Construction of Rational Points on Elliptic Curves over Finite Fields

Andrew Shallue[1] and Christiaan E. van de Woestijne[2]

[1] University of Wisconsin-Madison, Mathematics Department
480 Lincoln Dr, Madison, WI 53706-1388 USA
`shallue@math.wisc.edu`
[2] Universiteit Leiden, Mathematisch Instituut
Postbus 9512, 2300 RA Leiden, The Netherlands
`cvdwoest@math.LeidenUniv.nl`

Abstract. We give a deterministic polynomial-time algorithm that computes a nontrivial rational point on an elliptic curve over a finite field, given a Weierstrass equation for the curve. For this, we reduce the problem to the task of finding a rational point on a curve of genus zero.

1   Introduction

Elliptic curves over finite fields have been in the centre of attention of cryptographers since the invention of ECC, and in that of number theorists for a much longer time. It is not very hard to show that, unless the base field is extremely small, such curves always have rational points other than O, the point at infinity. However, it is a different question how to construct such rational points efficiently.

Until now, this was possible only using an obvious probabilistic method: given an equation for the curve, substitute random values for all coordinates but one and see if the remaining univariate equation can be solved for the last coordinate. If so, a probabilistic polynomial factorisation algorithm will give the last coordinate and a rational point has been found. The challenge for a deterministic algorithm has been up at least since 1985, when R. Schoof posed it in [8].

In a recent publication [11], however, M. Skałba proved that, given a cubic polynomial $f(x) = x^3 + Ax + B$ over a field F with characteristic unequal to 2 or 3, with $A \neq 0$, we have the identity

$$f(X_1(t^2))f(X_2(t^2))f(X_3(t^2)) = U(t)^2 \tag{1}$$

for some nonconstant univariate rational functions X_1, X_2, X_3, U over F. Such functions are given explicitly in his paper [11, Theorem 1]. We do not reproduce them here, as both their degree and their coefficients are large; if $X_1 X_2 X_3 = N/D$ for coprime polynomials N and D in $F[t]$, then $\deg N \leq 26$ and $\deg D \leq 25$, depending on the characteristic of F.

Now assume that $\mathbb{F}$ is a finite field and that the curve E is defined over $\mathbb{F}$ by the equation $y^2 = f(x)$, with f as above. The multiplicative group $\mathbb{F}^*$ is cyclic,

and therefore, as Skałba notes, if we specialise t in (1) to some value t_0 in $\mathbb{F}$, we find that at least one of the $f(X_i(t_0^2))$ is a square in $\mathbb{F}^*$. However, no efficient deterministic algorithm is known to date to take the square root.

In this paper, we show how to go on from this point to obtain a complete efficient deterministic algorithm for constructing rational points on curves given by cubic Weierstrass equations over finite fields. We will reprove Skałba's result to obtain, for the case of finite fields of odd characteristic, a parametrisation as in (1) that is *invertible* as a rational map (Lemmas 6 and 7 below).

The construction of this parametrisation in the case of odd characteristic rests on the ability to solve deterministically and efficiently equations of the form

$$ax^2 + by^2 = c \tag{2}$$

over finite fields, for which an algorithm will be given in Section 2 (Theorem 4).

In Section 2, we also give a deterministic algorithm that, given nonzero elements a_0, a_1, a_2 in a finite field such that their product is the square of a given element, computes a square root of one of them, in polynomial time. It is clear that such an algorithm is the missing step to construct a rational point on E, when an equation of the form (1) is given.

An analogon of (1) for finite fields of characteristic 2 will be used to obtain a point finding algorithm for elliptic curves in this case as well.

The main result is as follows:

Theorem 1. *There exists a deterministic algorithm that, given a finite field $\mathbb{F}$ of q elements and a cubic Weierstrass equation $f(x, y)$ over $\mathbb{F}$:*

- *(i) detects if $f(x, y)$ is singular, and if so, computes the singular points and gives a rational parametrisation of all rational points on the curve $f(x, y) = 0$;*
- *(ii) if $f(x, y)$ is nonsingular and $|\mathbb{F}| > 5$, computes an explicit rational map ρ from the affine line over $\mathbb{F}$ to an affine threefold V that is given explicitly in terms of the coefficients of f;*
- *(iii) given a rational point on the threefold V, computes a rational point on the elliptic curve $E : f(x, y) = 0$, in such a way that at least $(q-4)/8$ rational points on E are obtained from the image of the map ρ, and at least $(q-4)/3$ if $\mathbb{F}$ has characteristic 2;*

and performs all these tasks in time polynomial in $\log q$.

