $i = 1$ occurring anywhere, or $i = -1$ for all elements of the sequence, in which case q is not locally solvable.

Avoiding infinite recursion

It is not clear from the above that this algorithm will stop. In fact it is possible (and such cases can be constructed) for it to repeat an infinite number of times. To avoid this we need to keep an eye on the cumulative transformation matrix M.

First we note that if there exists u, v and $y \in R = \mathbb{F}_q[T]$ such that $y^2 = f(u, v)$ in R, we can assume that p does not divide both u and v; for if this were the case then we can replace the solution $(u : v : y)$ with $(p^{-1}u : p^{-1}v : p^{-2}y)$. At each recursive call of the algorithm we have a 2×2 matrix M with entries in R, and $\det(M) = p^e$ for some e. When we make a recursive call, we are checking whether a new equation $Y_1^2 = f_1(U_1, V_1)$ has a solution. If it does have a solution (u_1, v_1), then from the definition of M, this solution satisfies the equation:

$$M \begin{bmatrix} u_1 \\ v_1 \end{bmatrix} = \begin{bmatrix} u \\ v \end{bmatrix} \tag{6.8}$$

. So if $M \equiv 0 \pmod{p}$ then $u \equiv v \equiv 0 \pmod{p}$ - a contradiction. Therefore if during the algorithm M is identically zero modulo p, we do not have local

solvability.

6.2.2 Example

This is a worked example of applying the local-solvability algorithm to a quartic.

Let f be the following quartic over $\mathbb{F}_7(T)$:

$$f = (3T^6+T^5)U^4+(5T^5+5T^3)U^2V^2+(2T^6+5T^5+3T^4+5T^3+2T^2+2T)V^4. \tag{6.9}$$

This quartic has I and J invariants:

$$I = 2T^{12}+T^{11}+4T^{10}+6T^9+5T^7,$$
$$J = 4T^{17}+2T^{16}+6T^{15}+4T^{13}+T^{12}+3T^{11}+3T^{10}+T^9,$$

and discriminant:

$$\Delta = 4T^{18}(T+1)^2(T+3)^2(T+4)^2(T+5)(T^3+3T^2+4)^2(T^5+6T^4+5T^3+6T^2+T+1).$$

We only need to check local solvability at the "bad" places of f - the prime factors of Δ above. Here we check local solvability at $p = T$.

Firstly we note that the coefficients of f are all polynomials and have no square common factors, so we are able to call the algorithm with f as given above and $p = T$.

(1) Immediately we see that f is divisible by p so we proceed straight to step 2.

(2) Let $f_1 = \frac{1}{p}f$ and $\bar{f_1} = f_1$ reduced mod p. Reduction mod T is equivalent to evaluating the coefficients of f at $T = 0$ (recall the residue field $\mathbb{F}_q(T)/(p)$ is isomorphic to $\mathbb{F}_{q^{\deg(p)}}$), so we obtain $\bar{f_1} = 2V^4$, which has one repeated root $[1:0]$.

(3) The repeated root of $\bar{f}_1$ is at infinity so we repeat step 1 with the new function $f_2 = \frac{1}{p^2} f(U, TV)$. Re-name $f_2 = f$. The matrix of the (cumulative)transformation of variables is $M = \begin{pmatrix} 1 & 0 \\ 0 & T \end{pmatrix}$ and the new quartic $f = (3T^4 + T^3)U^4 + (5T^5 + 5T^3)U^2V^2 + (2T^8 + 5T^7 + 3T^6 + 5T^5 + 2T^4 + 2T^3)V^4$.

(4) Now f is divisible by p^2, so we replace f by $\frac{1}{T^2} f$ and start from step 1 again. This time $f = (3T^2 + T)U^4 + (5T^3 + 5T)U^2V^2 + (2T^6 + 5T^5 + 3T^4 + 5T^3 + 2T^2 + 2T)V^4$ and the matrix M is unchanged.

(5) f is divisible by p so we go to step 2 again. Dividing by p and reducing mod p gives $\bar{f}_1 = U^4 + 5U^2V^2 + 2V^4$ which factorizes as $\bar{f}_1 = (U^2 + UV + 3V^2)(U^2 + 6UV + 3V^2)$. So $\bar{f}_1$ has no roots in K. Hence f (and thus our original f) is not locally solvable at $p = T$. We do not need to check the other bad places of f now - the result at T is enough to show that f is not everywhere locally solvable.

Chapter 7

Polynomial lattices

This chapter introduces the notion of polynomial lattices and what is meant by a reduced basis for such lattices. The first work in this field was done by A. Lenstra in 1985 [Len]. Lenstra's definition of a reduced basis appeared independently in [MuSt] under the name "weak Popov form". The paper gives an algorithm for converting lattice basis-matrices over a polynomial ring into weak Popov form, and it is these definitions and methods we describe in this chapter.

7.1 Lattice-reduction for polynomial matrices - weak Popov form

7.1.1 Weak Popov form

A key part of the point-searching algorithm is the reduction of polynomial matrices to *weak Popov form*. A more thorough definition of weak Popov form is given in [MuSt], but may basically be summarised as follows:

Definition. A matrix with polynomial entries is said to be in weak Popov form if the rightmost entries of maximum degree in each row all lie in different columns.

One could give an equally valid definition by replacing "rightmost" with "leftmost" above. We use the former, as in [1].

Example.

$$M = \begin{pmatrix} 4t+7 & 4t+5 & 10t^4+3t^3+5t^2+2t+4 \\ 4t^4+3t^3+t^2+10t+10 & 8 & 5 \\ 4t^3+8t^2+4t+7 & 9t+8 & 4t^2+8t+4 \end{pmatrix}. \tag{7.1}$$

The above matrix, over $\mathbb{F}_{11}[t]$ is not in weak Popov form, since the rightmost entries for rows 2 and 3 both lie in the first column. However, by applying a finite sequence of elementary row operations we can reduce it to the following matrix, which is in the required form:

$$N = \begin{pmatrix} 8t^3+9t+4 & 5t^3+10t^2+3t+7 & 7t^2+4t+8 \\ 6t^3+8t^2+3t+10 & 2t^2+3t+8 & 7t^3+3t^2+7t+5 \\ 4t^3+8t^2+4t+7 & 9t+8 & 4t^2+8t+4 \end{pmatrix}. \tag{7.2}$$

Since every elementary row operation can be thought of as (left) multiplication by a matrix, we have $N = TM$. In this case, T is the following:

$$T = \begin{pmatrix} 1 & 8t+3 & 3t^2+7t \\ 0 & 1 & 10t \\ 0 & 0 & 1 \end{pmatrix}. \tag{7.3}$$

Four row operations were used by our algorithm in this example.

The use of weak Popov form matrices in our algorithm relies on a certain property of such matrices. This is a stronger version of a lemma in [MuSt]

(lemma 8.1). First we establish some notation:

P_i^M	the ith pivot-element of the matrix M,
$\deg(r)$	the maximum of the degree of the elements of the row r,
F	The base field of the polynomials - elements of M lie in $F[t]$.

P_i^M is the rightmost element of maximum degree in row i. We can now state the key lemma:

Lemma 9. *Let M be an $m \times n$ matrix in weak Popov form with no rows all zero. Let $\mathbf{r}_i \quad (1 \leq i \leq m)$ denote the ith row of M. If $\mathbf{r} = \sum_{i=1}^{m} d_i \mathbf{r}_i$ lies in the $F[t]$-span of the rows $\mathbf{r}_i$ then:*

$$\deg(\mathbf{r}) = \max_{1 \leq i \leq m} \{\deg(d_i) + \deg(P_i^M) \mid d_i \neq 0\}. \tag{7.4}$$

To prove this lemma, we require the following two sub-lemmas:

Lemma 10. *Let M be a $2 \times n$ matrix in weak Popov form with no rows all zero. Let $\mathbf{r}_1$ and $\mathbf{r}_2$ denote the rows of M. For $d_i \in F[t]$ not all zero $(i = 1, 2)$ the following holds:*

$$\deg(d_1\mathbf{r}_1 + d_2\mathbf{r}_2) = \max_{i=1,2}\{\deg(d_i) + \deg(\mathbf{r}_i) \mid d_i \neq 0\}. \tag{7.5}$$

Lemma 11. *Let M be an $m \times n$ matrix in weak Popov form. Let $d_i \quad (2 \leq i \leq m)$ be $m - 1$ elements of $F[t]$. Then the $2 \times n$ matrix consisting of the rows $\mathbf{r}_1$ and $\sum_{i=2}^{m} d_i \mathbf{r}_i$ is in weak Popov form, where $\mathbf{r}_i$ are the rows of M.*

Lemma 10 is simply Lemma 9 in the specific case $m = 2$.