From the proofs in this paper, such an algorithm can be explicitly constructed; the running time of this algorithm is not much worse than that of a probabilistic point generation algorithm. We plan to give an explicit algorithm, with detailed running time bounds, in a forthcoming publication.

After Section 2 on how to solve diagonal quadratic equations, we give some generalities on Weierstrass equations in Section 3 and show how to parametrise the solutions of a *singular* Weierstrass equation in Section 4. The nonsingular case, where the given equation indeed defines an elliptic curve, is split into two cases: in Section 5, we prove Theorem 1 for base fields of odd characteristic, whereas base fields of characteristic 2 are considered in Section 6.

2   Quadratic Equations

Before turning to cubic equations, we first give the necessary algorithms for solving quadratic equations. Theorem 3 is concerned with taking square roots, while Theorem 4 is about equations of the form (2). These results, which are taken from the second author's Ph.D. thesis [12], are new and deterministic efficient algorithms have been unknown to date.

We write $v_2(a)$ to denote the number of factors 2 in a nonzero integer a; if a is a nonzero element of a finite field $\mathbb{F}$, we write $\mathrm{ord}(a)$ to denote the order of a in the multiplicative group $\mathbb{F}^*$.

Lemma 2. *There exists a deterministic algorithm that, given a finite field $\mathbb{F}$ of q elements, and nonzero elements a and z of $\mathbb{F}$ such that either*

(i) $v_2(\mathrm{ord}\,a) < v_2(\mathrm{ord}\,z)$, or
(ii) $\mathrm{ord}\,a$ is odd,

computes a square root of a, in time polynomial in $\log q$.

Proof. We construct a deterministic adaptation of the Tonelli-Shanks algorithm; for the latter, see Section 1.5.1 in [5], for example.

It is easy to prove that to compute a square root of a nonzero element $a \in \mathbb{F}$, it is sufficient to have a generator z of the 2-Sylow subgroup of $\mathbb{F}^*$. Usually, such a generator is obtained by guessing a nonsquare element n and computing $z = n^u$, where we write $q - 1 = 2^e \cdot u$ such that u is an odd integer; this is the only probabilistic part of the Tonelli-Shanks algorithm.

The proof is as follows: a^u is in the 2-Sylow subgroup, and hence there exists an integer k such that $z^k = a^u$. The integer k is even if and only if a is a square in $\mathbb{F}$; furthermore, it is clear that $z^{k/2}$ is a square root of a^u, and from a square root of a^u it is easy to compute a square root of a, because u is odd and hence a^{u+1} is an obvious square. Thus, the real task of the Tonelli-Shanks algorithm is the computation of the integer k.

However, the only thing that is used about z is the fact that

$$a^u = z^k$$

for some even integer k; and for such a k to exist, it is only necessary that either $a^u = 1$, or the group generated by a^u inside 2-Syl $\mathbb{F}^*$ is strictly contained in the group generated by z. But these conditions correspond to our assumptions $v_2(\mathrm{ord}\,a) = 0$ and $v_2(\mathrm{ord}\,a) < v_2(\mathrm{ord}\,z)$, respectively. Therefore, if instead of a 2-Sylow subgroup generator we use any element whose order contains enough factors 2, the Tonelli-Shanks algorithm as given in [5] works just as well, while the nondeterministic part of guessing a nonsquare element is eliminated.   ◆

Theorem 3. *There exists a deterministic algorithm that, given a finite field $\mathbb{F}$ of q elements, and nonzero elements a_0, a_1, a_2, b of $\mathbb{F}$ such that $a_0 a_1 a_2 = b^2$, returns an i in $\{0, 1, 2\}$ and a square root of a_i, in time polynomial in $\log q$.*

Proof. After changing the order of the a_i, we may assume that $v_2(\operatorname{ord} a_0) \geq v_2(\operatorname{ord} a_1) \geq v_2(\operatorname{ord} a_2)$. If $v_2(\operatorname{ord} a_0) > v_2(\operatorname{ord} a_1)$, then by Lemma 2 we may use a_0 as a substitute for a 2-Sylow subgroup generator, and compute a square root of a_1; and if $v_2(\operatorname{ord} a_1) > v_2(\operatorname{ord} a_2)$, the same holds for a_1 and a_2.