Proof of Lemma 10

The statement clearly holds if either d_1 or d_2 is zero. Assume now that neither is zero.

Let $\mathbf{r}_1^* = d_1\mathbf{r}_1$, $\mathbf{r}_2^* = d_2\mathbf{r}_2$. Denote by M^* the matrix with rows $\mathbf{r}_1^*$ and $\mathbf{r}_2^*$. M^* is in weak Popov form since multiplication by a non-zero element of $\mathbb{F}[t]$ does not change the pivot index of a row. Without loss of generality $\deg(\mathbf{r}_1^*) \geq \deg(\mathbf{r}_2^*)$. Let $\deg(\mathbf{r}_1^*) = d$.

We wish to prove that $\deg(\mathbf{r}_1^* + \mathbf{r}_2^*) = d$. If $\deg(\mathbf{r}_2^*) < d$ then the element of $\mathbf{r}_2^*$ in the same column as $P_1^{M^*}$ has degree $< d$. Therefore there exists an element of $\mathbf{r}_1^* + \mathbf{r}_2^*$ of degree d, and $\deg(\mathbf{r}_1^* + \mathbf{r}_2^*) \geq d$. If $\deg(\mathbf{r}_2^*) = d$ then let $I = \max\{I_1^{M^*}, I_2^{M^*}\}$ where I_i^M denotes the pivot index of the ith row of M. The Ith element of $\mathbf{r}_1^* + \mathbf{r}_2^*$ has degree d, therefore $\deg(\mathbf{r}_1^* + \mathbf{r}_2^*) \geq d$.
But it is also true that $\deg(\mathbf{r}_1^* + \mathbf{r}_2^*) \leq d$ since:

$$\begin{aligned}\deg(\mathbf{r}_1^* + \mathbf{r}_2^*) &= \max\{\deg(a_i + b_i) : a_i \in \mathbf{r}_1^*, b_i \in \mathbf{r}_2^*\} \\ &\leq \max \bigcup_{i=1}^{n} \{\deg(a_i), \deg(b_i)\} = \max\{\deg(\mathbf{r}_1^*), \deg(\mathbf{r}_2^*)\} = d.\end{aligned}$$

Therefore $\deg(\mathbf{r}_1^* + \mathbf{r}_2^*) = d$ and the lemma is proved.

Proof of Lemma 11

Since multiplication of a row by a scalar ($\in F[t]$) does not change the position of its pivot element (its *pivot index*) it is required to prove that if two rows have pivot index I_1 and I_2 where $I_1 \neq I_2$ then the pivot index of their sum is either I_1 or I_2. The lemma then follows by induction.
Let $\mathbf{r}_1$ and $\mathbf{r}_2$ be the two rows. Without loss of generality we can assume $\deg(\mathbf{r}_2) \leq \deg(\mathbf{r}_1) = d$. We now consider two cases:

Case 1: $\deg(\mathbf{r}_2) < \deg(\mathbf{r}_1)$

Let $\mathbf{r}_1 = [x_1, \ldots, x_n]$, $\quad \mathbf{r}_2 = [y_1, \ldots, y_n]$. Let $k = I_1$, so x_k has degree d and is a pivot element. Since $\deg(\mathbf{r}_2) < \deg(\mathbf{r}_1) = d$ it follows that $\deg(x_k + y_k) = d$.

Now if $i > k$ then $\deg(x_i + y_i) \leq \max\{\deg(x_i), \deg(y_i)\} < d$; while if $i < k$ then $\deg(x_i + y_i) \leq \max\{\deg(x_i), \deg(y_i)\} \leq d$. Hence $x_k + y_k$ is a pivot element.

Case 2: $\deg(\mathbf{r}_2) = \deg(\mathbf{r}_1)$

Let x_k be the rightmost pivot element, so $\deg(x_k) = d$, then we have $\deg(x_k + y_k) = d$. Now if $i > k$ then $\deg(x_i + y_i) \leq \max\{\deg(x_i), \deg(y_i)\} < d$; while if $i < k$ then $\deg(x_i + y_i) \leq \max\{\deg(x_i), \deg(y_i)\} \leq d$. Hence $x_k + y_k$ is a pivot element.

We are now able to complete the proof of Lemma 9.

Proof of Lemma 9

Let $\mathbf{r} = \sum_{i=1}^{m} d_i \mathbf{r}_i$. Let k be an integer such that $1 \leq k \leq m$ and $d_k \neq 0$. Without loss of generality $k = m$. Let $\mathbf{r}^* = \sum_{i=1}^{m-1} d_i \mathbf{r}_i$. If we assume for induction that Lemma 9 holds for $l \times n$ matrices with $l < m$ then

$$\deg(\mathbf{r}^*) = \max\{\deg(d_i) + \deg(P_i^M) \mid d_i \neq 0\}.$$

Let M^* be the matrix with first row $\mathbf{r}^*$ and second row $\mathbf{r}_m$. By Lemma 9, M^* is in weak Popov form. Hence by applying Lemma 10 to M^* it follows that

$$\deg(\mathbf{r}) = \max\{\deg(\mathbf{r}^*), \deg(d_m) + \deg(\mathbf{r}_m)\} = \max\{\deg(\mathbf{r}^*), \deg(d_m) + \deg(P_m^M)\}$$

$$= \max_{1 \leq i \leq m} \{\deg(d_i) + \deg(P_i^M) \mid d_i \neq 0\}. \qquad \square$$

7.1.2 Algorithm for transforming matrices to weak Popov form

A more detailed account of the weak Popov form algorithm than is necessary here is given in [MuSt]. The key point is that a matrix with polynomial entries

can be transformed into weak Popov form by application of a finite number of *elementary row operations*. An elementary row operation is one of the following:

(i) swapping two rows,

(ii) adding a scalar (polynomial) multiple of one row to another.

The first type is not needed because swapping two rows clearly makes no difference to whether a matrix is in weak Popov form or not. Only the second type of row operation is required. Such row operations can be accomplished by (left) multiplication by a unimodular transformation matrix as follows:

Let M be an $m \times n$ matrix with entries in the ring R. The matrix obtained by adding x times the ith row to the jth row of M ($i \neq j$) is TM where $t_{j,i} = x$, all diagonal entries are 1 and all other entries zero. The determinant of such matrices is always 1.

The algorithm proceeds as follows (omitting details of how to decide which row operation to use at each step):

ALGORITHM 7.1.1: **Weak Popov Form**

INPUT: M: an $m \times n$ matrix defined over $F[t]$.

OUTPUT: N, T where N is an $m \times n$ matrix in weak Popov form and $N = TM$.

1. N:= copy(M);
2. T:= $I_{n \times n}$ (identity matrix);
3. WHILE N is not in weak Popov form DO
4. U:= unimodular matrix of appropriate row operation;
5. N:= UN; T:= UT;
6. END WHILE;
7. RETURN N, T;

This algorithm has been fully implemented in MAGMA .

7.2 Polynomial Lattices

Since the main part of the search for points on quadric intersections (both over $\mathbb{Z}$ and over $\mathbb{F}_q[T]$) involves the construction and analysis of many lattices, it is important to first establish what we mean by a lattice.

7.2.1 Lattices over $\mathbb{Z}$

A *lattice in* $\mathbb{Z}^n$ is a (free) $\mathbb{Z}$-submodule of $\mathbb{Z}^n$. More generally a lattice in $\mathbb{R}^n$ is a discrete $\mathbb{Z}$-submodule of $\mathbb{R}^n$. The *rank* of a lattice L is the minimal number of generators of L. If a set of vectors $B = \{\mathbf{b}_1, \ldots, \mathbf{b}_r\}$ are $\mathbb{Z}$-linearly independent and generate L, then L has rank r and B is said to be a *basis* of L. If $r = n$ then L is said to have full rank. Associated to the lattice basis B is its *basis matrix*. This is the $r \times n$ matrix whose rows are the basis vectors $\mathbf{b}_1, \ldots, \mathbf{b}_r$.

A lattice L_2 is called a *sublattice* of a lattice L_1 (notation: $L_2 \subseteq L_1$) if L_2 is a free $\mathbb{Z}$-submodule of L_1.

When working with lattices in $\mathbb{Z}^n$ to find points on quadric-intersections it is necessary to perform LLL-reduction on the lattice basis. The reduced basis has special properties which allow us to search for points by looking at "short" vectors in the lattice.

7.2.2 Lattices over $F[T]$

Just as we may define a lattice in $\mathbb{R}^n$ we may also formulate an equivalent definition of a lattice in K^n where K is a function field (the field-of-fractions

of a polynomial ring $F[T]$ where F is a field). In this case the polynomial ring $F[T]$, which we denote by R, plays the role of $\mathbb{Z}$.

As in the real case we can think of K^n as an R-module. We define a lattice as follows:

A lattice in $F(T)^n$ (resp. $F[T]^n$) is a free $F[T]$-submodule of $F(T)^n$ (resp. $F[T]^n$).

Rank, basis and sublattice are defined analogously to the $\mathbb{Z}$ case.

There is no direct analogue of LLL-reduction in this case. However, for sub-lattices of R^n (i.e. lattices whose basis matrix has polynomial entries) the basis-matrix can be transformed into weak Popov form. When searching for points on quadric-intersections, reducing a lattice basis to weak Popov form gives it the necessary properties for a short-vector point search. This will be explained in more detail later, in section 8.6.

Chapter 8

Searching for points on an intersection of two quadrics

This chapter will describe a method used for searching for points (up to a certain degree bound) on quadric intersections over function fields. The method is in two main stages. Firstly, for each point P on the curve modulo some irreducible p, the construction of a lattice in which lie all points congruent to P mod p. Secondly, these lattices are reduced to weak Popov form, and using Lemma 1 we search for small vectors in the lattice.

8.1 Lattice methods over $\mathbb{Q}$

The algorithm described in this chapter comes from various similar methods applied to curves over $\mathbb{Q}$, that can be adapted to work for curves over function fields by using the notion of lattice basis reduction described in the previous chapter.

The basic idea is, given a curve $\mathcal{C}$ in $\mathbb{P}^3$ and a prime p such that $\mathcal{C}$ has good-reduction at p, to list all points $\tilde{P}$ on $\mathcal{C}$ mod p and for each $\tilde{P}$ to construct a lattice containing all points on $\mathcal{C}$ congruent to $\tilde{P}$ mod p. We then search for short vectors in the lattice to find points on $\mathcal{C}$ of height up to a bound H. This was first suggested to John Cremona by Roger Heath-Brown in 1999 [H-B].

Let $\mathbf{x}$ be a lift of $\tilde{P}$ to $\mathbb{Z}^4$ and $\mathbf{e}_1, \ldots, \mathbf{e}_4$ be the unit vectors of $\mathbb{Z}^4$. The various lattice-methods formulated so far are as follows:

The mod p lattice

The set of all points in $\mathbb{Z}^4$ that reduce to $\mathbf{x}$ mod p forms a lattice of index p^3 with basis $\{\mathbf{x}, p\mathbf{e}_2, p\mathbf{e}_3, p\mathbf{e}_4\}$. This lattice is not useful as we would need $p > cH^{4/3}$ for some constant c, giving a run-time of $O(H^{4/3})$.

The mod p^2 lattice

This method was used by Tom Womack in his thesis [Wom].
Lift $\mathbf{x}$ to a solution mod p^2. The set of all points in $\mathbb{Z}^4$ that reduce to $\mathbf{x}$ mod p^2 forms a lattice of index p^5 with basis $\{\mathbf{x}, p\mathbf{u}, p^2\mathbf{e}_3, p^2\mathbf{e}_4\}$ where $\mathbf{u} = (0, 1, u_3, u_4)$ is a solution to $\mathcal{C}$ mod p. All points in this lattice are solutions mod p^2. This method has a run time of $O(H^{4/5})$.

Heath-Brown's method

Lift $\mathbf{x}$ in all possible ways to solutions $\mathbf{y}$ mod p^2. There are exactly p of these lifts for each $\mathbf{x}$ so we get p lattices for each $\mathbf{x}$. The next step is to lift each $\mathbf{y}$ to a solution mod p^4. Then the set of all points in $\mathbb{Z}^4$ that reduce to $\mathbf{y}$ mod p^3 forms a lattice of index p^{10} with basis $\{\mathbf{x}, p^2\mathbf{u}, p^4\mathbf{e}_3, p^4\mathbf{e}_4\}$ where $\mathbf{u} = (0, 1, u_3, u_4)$ is a solution to $\mathcal{C}$ mod p^2. All points in this lattice are solutions to $\mathcal{C}$ mod p^4. Here

we need $p > cH^{2/5}$ but there are $O(p^2)$ lattices rather than $O(p)$ so the run time is $O(H^{4/5})$.

Watkins' method

Watkins constructs a lattice of index p^6 which he claims [Wat] to be precisely of the set of lifts of $\mathbf{x}$ that are solutions to $\mathcal{C}$ mod p^3. This is not true, however we only need that the lattice contains all such solutions which it indeed does. Using this lattice gives a run time of $O(H^{2/3})$.

Watkins calls this method the "Elkies ANTS-IV Algorithm", referring to a paper by Elkies [Elk]. Elkies' method uses real approximations rather than p-adic ones.

Long's method

Rachel Long's approach [Lon] is to construct a lattice for each mod p^2 solution like Heath-Brown, but in this case the lattices have index p^{12} and consist of solutions to $\mathcal{C}$ mod p^6 that are congruent mod p^2 to one of the mod p^2 solutions. We need $p > cH^{1/3}$ here but there are $O(p^2)$ lattices as in Heath-Brown's method so the run time is $\mathcal{O}(H^{2/3})$. Replacing p^2 by p in Long's algorithm constructs the same lattice as that given by Watkins.

8.2 The method of undetermined coefficients

The following is a fairly "brute-force" method of finding points on curves over function fields up to a certain height bound, but it will always work in theory. The basic idea is as follows:

Suppose we have a projective curve $\mathcal{C}$ defined over a function field $F(T)$.

After choosing a suitable model for $\mathcal{C}$ we have a set of equations defining $\mathcal{C}$:

$$\begin{cases} f_1(X_1, \ldots, X_m) = 0 \\ f_2(X_1, \ldots, X_m) = 0 \\ \vdots \\ f_n(X_1, \ldots, X_m) = 0 \end{cases} \tag{8.1}$$

where f_i are homogeneous of equal degree.

We define the *degree* of a point $P \in \mathcal{C}(F(T))$ as follows: $\deg(P)$ is the maximum of the degrees of the polynomials $g_i(T)$ where P is written as $(g_1(T) : \ldots : g_m(T))$ and $g_1, \ldots, g_m$ are coprime.