Thus, consider the case where $v_2(\operatorname{ord} a_0) = v_2(\operatorname{ord} a_1) = v_2(\operatorname{ord} a_2)$. Then it follows that, say, $a_0 a_1$ has fewer factors 2 in its order than a_2, and we can compute $\sqrt{a_0 a_1}$; but by the given relation among the a_i, we have $a_0 a_1 = a_2/b^2$, and we compute a square root of a_2. ♦

Theorem 4. *There exists a deterministic algorithm that, given a finite field $\mathbb{F}$ of q elements, and nonzero elements a, b, c of $\mathbb{F}$, computes $x, y \in \mathbb{F}$ such that*

$$ax^2 + by^2 = c.$$

Proof. We may of course assume that $c = 1$. Now if $v_2(\operatorname{ord}(a)) > v_2(\operatorname{ord}(b))$, we can use the algorithm in Lemma 2 to take a square root of b, and the problem is solved by taking $x = 0$ and $y = 1/\sqrt{b}$; and analogously if b has the larger order. If $v_2(\operatorname{ord}(a)) = v_2(\operatorname{ord}(b)) =_{\operatorname{def}} w$, we distinguish three cases: $w = 0$, $w = 1$, and $w > 1$.

If $w = 0$, then we can still compute square roots of both a and b by means of Lemma 2, and we are done. If $w > 1$, then $v_2(\operatorname{ord}(-ab)) < w$, so that, after computing $\sqrt{-ab}$, we may assume $b = -a$. The equation $ax^2 - ay^2 = 1$ is easily solved by putting $x + y = 1$ and $x - y = 1/a$ and solving the linear system.

The case $w = 1$ is the hardest. Both $-a$ and $-b$ have odd order, so we may take their square roots by Lemma 2 and obtain the equation $-x^2 - y^2 = 1$. One sees that this is equivalent to

$$x^2 + y^2 + z^2 = 0.$$

For this, we developed a fast algorithm in Section 5.5 of [12]. A slower, but also deterministic, algorithm for this problem can be found in [4], and also the algorithm given in the second proof of Corollary 1 in [11] can be adapted to this case, by using Lemma 2. ♦

Remarks. It is well known that (2) is always solvable; this follows already from the fact that the cardinalities of the sets $\{ax^2 \mid x \in \mathbb{F}\}$ and $\{c - by^2 \mid y \in \mathbb{F}\}$ add up to more than q, and therefore these sets must meet.

The algorithm for solving (2) given above is a special case of the main algorithm from [12]; this algorithm can solve diagonal equations of the form

$$a_1 x_1^n + \ldots + a_n x_n^n = b$$

over finite fields.

In finite fields of characteristic 2, the above results are trivial, since all elements have odd order. However, over such fields many quadratic equations cannot be reduced to the diagonal form (2), and this yields new difficulties. We refer to Section 6 for a discussion of this case.

3   Weierstrass Equations

Let $\mathbb{F}$ be a finite field, let q be its number of elements, and let E be the affine curve given by the *Weierstrass equation*

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

where the a_i are in $\mathbb{F}$. The curve E also has one point at infinity, with homogeneous coordinates $(0 : 1 : 0)$, which is called $\mathcal{O}$.

If E is nonsingular, then the projective closure $\tilde{E}$ of E is a smooth projective curve of genus 1 over $\mathbb{F}$ with a specified rational point, so it is an *elliptic curve* over $\mathbb{F}$, and every elliptic curve over $\mathbb{F}$ may be given in this way [10, Proposition III.3.1]. The set of rational points on $\tilde{E}$ has a natural abelian group structure, with the point $\mathcal{O}$ as identity element.

We will be interested in methods to construct rational points on $\tilde{E}$ other than $\mathcal{O}$, or to show that no other points exist. By Hasse's bound [10, V.1.4], we know that the number N of rational points on $\tilde{E}$ satisfies

$$|q + 1 - N| \le 2\sqrt{q}.$$

From this, it is easily verified that $\tilde{E}$ has at least 2 rational points whenever $q \ge 5$. On the other hand, if $q \le 4$, curves over $\mathbb{F}$ exist with only the trivial rational point $\mathcal{O}$, such as the curve $y^2 = x^3 - x - 1$ over $\mathbb{F}_3$, and the curve $y^2 + y = x^3 + \alpha$ over $\mathbb{F}_4 = \mathbb{F}_2(\alpha)$.

Normal Forms. The equation (3) may be simplified depending on the characteristic of the base field. We give these forms in detail as we will use their properties later on; these formulas are given in Section III.1 and Appendix A of [10].