Suppose we are searching for points $(x_1 : \ldots : x_m) \in \mathcal{C}(F(T))$ of degree no greater than d. Then we may assume that such points have coordinates in $\mathbb{F}[T]$ since $\mathcal{C}$ is a projective curve. We introduce $m(d+1)$ new variables a_{ij} for $1 \leq i \leq m$, $0 \leq j \leq d$ and make the substitution:

$$\mathbf{X} = \left(\sum_{j=0}^{d} a_{1j} T^j, \sum_{j=0}^{d} a_{2j} T^j, \ldots, \sum_{j=0}^{d} a_{mj} T^j \right). \tag{8.2}$$

Substituting this into the polynomials in (8.1) we get a set of polynomials $F_1(T), \ldots, F_n(T)$ where $F_i(T) = f_i(\mathbf{X})$. For $\mathbf{X}$ to satisfy (8.1) these polynomials must all be identically zero so their coefficients must all be zero. If S is the set of all coefficients of $F_1, \ldots, F_n$ then S is a set of polynomials s_i in the variables a_{ij}. Setting these all to zero now gives us a set of equations defining a variety $\mathcal{V}$ over the field F:

$$\begin{cases} s_1(a_{11}, \ldots, a_{1d}, a_{21}, \ldots, a_{2d}, \ldots, a_{m1}, \ldots, a_{md}) = 0 \\ s_2(a_{11}, \ldots, a_{1d}, a_{21}, \ldots, a_{2d}, \ldots, a_{m1}, \ldots, a_{md}) = 0 \\ \vdots \\ s_k(a_{11}, \ldots, a_{1d}, a_{21}, \ldots, a_{2d}, \ldots, a_{m1}, \ldots, a_{md}) = 0 \end{cases} \tag{8.3}$$

where $k = \#S$. All one has to do now is find points on $\mathcal{V}$. $\mathcal{V}$ will in general have positive dimension and consist of points, curves, surfaces etc. The irreducible components of dimension 0 will map to degree d points on $\mathcal{C}$, dimension 1 components to degree $d-1$ points etc. The various methods of finding the components of $\mathcal{V}$ all involve tedious computation using Groebner bases, and most approaches via computer stall in an explosion of variables and equations around the $d = 5$ mark. Thankfully, the only occasions we need to use this method are for $d = 1, 2$ and these are described in 8.7. In this case F is a finite field so MAGMA can be asked straight for the points on $\mathcal{V}$.

One advantage of this method is that it does not assume anything about $\mathcal{C}$ beforehand.

8.3 Quadric Intersections

Definition. Let K be a field. A quadric surface over K is a surface in $\mathbb{P}^3(K)$ given by an equation in the following form:

$$aW^2+bX^2+cY^2+dZ^2+eWX+fWY+gWZ+hXY+iXZ+jYZ=0 \quad (8.4)$$

with the coefficients a to j in K.

If V is the vector of variables $[W, X, Y, Z]$ then provided K has characteristic $\neq 2$, the above equation can be written as $VMV^T = 0$ where M is the following symmetric matrix:

$$M = \begin{pmatrix} 2a & e & f & g \\ e & 2b & h & i \\ f & h & 2c & j \\ g & i & j & 2d \end{pmatrix}. \quad (8.5)$$

Note that this matrix is far from unique, even up to scalar multiplication. For example, it is sometimes equally convenient to write it as an upper-triangular matrix over K.

Definition. A Quadric Intersection over a field K is a curve in $\mathbb{P}^3(K)$ given by a pair of independent equations in the same form as (8.4), or equivalently by two symmetric matrices with entries in K.

Definition. A quadric intersection $\mathcal{C}$ is said to have *good reduction* at a prime p if the associated quartic given by $\det(uA - vB)$ has no repeated roots mod p, where A and B are the matrices defining $\mathcal{C}$.

8.4 Finding all solutions mod p

First we establish the following notation:

$$\begin{aligned} R &: \text{the polynomial ring } \mathbb{F}_q[t] \text{ where } 2 \nmid q, \\ K &: \mathbb{F}_q(t) \text{ - the field-of-fractions of } R, \\ p &: \text{an irreducible in } R \text{ of degree } \geq 2, \\ \mathfrak{p} &: \text{the prime ideal of } R \text{ generated by } p, \\ k_{\mathfrak{p}} &: \text{the quotient-ring } R/\mathfrak{p}. \end{aligned}$$

Let $\mathcal{C}$ be a quadric intersection in $\mathbb{P}^3(K)$ with defining polynomials F_a and F_b, that has good reduction at p. Let $\tilde{F}_a$ and $\tilde{F}_b$ denote the polynomials derived from F_a and F_b by reducing the coefficients mod p. Let $\tilde{\mathcal{C}}$ denote the reduced curve in $\mathbb{P}^3(k_{\mathfrak{p}})$.

Note that since we are working with rings over finite fields, the quotient-ring $k_{\mathfrak{p}}$ is itself a finite field of size $q^{\deg(p)}$. It follows that $\tilde{\mathcal{C}}$ has a finite number of points.

Points on $\tilde{\mathcal{C}}$ can be grouped into four types: $(0:0:0:1)$, $(0:0:1:\star)$, $(0:1:\star:\star)$ and $(1:\star:\star:\star)$ where $\star$ denotes an element of $k_{\mathfrak{p}}$. For the first type, it is a trivial matter to check whether $(0:0:0:1)$ lies on $\tilde{\mathcal{C}}$.

For points of the second type we introduce a new variable u and evaluate $\tilde{F}_a$ and $\tilde{F}_b$ at $(0:0:1:u)$. This gives two quadratic polynomials in u and points on $\tilde{\mathcal{C}}$ of the second type correspond to common roots of these polynomials. A similar method is used for points of the third type, except this time two variables u and v are needed, and we solve two simultaneous quadratics in u and v, with solutions corresponding to points on $\tilde{\mathcal{C}}$.

For points of the fourth type, we cannot extrapolate this method further since we would have three variables but only two equations. Instead we fix the second coordinate to a particular value in $k_{\mathfrak{p}}$ and proceed as for points of the third type. This must be repeated for every element of $k_{\mathfrak{p}}$, of which there are $q^{\deg(p)}$.

The computer algebra package MAGMA has a built in function for finding all points on a variety defined over a finite field. However, for quadric intersections at least, the above method is much faster. The following is a summary of this method, notation as stated above:

ALGORITHM 8.4.1: **Points mod p**

`INPUT:` $\mathcal{C}$: `A quadric intersection over` $\mathbb{F}_q(t)$, p: `an irreducible in` $\mathbb{F}_q[t]$

`OUTPUT: A list of points on` $\mathcal{C}$ `mod` p

1. $\tilde{\mathcal{C}}$`:= Reduce` $\mathcal{C}$ `mod` p`;`

   $\tilde{F}_a, \tilde{F}_b$`:= Defining polynomials of` $\tilde{\mathcal{C}}$`;`

   S`:= empty-set;`

2. `IF` $\tilde{F}_a(0,0,0,1) = 0$ `AND` $\tilde{F}_b(0,0,0,1) = 0$ `THEN`

3. S:= $S \cup \{(0:0:0:1)\}$;

4. END IF;

5. IF $\tilde{F}_a(0,0,1,X) = 0,\ \tilde{F}_b(0,0,1,X) = 0$ has solutions x_1, x_2 THEN

6. S:= $S \cup \{(0:0:1:x_1),(0:0:1:x_2)\}$;

7. END IF;

8. IF $\tilde{F}_a(0,1,X,Y) = 0,\ \tilde{F}_b(0,1,X,Y) = 0$ has solutions $(x_1,y_1),(x_2,y_2),(x_3,y_3),(x_4,y_4)$ THEN

9. append S:= $S \cup \{(0:1:x_1:y_1),(0:1:x_2:y_2),(0:1:x_3:y_3),(0:1:x_4:y_4)\}$;

10. END IF;

11. FOR x in $k_{\mathfrak{p}}$ DO

12. IF $\tilde{F}_a(1,x,Y,Z) = 0,\ \tilde{F}_b(1,x,Y,Z) = 0$ has solutions $(y_1,z_1),(y_2,z_2),(y_3,z_3),(y_4,z_4)$ THEN

13. S:= $S \cup \{(1:x:y_1:z_1),(1:x:y_2:z_2),(1:x:y_3:z_3),(1:x:y_4:z_4)\}$;

14. END IF;

15. END FOR;

16. RETURN S;

8.5 Constructing the lattice

For each point P on $\mathcal{C}$ mod p our aim is to construct a lattice in R^4 such that all points on $\mathcal{C}$ that reduce to P mod p lie in the lattice. The lattice will be constructed so as to be defined by a matrix in the following form:

$$\begin{pmatrix} 1 & \star & \star & \star \\ 0 & p & p\star & p\star \\ 0 & 0 & p^2 & p^2\star \\ 0 & 0 & 0 & p^3 \end{pmatrix}$$

where $\star$ denotes an element of R. The rows of the matrix form the basis of the lattice.

The preliminary step is to choose a prime p that is large enough (has large enough degree). The choice of p will b discussed later in 8.6.1.

8.5.1 Step 1

We now have chosen a prime p such that $\mathcal{C}$ has good reduction at p. Let $F_a = 0$ and $F_b = 0$ be the equations defining the quadric intersection $\mathcal{C}$. Let $\mathbf{F}$ be the column vector of the defining polynomials F_a and F_b. Let $\nabla\mathbf{F}$ be the Jacobian matrix of $\mathbf{F}$ so:

$$\nabla\mathbf{F} = \begin{pmatrix} \frac{\partial F_a}{\partial W} & \cdots & \frac{\partial F_a}{\partial Z} \\ \frac{\partial F_b}{\partial W} & \cdots & \frac{\partial F_b}{\partial Z} \end{pmatrix}. \tag{8.6}$$

Let $\tilde{\mathcal{C}}$ be the reduced curve $\mathcal{C}$ mod p, and $\tilde{P}$ be a point on $\tilde{\mathcal{C}}$ that we have found using the algorithm "Points mod p" described in 8.4. $\nabla\mathbf{F}(\tilde{P})$ has rank 2 since $\tilde{\mathcal{C}}$ is smooth at $\tilde{P}$. After permuting variables if necessary we may assume that $\tilde{P}$ may be lifted to $\mathbf{x} \in R^4$, $\mathbf{x} = (1, x_2, x_3, x_4)$.

Since p is a prime where $\mathcal{C}$ has good reduction, we may assume that the submatrix of $\nabla\mathbf{F}$ consisting of the second, third and fourth columns, has rank 2. After a second permutation of variables if necessary we may now assume that the matrix:

$$M(\mathbf{X}) = \begin{pmatrix} \frac{\partial F_a}{\partial Y} & \frac{\partial F_a}{\partial Z} \\ \frac{\partial F_b}{\partial Y} & \frac{\partial F_b}{\partial Z} \end{pmatrix} \tag{8.7}$$

is non-singular mod p at $\tilde{P}$.

8.5.2 Step 2

Next we lift $\mathbf{x}$ to a point mod p^3, so $\mathbf{F}(\mathbf{x}) \equiv 0 \pmod{p^3}$.

Since $\tilde{C}$ is non-singular at $\tilde{P}$, by Hensel's Lemma we can lift $\tilde{P}$ to a point mod p^k for $k \geq 1$. First we lift $\mathbf{x}$ to a point mod p^2: Set $\mathbf{x}' = \mathbf{x} + p(0, 0, a, b)$

where

$$\begin{pmatrix} a \\ b \end{pmatrix} = -M(\mathbf{x})^{-1} \begin{pmatrix} F_a(\mathbf{x})/p \\ F_b(\mathbf{x})/p \end{pmatrix} \mod p \tag{8.8}$$

so now $\mathbf{F}(\mathbf{x}') \equiv 0 \mod p^2$. Then set $\mathbf{x}'' = \mathbf{x}' + p^2(0,0,c,d)$ where

$$\begin{pmatrix} c \\ d \end{pmatrix} = -M(\mathbf{x}')^{-1} \begin{pmatrix} F_a(\mathbf{x}')/p^2 \\ F_b(\mathbf{x}')/p^2 \end{pmatrix} \mod p \tag{8.9}$$

so $\mathbf{F}(\mathbf{x}'') \equiv 0 \mod p^3$.

The above works by examining the Taylor expansion of $\mathbf{F}$:

$$\begin{aligned} \mathbf{F}(\mathbf{x}') &= \mathbf{F}(\mathbf{x} + p(0,0,a,b)) \\ &\equiv \mathbf{F}(bx) + \nabla\mathbf{F}(\mathbf{x}) \cdot (0,0,a,b) \mod p^2 \\ &\equiv \mathbf{F}(\mathbf{x}) + pM(\mathbf{x}) \begin{pmatrix} a \\ b \end{pmatrix} \\ &\equiv 0 \mod p^2. \end{aligned}$$

Similarly $\mathbf{F}(\mathbf{x}'') \equiv 0 \mod p^3$.

8.5.3 Step 3

We adjust our notation so $\mathbf{x}''$ is now simply called $\mathbf{x} = (1, x_2, x_3, x_4)$. Then $\mathbf{F}(\mathbf{x}) \equiv 0 \mod p^3$.

Let $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4$ be the four unit vectors of R^4. Let $L_1 = \langle \mathbf{x}, p\mathbf{e}_2, p\mathbf{e}_3, p\mathbf{e}_4 \rangle$ be the lattice of all lifts of $\tilde{P}$. Next we construct a lattice L_2 of index p^5 in R^4 that is a sublattice of L_1 consisting precisely of the vectors $\mathbf{y} \in L_1$ such that $\mathbf{F}(\mathbf{y}) \equiv 0 \mod p^2$.

Set $u_1 = 0$, $u_2 = 1$ and $\begin{pmatrix} u_3 \\ u_4 \end{pmatrix} \equiv -M(\mathbf{x})^{-1} \begin{pmatrix} \frac{\partial \mathbf{F}}{\partial X}(\mathbf{x}) \\ \frac{\partial \mathbf{F}}{\partial X}(\mathbf{x}) \end{pmatrix} \mod p^2$, so that

$\nabla \mathbf{F}(\mathbf{x}) \cdot \mathbf{u} \equiv 0 \mod p^2$, and let L_2 be the lattice with basis:

$$L_2 = \langle \mathbf{x}, p\mathbf{u}, p^2\mathbf{e}_3, p^2\mathbf{e}_4 \rangle. \tag{8.10}$$

Lemma 12. *The primitive vectors in L_2 are precisely the primitive vectors $\mathbf{y} \in L_1$ such that $\mathbf{F}(\mathbf{y}) \equiv 0 \mod p^2$. By primitive we mean that not all coordinates are divisible by p.*

Proof. Let $\mathbf{y} = \lambda\mathbf{x} + p\mathbf{t} \in L_1$ with $\mathbf{t} = (0, t_2, t_3, t_4)$, $t_i \in R$, $\lambda \in R$, $\lambda \notin pR$ (primitivity). Then

$$\begin{aligned} \mathbf{F}(\mathbf{y}) &\equiv \mathbf{F}(\lambda\mathbf{x}) + \nabla\mathbf{F}(\lambda\mathbf{x}) \cdot p\mathbf{t} \mod p^2 \\ &\equiv \lambda^2\mathbf{F}(\mathbf{x}) + \lambda p \nabla\mathbf{F}(\mathbf{x}) \cdot \mathbf{t} \mod p^2 \\ &\equiv 0 + \lambda p \nabla\mathbf{F}(\mathbf{x}) \cdot \mathbf{t} \mod p^2. \end{aligned} \tag{8.11}$$

So $\mathbf{F}(\mathbf{y}) \equiv 0 \mod p^2 \Leftrightarrow \nabla\mathbf{F}(\mathbf{x}) \cdot \mathbf{t} \equiv 0 \mod p \Leftrightarrow \mathbf{t} = t_2\mathbf{u} \mod p \Leftrightarrow \mathbf{y} \in L_2$ □

Higher order Taylor expansions

We will need more terms of the Taylor expansion of $\mathbf{F}$ to carry out the next step. We use the following fact:

$$\mathbf{F}(\mathbf{x} + \mathbf{h}) = \mathbf{F}(\mathbf{x}) + \nabla\mathbf{F}(\mathbf{x}) \cdot \mathbf{h} + \mathbf{H}(\mathbf{x})(\mathbf{h}) + \cdots \tag{8.12}$$

where $\mathbf{H}(\mathbf{x})(\mathbf{h}) = \frac{1}{2}\begin{pmatrix} \mathbf{h}^T H_a(\mathbf{x})\mathbf{h} \\ \mathbf{h}^T H_b(\mathbf{x})\mathbf{h} \end{pmatrix}$ and H_a is the Hessian of F_a given by $H_a = \left(\frac{\partial^2 F_a}{\partial X_i \partial X_j}\right)$ (similarly for H_b). Note that when $\deg(F) = 2$ as is the case here, we have $\frac{1}{2}\mathbf{h}^T H(\mathbf{x})\mathbf{h} = F(\mathbf{h})$.