If the characteristic of $\mathbb{F}$ is not 2 or 3, then a linear change of coordinates transforms (3) into

$$y^2 = x^3 + Bx + C =_{\text{def}} f(x). \tag{4}$$

For this form of the equation, the important associated quantities Δ (the *discriminant*) and j (the *j-invariant*) are easily computed: we have

$$\Delta = -16(4B^3 + 27C^2), \qquad j = -1728(4B)^3/\Delta.$$

Now E is singular if and only if $\Delta = 0$, and thus if and only if the right hand side $f(x)$ of (4) has a repeated zero; it has j-invariant 0 if and only if $\Delta \ne 0$ and $B = 0$.

In characteristic 3, we must admit a third coefficient; we can transform (3) into

$$y^2 = x^3 + Ax^2 + Bx + C =_{\text{def}} f(x), \tag{5}$$

with associated quantities

$$\Delta = A^2B^2 - A^3C - B^3, \qquad j = A^2/\Delta.$$

Again, E is singular if and only if f has a double zero. Also, we find that for a nonsingular equation we have $j = 0$ if and only if $A = 0$.

In characteristic 2, no coefficient of (3) can be omitted in all cases. However, we can obtain one of the following two normal forms, depending on whether a_1 is zero:

$$Y^2 + a_3 Y = X^3 + a_4 X + a_6 \qquad\qquad \text{if } a_1 = 0 \text{ initially,} \qquad (6)$$

$$Y^2 + XY = X^3 + a_2 X^2 + a_6 \qquad\qquad \text{if } a_1 \neq 0 \text{ initially.} \qquad (7)$$

In these normal forms, we have $\Delta = (a_3)^4$ and $\Delta = a_6$, respectively, which gives an easy criterion for singularity of E. Furthermore, for nonsingular equations, the two cases correspond to j being respectively zero or nonzero.

4   Singular Weierstrass Equations

For completeness, we show how to detect deterministically whether E is singular and, if it is, how to find points on it. We continue to assume that $\mathbb{F}$ is a finite field, although the only thing we really use in this section is the assumption that the base field is perfect.

If the singularity test is positive, the projective closure $\tilde{E}$ has genus 0 and a unique singular point, which is rational over $\mathbb{F}$ provided $\mathbb{F}$ is perfect. We can use this point to find a rational parametrisation of all *nonsingular* points on E. It follows that the construction of rational points on a singular E is easy. Furthermore, the constructions given below give rise to efficient deterministic algorithms, whenever the operations of the field $\mathbb{F}$ are deterministically and efficiently computable, including the operation of taking a pth root if char $\mathbb{F} = p$.

We distinguish the cases of characteristic equal to 2 and unequal to 2.

Odd Characteristic. Let char $\mathbb{F}$ be unequal to 2, and let E be given by $y^2 = f(x)$ for some cubic polynomial f over $\mathbb{F}$. If f has a double zero x_2, then $(x_2, 0)$ is the unique singular point on E. Such a double zero must be in $\mathbb{F}$, as f has degree 3, and also the third zero of f must be rational.

Let $d = \gcd(f, f')$, where f' is the derivative of f. If d is constant, then f does not have a double zero and E is nonsingular. If d is linear, then its unique zero gives the double zero x_2. If d is quadratic, then char $\mathbb{F} \neq 3$ and f has a triple zero, which is equal to the unique zero of the linear polynomial d', the derivative of d. If d is cubic, then char $\mathbb{F} = 3$ and f has a triple zero $x_2 = \sqrt[3]{C} = C^{3^{m-1}}$, where m is the order of 3 modulo $|\mathbb{F}| - 1$.

Assume E is singular; by an $\mathbb{F}$-linear change of variables, we may assume that the singularity is at $(0, 0)$, and hence E is given by $y^2 = x^3 + Ax^2$ for some $A \in \mathbb{F}$. Now we parametrise E by projecting lines from the singular point: any such line has the form $y = \ell x$ with $\ell \in \mathbb{F}$, and it intersects E twice in $(0, 0)$ and once more in $(\ell^2 - A, \ell^3 - A\ell)$. This provides a rational parametrisation of E, which is clearly computable efficiently and deterministically.

Characteristic 2. Now let char $\mathbb{F}$ be 2, and let E be given by the generic cubic Weierstrass equation (3). We have $\frac{\partial}{\partial y} = a_1 x + a_3$, and hence E can be singular in two ways.