8.5.4 Step 4

The final step is to construct a lattice $L_3 \subset R^4$ of index p^6 such that every primitive lift of $\tilde{P}$ that is a solution to $\mathcal{C} \bmod p^3$ lies in L_3. This time however

the lattice does not consist precisely of such points, it only contains them all.

Write $\mathbf{H}(\mathbf{x})(\mathbf{u}) = \frac{1}{2}\begin{pmatrix}\mathbf{u}^T H_a(\mathbf{x})\mathbf{u}\\ \mathbf{u}^T H_b(\mathbf{x})\mathbf{u}\end{pmatrix}$. In our case, since $\deg(F_a) = \deg(F_b) = 2$ we have $\mathbf{H}(\mathbf{x})(\mathbf{u}) = \mathbf{F}(\mathbf{u})$. Let $\begin{pmatrix}a\\ b\end{pmatrix} \equiv M(\mathbf{x})^{-1}\mathbf{H}(\mathbf{x})(\mathbf{u}) \mod p$. Then we claim that $\begin{pmatrix}a\\ b\end{pmatrix} \not\equiv \begin{pmatrix}0\\ 0\end{pmatrix} \mod p$. For if this were not the case then it would follow that $\mathbf{F}(\mathbf{u}) \equiv 0 \mod p$, and hence $\mathbf{F}(\lambda\mathbf{x} + \mu\mathbf{u}) \equiv 0 \mod p \ \forall \lambda, \mu$ which would mean that $\tilde{\mathcal{C}}$ contains a line (since $\mathbf{x}$ and $\mathbf{u}$ are independent) and this cannot occur, since we have good reduction at p.

After making our last available permutation of variables if necessary, we assume that $a \not\equiv 0 \mod p$. Set $\mu \equiv a^{-1}b \mod p$ and let L_3 be the following lattice:

$$L_3 = \langle \mathbf{x}, p\mathbf{u}, p^2\mathbf{u}', p^3\mathbf{e}_4 \rangle \tag{8.13}$$

where $\mathbf{u}' = (0, 0, 1, \mu)$. This is a lattice of index p^6 in R^4. We now show that every mod p^3 lift of $\tilde{P}$ lies in L_3:

Proposition 13. *Every primitive* $\mathbf{w} \in L_1$ *satisfying* $\mathbf{F}(\mathbf{w}) \equiv 0 \mod p^3$ *lies in* L_3.

Proof. Such a $\mathbf{w}$ is clearly in L_2 so $\mathbf{w} = \lambda\mathbf{x} + p\gamma\mathbf{u} + p^2\mathbf{v}$ where $\mathbf{v} = (0, 0, c, d)$, $\lambda, \gamma, c, d \in R$ and $p \nmid \lambda$. We need to show that $d \equiv \mu c \mod p$, that is $ad \equiv bc \mod p$.

We have

$$\begin{aligned} 0 &\equiv \mathbf{F}(\mathbf{w}) \\ &\equiv \mathbf{F}(\lambda\mathbf{x}) + \nabla\mathbf{F}(\lambda\mathbf{x}) \cdot (p\gamma\mathbf{u} + p^2\mathbf{v}) + \mathbf{H}(\lambda\mathbf{x})(p\gamma\mathbf{u} + p^2\mathbf{v}) \mod p^3 \\ &\equiv p^2\lambda\nabla\mathbf{F}(\mathbf{x}) \cdot \mathbf{v} + p^2\gamma^2\mathbf{H}(\mathbf{x})(\mathbf{u}) \end{aligned} \tag{8.14}$$

using $\mathbf{F}(\mathbf{x}) \equiv 0 \mod p^3$ and $\nabla\mathbf{F}(\mathbf{x}) \cdot \mathbf{u} \equiv 0 \mod p^2$ and the fact that $\nabla\mathbf{F}$ is homogeneous of degree 1, and $\mathbf{H}$ is constant.

It follows that:

$$\begin{aligned}
& \lambda\nabla\mathbf{F}(\mathbf{x}) \cdot \mathbf{v} + \gamma^2\mathbf{H}(\mathbf{x})(\mathbf{u}) \equiv 0 \mod p \\
\Rightarrow \quad & \lambda M(\mathbf{x}) \begin{pmatrix} c \\ d \end{pmatrix} + \gamma^2 M(\mathbf{x}) \begin{pmatrix} a \\ b \end{pmatrix} \equiv \begin{pmatrix} 0 \\ 0 \end{pmatrix} \\
\Rightarrow \quad & \lambda \begin{pmatrix} c \\ d \end{pmatrix} + \gamma^2 \begin{pmatrix} a \\ b \end{pmatrix} \equiv \begin{pmatrix} 0 \\ 0 \end{pmatrix} \\
\Rightarrow \quad & \lambda(d - \mu c) \equiv -\gamma^2(b - \mu a) \equiv 0 \\
\Rightarrow \quad & d - \mu c \equiv 0 \mod p.
\end{aligned} \tag{8.15}$$

□

Let M_3 be the 4×4 matrix whose rows consist of the basis vectors of L_3 as given in (8.13). In order to apply Lemma 9, this matrix will need to be converted to weak Popov form. Since this transformation is carried out through multiplication by a unimodular matrix, the resulting weak Popov form matrix defines the same lattice. We will need to use the following fact:

Proposition 14. *Let M be an $n \times n$ matrix over $F(t)$ where F is a field. If M is in weak Popov form and $\det(M) = D$ then*

$$\deg(D) = \sum_{i=1}^{n} \deg(P_i^M)$$

where P_i^M denotes the ith pivot-element of M.

Proof. Since the determinant of an $n \times n$ matrix is a sum of $n!$ terms each consisting of a product of n matrix entries no two of which share a row or column, it is clear that one term must be the product of the pivots. All other terms must contain at least one entry to the right of the pivot in its row and

will therefore have strictly smaller degree than the pivot-product. The result follows. □

8.6 Using the lattice

8.6.1 Properties of the lattice basis

We now have, for each point P on $\mathcal{C}$ (mod p), a lattice L_P defined by a matrix M_P in weak Popov form, containing all points on $\mathcal{C}$ that are congruent to P mod p. In the description that follows we omit the subscript $_P$ from L and M as the method will be applied to each lattice independently.

The rows of M are the basis of L. Let $\mathbf{b}_i$ $(1 \leq i \leq 4)$ denote the rows of M and $d_i := \deg(\mathbf{b}_i)$ denote the degree of the ith pivot element of M. After permuting the rows we can assume that $d_1 \leq d_2 \leq d_3 \leq d_4$. By proposition 14 we have:

$$d_1 + d_2 + d_3 + d_4 = 6\deg(p). \tag{8.16}$$

We wish to find all points on $\mathcal{C} \cap L$ whose degree is no greater than a desired bound, denoted by H. Let $\mathbf{v} = \sum_{i=1}^{4} a_i \mathbf{b}_i$ be a vector in the lattice. By lemma 9:

$$\deg(\mathbf{v}) = \max_{1 \leq i \leq 4} \{\deg(a_i) + d_i\}. \tag{8.17}$$

So for $\mathbf{v}$ to have degree $\leq H$ we only need to consider, for each i, scalars a_i whose degree is no greater then $H - d_i$. In particular if one or more of the rows of M have degree greater then H then these rows are redundant. With less rows to consider, the algorithm for finding points on $\mathcal{C}$ in each lattice is faster, however this is achieved by choosing p to be of a higher degree, which increases the number of lattices that need to be analysed. For the algorithm described

below we choose p so that $d_4 > H$. This is the default setting in the MAGMA implementation, although one can, if desired, change the parameter H_0 (the degree of p) to be a different value.