The first is to have $a_1 = a_3 = 0$; we get $\frac{\partial}{\partial x} = x^2 + a_4$, and the singular point will be $(\sqrt{a_4}, \sqrt{a_6})$, which we move to $(0,0)$ by a translation. We have already seen that the equation becomes $y^2 = x^3 + Ax^2$ for some $A \in \mathbb{F}$. We parametrise E just as in the case of characteristic not 2, and find that $\ell \mapsto (\ell^2 + A, \ell^3 + A\ell)$ is a rational parametrisation, computable efficiently and deterministically.

The second has $a_1 \neq 0$, and there we may assume $a_3 = a_4 = 0$ and $a_1 = 1$ by linear substitutions; by the equation $0 = \frac{\partial}{\partial x} = y + x^2$, we find that E has a singularity at $(0,0)$ if and only if in addition $a_6 = 0$, and that E is nonsingular otherwise. Assume E is singular; we now get the equation $y^2 + xy = x^3 + Ax^2$, for some $A \in \mathbb{F}$. The same way of parametrising shows that $\ell \mapsto (\ell^2 + \ell + A, \ell^3 + \ell^2 + A\ell)$ is a rational parametrisation, computable efficiently and deterministically.

Remark. For a singular Weierstrass equation, there even exists a parametrisation that is also a group homomorphism, but this map uses another affine patch of the equation and need not always be defined over the base field $\mathbb{F}$ (see Proposition 2.5 in [10]).

5   Elliptic Curves in Odd Characteristic

In this section, we prove Theorem 1 under the assumption that the base field $\mathbb{F}$ is a finite field of odd characteristic and that E is the curve given by a nonsingular Weierstrass equation (4) or (5). In particular, we let f be a cubic monic polynomial over $\mathbb{F}$ without double zeros. The considerations up to Lemma 7 in fact work over any field of characteristic not 2.

Let V denote the threefold

$$f(x_1)f(x_2)f(x_3) = y^2, \tag{8}$$

which, geometrically speaking, is the quotient of $E \times E \times E$ by the action of a Klein 4-group of automorphisms, namely those automorphisms that act as -1 on two components and as the identity on the third. We will obtain an explicit birational map from the affine line to a curve on V; see Lemmas 6 and 7 below.

Let $R = \mathbb{F}[x]/(f)$ be the residue class ring of polynomials over $\mathbb{F}$ modulo f; as f has no multiple zeros, the ring R is a finite étale algebra over $\mathbb{F}$ (cf. [3], Section V.6, especially Theorem 4 in V.6.7, and Section V.8). We denote by θ the class of x modulo f; thus θ generates R as an $\mathbb{F}$-algebra. If g is a polynomial in $\mathbb{F}[x]$ of degree d, then the *homogenisation* $g^{\mathrm{hom}} \in \mathbb{F}[x,y]$ of g is defined to be $y^d g(x/y)$.

Lemma 5. *For any $u, v, w \in \mathbb{F}$ satisfying $u + v + w + A = 0$, we have*

$$f(u)f(v)f(w) = (uv + uw + vw - B)^3 f\left(\frac{uvw + C}{uv + uw + vw - B}\right). \tag{9}$$

Proof. Let $\phi : \mathbb{F}^3 \to R$ be the map sending (u, v, w) to $(u - \theta)(v - \theta)(w - \theta)$. For any $u, v, w \in \mathbb{F}$, we have

$$\phi(u,v,w) = uvw - (uv + uw + vw)\theta + (u + v + w)\theta^2 - \theta^3$$
$$= (uvw + C) - (uv + uw + vw - B)\theta + (u + v + w + A)\theta^2, \quad (10)$$

because $f(\theta) = 0$.

Let H_A be the subspace of $\mathbb{F}^3$ of triples (u,v,w) satisfying $u + v + w + A = 0$. Then ϕ maps H_A into the subspace R_{lin} of R of elements that are *linear* in θ. Now if $\alpha - \beta\theta \in R_{\text{lin}}$, with $\alpha, \beta \in \mathbb{F}$, then we have

$$\text{Norm}_{R/\mathbb{F}}(\alpha - \beta\theta) = f^{\text{hom}}(\alpha, \beta) = \beta^3 f\left(\frac{\alpha}{\beta}\right). \quad (11)$$

In particular, $\text{Norm}(\alpha - \theta) = f(\alpha)$. Thus by taking norms, equation (10) is mapped to (9). ♦

Remark. The formula in the Lemma is a bit misleading in the sense that u, v, w will not perform the same functions as x_1, x_2, x_3 in (8).