Choice of p

Given a quadric-intersection $\mathcal{C}$ the first step towards finding all points on $\mathcal{C}$ of degree $\leq H$ is to find all points on the reduced-mod-p curve $\tilde{C}$, as described in 1.2.2. Since constructing the lattice L_P from each such point P takes time $O(1)$ the timing of the overall algorithm is dependent on our choice of p itself. Two factors need to be taken into consideration. Firstly and most importantly is the degree of p, as this directly determines both how many lattices need to be constructed and how long each takes to analyse. Secondly, once the degree of p has been decided, a further refinement can be made by choosing p such that the reduced curve $\tilde{C}$ has as few points as possible.

Using the notation established in 1.2.2, $\tilde{C}$ is a curve of genus one over a finite field of size $q^{\deg(p)}$, so by Hasse's Theorem the number of points on $\tilde{C}$ is $O(q^{deg(p)})$. The number of lattices that have to be constructed therefore increases exponentially with the degree of p.

Proposition 15. *Let H be a positive integer. Let L, M, $\mathbf{b}_i$ and d_i be as described in 1.2.4, with $d_1 \leq d_2 \leq d_3 \leq d_4$. If $\deg(p) > \frac{2H}{3}$ then every vector in L of degree less than or equal to H is contained in the span of $\mathbf{b}_1$, $\mathbf{b}_2$ and $\mathbf{b}_3$.*

Proof. Let L' denote the span of $\mathbf{b}_1$, $\mathbf{b}_2$ and $\mathbf{b}_3$.

If $\deg(p) > \frac{2H}{3}$ then $6\deg(p) > 4H$. It follows from (1.8) that $\sum_{i=1}^{4} d_i > 4H$. Since $d_4 \geq d_i$ for $i < 4$, it follows that $d_4 \geq \frac{1}{4}\sum_{i=1}^{4} d_i$, hence $d_4 > H$.

By Lemma 9, since $d_4 > H$, any vector in L that is not in L' has degree

greater than H, therefore every vector in L of degree at most H lies in L'.

□

The lower bound of $\frac{2H}{3}$ on $\deg(p)$ cannot be improved on, for while it is possible to have $d_4 > H$ if $\deg(p)$ is smaller, this cannot be guaranteed. For example if $\deg(p) = \frac{2H}{3}$ then the d_i can all equal H.

The above proposition allows us to choose p so that we only need to search for points on $\mathcal{C}$ in the span of three vectors in the lattice, which improves the running-time of the program considerably (more on this later).

8.6.2 The search algorithm

We now have a lattice $L \subset \mathbb{F}_q[t]^4$ of index p^6 together with a matrix M, in weak Popov form whose rows $\mathbf{b}_1, \ldots, \mathbf{b}_4$ are the basis vectors of L. All that remains is to search for vectors in L of bounded degree that are points on the quadric-intersection $\mathcal{C}$. We assume now that p has been chosen so that $\mathbf{b}_4$ has degree greater than H.

The method used to find points on $\mathcal{C}$ depends now on whether M has one, two or three "good" rows. A "good" row is one whose degree is $\leq H$. In L, only a linear combination of the good rows can have degree $\leq H$ by Lemma 1.

One good row

If there is only one good row, then all that needs to be checked is whether this row, as an element of $\mathbb{P}^3(\mathbb{F}_q(t))$ is a point on $\mathcal{C}$. Since $\mathcal{C}$ is a projective curve, scalar multiplication (by an element of $\mathbb{F}_q[t]$) does not change the point.

Two good rows

In this case a candidate for a point on $\mathcal{C}$ has the form $u\mathbf{b}_1 + v\mathbf{b}_2$ where u and v are elements of $\mathbb{F}_q[t]$. Substituting this general point into the equations defining $\mathcal{C}$ we get

$$F(u\mathbf{b}_1 + v\mathbf{b}_2) = 0 \tag{8.18}$$

where F is one of the two defining polynomials of $\mathcal{C}$. Since scaling u and v by the same element of $\mathbb{F}_q(t)$ does not change $u\mathbf{b}_1 + v\mathbf{b}_2$ as a point in projective space, the above equation is a homogeneous quadratic in the variable $[u : v]$ and can be easily solved. To find points on $\mathcal{C}$ we therefore look for common solutions to $F_a(u\mathbf{b}_1 + v\mathbf{b}_2) = 0$ and $F_b(u\mathbf{b}_1 + v\mathbf{b}_2) = 0$. Note that here we have not placed any restrictions on the degrees of u and v, so this method may find points outside the degree bound H.

Three good rows

Similar to the two-good-rows case, a candidate for a point on $\mathcal{C}$ has the form $u\mathbf{b}_1 + v\mathbf{b}_2 + w\mathbf{b}_3$ where u and v lie in $\mathbb{F}_q[t]$. However, we can now proceed in two different ways. One method is to continue in a similar way to the previous case, resulting in two quadratic polynomials in $[u : v : w]$ defining the intersection of two conics in $\mathbb{P}^2(\mathbb{F}_q[t])$. One can then either ask Magma for the points directly (since this has dimension zero) or solve the equations using the conic-solving (see below) methods described in chapter 5. As in the previous case, we may find points outside the desired degree bound.

An alternative way of proceeding is to put restrictions on the degrees of the polynomials u, v, and w. We write these as $u = u_0 + u_1 t + \cdots + u_{H-d_1} t^{H-d_1}$ and similarly for v and w. The coefficients u_i, v_i and w_i represent elements

of $\mathbb{F}_q$. For the purposes of MAGMA they are variables in a multivariate polynomial ring over $\mathbb{F}_q$. Substituting the general point $u\mathbf{b}_1 + v\mathbf{b}_2 + w\mathbf{b}_3$ into the equations for $\mathcal{C}$ yields two polynomials in t with coefficients that are polynomials in $\mathbb{F}_q[u_0, \ldots, u_{H-d_1}, v_0 \ldots, v_{H-d_2}, w_0, \ldots, w_{H-d_3}]$ - a polynomial ring with $3+3H-(d_1+d_2+d_3)$ variables. For the general point to lie on $\mathcal{C}$ we require that these coefficients equal zero. Equating all of these polynomials to zero results in a set of equations defining a variety over $\mathbb{F}_q$. Since this is a variety over a finite field one can then ask MAGMA for its points. Substituting these back into u, v and w gives points on $\mathcal{C}$.

This method can also be used when there are four good rows (though this case only arises if H_0 (the degree of p) is chosen manually to be $\leq \frac{2H}{3}$).

The three-good-rows case via conic parametrization

Since we can solve and parametrize conics over function fields (see chapter 5), the case where M has three rows of degree $\leq H$ can be handled in the following way.

Let $P(u, v, w) = u\mathbf{b}_1 + v\mathbf{b}_2 + w\mathbf{b}_3$ $(u, v, w \in \mathbb{F}_q[t])$ be a general point of L in the span of the first three basis vectors (rows of M). If there is a point $(u_0 : v_0 : w_0) \in \mathbb{P}^2(\mathbb{F}_q(T))$ such that $P(u_0, v_0, w_0)$ lies on $\mathcal{C}$ then u_0,v_0 and w_0 provide a solution to the following pair of equations:

$$\begin{cases} F_a(P(u, v, w)) & = 0 \\ F_b(P(u, v, w)) & = 0. \end{cases} \tag{8.19}$$

The above two equations define a pair of conics in $\mathbb{P}^2(\mathbb{F}_q(t))$. Let C_a and C_b denote the conics defined by $F_a(P)$ and $F_b(P)$ respectively. If neither conic has any points then the lattice can give rise to no points on $\mathcal{C}$. Otherwise we solve and parametrize the first (though either will do) conic. Solving $F_a(P(u, v, w)) =$

0 and parametrizing we obtain three quadratics $q_i(X,Y)$ $(i = 1,2,3)$ such that $(q_1(x,y) : q_2(x,y) : q_3(x,y))$ is a point on C_a for any $(x : y)$ on the projective line over $\mathbb{F}_q(T)$.

To obtain a point on $\mathcal{C}$ we require a point to lie on both C_a and C_b. To find such points we substitute the parametrization $(q_1 : q_2 : q_3)$ into the polynomial defining C_b. This then gives $F_b(P(q_1, q_2, q_3))$, which is a homogeneous quartic in the two variables X, Y. It is then simply a matter of finding the roots (if any) of the quartic. If the quartic has a root $(x : y)$ then $P(q_1(x,y), q_2(x,y), q_3(x,y))$ is a point on the quadric-intersection $\mathcal{C}$.

8.7 Special treatment of the cases $H = 0$ and $H = 1$

Since the first step of the method described in this chapter is to reduce the curve $\mathcal{C}$ mod p, it soon becomes obvious that in the case $H = 0$, the "best" degree for p to have as 0. Since we cannot reduce by a constant in $\mathbb{F}_q[t]$ a separate treatment is required. Some particular features of MAGMA also dictate that we must also treat $H = 1$ as a special case.

A couple of options are available: the first would be to increase the degree of p to 2 and then continue as for the $H \geq 2$ cases. This will be faster than if we were searching for points of degree up to 2 say, since the number of "good" rows in each lattice-basis-matrix will usually be smaller.

Alternatively we define a general point as $P = (w : x : y : z)$ for points of degree 0, and $P = (w_0 + w_1 t : x_0 + x_1 t : y_0 + y_1 t : z_0 + z_1 t)$ for points of degree ≤ 1. Substituting this into the polynomials defining $\mathcal{C}$ we obtain two

polynomials $F_a(P)$ and $F_b(P)$. These can be viewed as polynomials in t over the multivariate polynomial ring whose variables are the coefficients w, x, y, z etc. Setting all coefficients of powers of t to zero gives a set of equations defining a variety over $\mathbb{F}_q$. The points of this variety then yield points on $\mathcal{C}$ after substitution into the general point P.

Chapter 9

Possible directions for further work

9.1 When E has no 2-torsion

The two main algorithms give an explicit descent method for calculating rank-bounds on elliptic curves but only work when the curves have full 2-torsion defined over the base-field $K = \mathbb{F}_q(T)$. An obvious next step to take would be to seek methods that apply to *all* elliptic curves over K. One possible approach would be to extend K to an algebraic function field L over which E has a two torsion point and try to adapt the methods given in this thesis to work over such fields.

9.2 More general function fields

This thesis has been only concerned with function fields of the type $\mathbb{F}_q(T)$. It would be a logical step to seek methods for calculating the Mordell Weil group of elliptic curves over more general function fields, e.g. $\mathbb{Q}(T)$, $\mathbb{C}(T)$ or algebraic function fields. Over $\mathbb{Q}(T)$ one immediately encounters the problem that $\mathbb{Q}(T)(S,2)$ is infinite, and over $\mathcal{C}(T)$ the problem is that every homogeneous space constructed using the methods given in this thesis always has local points since $\mathbb{C}$ is algebraically closed. Kuwata [Kuw] has shown how certain elliptic curves (constructed from $K3$-surfaces) over $\mathbb{C}(T)$ or $\bar{\mathbb{Q}}(T)$ can be handled.

9.3 Generators for the Mordell-Weil group

This thesis has shown how the Mordell Weil group and the rank of an elliptic curve over $K = \mathbb{F}_q(T)$ may be calculated as an abstract group, but though we explicitly find independent points on E these do not necessarily generate $E(K)$, only a subgroup of finite index. This problem of saturation is another possible extension of the work carried out here.

Appendix A

Tables of collected run-time data

The first table below gives the CPU time taken to find all points of degree ≤ 5 on quadric intersections over base-fields $\mathbb{F}_p(T)$ for various p. The quadric intersections all arise from the following equations over $\mathbb{Z}(T)$:

$$\begin{cases} 6{Z_1}^2 + 6T^2{Z_2}^2 + (T+6){Z_4}^2 = 0 \\ (6T^2+6){Z_1}^2 + (T^4+T^2){Z_3}^2 + T{Z_4}^2 = 0. \end{cases} \tag{A.1}$$

Each quadric intersection over $\mathbb{F}_p(T)$ is formed by mapping the coefficients into $\mathbb{F}_p(T)$ in the obvious way. In the case $p = 5$ the curve is singular.

The second table gives the time taken to calculate the rank of the following elliptic curve over $\mathbb{F}_{31}(T)$ using descent via two isogeny with varying height bound H:

$$Y^2 = (X + 14T + 16)(X^2 + (17T + 4)X + 18). \tag{A.2}$$

Figure A.1: Timing for point search with varying base field

p	CPU time (seconds)
3	2.53
7	91.44
11	536.57
13	1079.48
17	3550.73
19	5058.24
23	12005.47
29	30507.32
31	40735.68
37	73334.18
41	111500.05
43	150164.44

Figure A.2: Timing for descent via 2-isogeny with varying H.

H	CPU time (seconds)
2	0.82
3	1.98
4	10.26
5	18.54
6	1340.04
7	19759.06

Bibliography

[Cr1] J. Cremona, *Algorithms for Modular Elliptic Curves*. Cambridge University Press (1997).

[Cr2] `http://www.warwick.ac.uk/~masgaj/ftp/progs/magma/`

[Cr3] J. Cremona, *Higher Descents on Elliptic Curves* (1997) `http://www.warwick.ac.uk/~masgaj/papers/d2.ps`

[CrRo] J. Cremona, D. Roberts, *Applications of polynomial lattices to point-finding over rational function fields*. (2007)

[CrRu] J. Cremona, D. Rusin, *Efficient solution of rational conics*. Math. Comp. **72** (2003), no. 243, pp. 1417-1441.

[Elk] N. D. Elkies, *Rational points near curves and small nonzero $|x^3 - y^2|$ via lattice reduction*, In ANTS-IV proceedings (W. Bosma, ed), LNCS **1838**. Springer-Verlag (2000), pp. 33-63.

[H-B] D. R. Heath-Brown, Personal communication, 4/11/1999

[Kuw] M. Kuwata, *Mordell-Weil Groups and Elliptic K3 Surfaces*. PhD thesis, Brown University, 1989.

[Len] A. K. Lenstra, *Factoring mutilvariate polynomials over finite fields*. Journal of Computer and System Sciences **30** (1985), pp. 235-248.

[Lon] R. Long, *The Algorithmic Solution of Simultaneous Diophantine Equations*. PhD thesis, Oxford Brookes University, 2006.

[Mag] `http://magma.maths.usyd.edu.au/magma/htmlhelp/MAGMA.htm`

[MSS] J. Merriman, S. Siksek, N. Smart, *Explicit 4-descents on an elliptic curve*. Acta Arithmetica **LXXVII.4** (1996), pp. 358-404.

[MuSt] T. Mulders, A. Storjohann, *On lattice reduction for polynomial matrices*. Journal of Symbolic Computation **35** (2003), pp. 377-401.

[Ros] M. Rosen, *Number Theory in Function Fields*. Springer, Graduate Texts in Mathematics **210** (2002).

[Sil1] J. Silverman, *The Arithmetic of Elliptic Curves*. Springer, Graduate Texts in Mathematics **106** (2002).

[Sil2] J. Silverman, *Advanced Topics in the Arithmetic of Elliptic Curves*. Springer, Graduate Texts in Mathematics **151** (1994).

[Sim1] D. Simon, *Sur la paramétrisation des solutions des équations quadratiques*, Journal de Théorie des Nombres de Bordeaux **18** (2006), no 1, pp.265-283.

[vHCr] M. van Hoeij, J. Cremona, *Solving conics over function fields*. Journal de Théorie des Nombres de Bordeaux **18** (2006), pp. 595-606.

[Wat] M. Watkins, *Searching for points with the Elkies ANTS-IV algorithm*. `http://www.maths.bris.ac.uk/~mamjw/papers/padic.ps`

[Wom] T. Womack, *Explicit descent on elliptic curves*. PhD thesis, University of Nottingham, 2003.