Lemma 6. *Put $h(u,v) = u^2 + uv + v^2 + A(u + v) + B$, and define*

$$S : y^2 h(u,v) = -f(u), \quad (12)$$
$$\psi : (u,v,y) \mapsto \left(v, -A - u - v, u + y^2, f(u + y^2)h(u,v)\,y^{-1}\right). \quad (13)$$

Then ψ is a rational map from the surface S to V that is invertible on its image.

Proof. We break the symmetry in (9) by putting $w = -A - u - v$. We find $uv + uw + vw - B = -u^2 - uv - v^2 - A(u + v) - B = -h(u,v)$, and $uvw + C = u(uv + uw + vw - B) + u^2(-v - w) + uB + C = -uh(u,v) + f(u)$.

Now let (u,v,y) be a rational point on S such that $f(u) \neq 0$; it follows that $y \neq 0$ and $h(u,v) \neq 0$, as well. Then applying Lemma 5 with u, v, and $-A - u - v$ and using the equation of S twice gives us

$$f(v)f(-A-u-v)y^2 = h(u,v)^2 f\left(\frac{-uh(u,v) + f(u)}{-h(u,v)}\right) = h(u,v)^2 f(u + y^2). \quad (14)$$

We multiply by $f(u + y^2)$ and divide by y^2 to see that we have a rational point on the threefold V.

From the definition of the map, it is clear that u, v, y can be computed from the image of (u,v,y) on V, so that ψ is invertible on its image. ♦

Lemma 7. *There exists a deterministic algorithm that, given a finite field $\mathbb{F}$ of q elements, where q is odd, a nonsingular cubic Weierstrass equation $y^2 = f(x)$ over $\mathbb{F}$, and an element $u \in \mathbb{F}$ such that*

$$f(u) \neq 0 \quad and \quad \tfrac{3}{4}u^2 + \tfrac{1}{2}Au + B - \tfrac{1}{4}A^2 \neq 0,$$

computes a rational map

$$\phi : \mathbb{A}^1 \to S$$

defined over $\mathbb{F}$ that is invertible on its image, in time polynomial in $\log q$. Here the surface S is as defined in (12).

Proof. Note that we may assume $A = 0$ whenever char $\mathbb{F} \neq 3$; this could facilitate reading the proof.

We fix a $u \in \mathbb{F}$ that satisfies the requirements given above; then the equation (12) of the surface S specialises to a nondegenerate quadratic equation

$$\left[y(v + \tfrac{1}{2}u + \tfrac{1}{2}A)\right]^2 + \left[\tfrac{3}{4}u^2 + \tfrac{1}{2}Au + B - \tfrac{1}{4}A^2\right]y^2 = -f(u), \qquad (15)$$

which is of the form (2) for the variables $z = y(v + \tfrac{1}{2}u + \tfrac{1}{2}A)$ and y.

Now use Theorem 4 to compute a rational point (z_0, y_0) on (15), and let $t \mapsto (\alpha(t), \beta(t))$ be the corresponding rational parametrisation of the conic (15), still for the variables (z, y) (see [9, Sect. 1.2] or [6, Sect. 1.1]). We have $v = z/y - u/2 - A/2$; therefore the map

$$\phi : t \mapsto \left(u, \frac{\alpha(t)}{\beta(t)} - \frac{u}{2} - \frac{A}{2}, \beta(t)\right) \qquad (16)$$

parametrises all rational points on S with the given u-coordinate, except $(u, z_0/y_0 - u/2 - A/2, y_0)$, because this point corresponds to $t = \infty$. ◆

After having given the ingredients of the construction of rational points on the threefold V, we ask ourselves how many rational points will be found in this way. The bound of $(q - 4)/16$ given by Lemma 9 can probably be improved.

Definition 8. *We define two points $P = (x_1, x_2, x_3, y)$ and $P' = (x_1', x_2', x_3', y')$ on V to be* disjoint *if the sets $\{x_1, x_2, x_3\}$ and $\{x_1', x_2', x_3'\}$ are disjoint.*

Lemma 9. *Let $\mathbb{F}$ be a finite field of q elements, let $u_0 \in \mathbb{F}$ satisfy the requirements of Lemma 7, and let $\phi : \mathbb{A}^1 \to S$ be the corresponding map. Let ψ be the map from Lemma 6.*

Then there is a subset $T \subseteq \mathbb{F}$ of cardinality at least $(q - 4)/16$, such that for all distinct $t, t' \in T$, the points $\psi \circ \phi(t)$ and $\psi \circ \phi(t')$ are disjoint.

Proof. Let u_0 be as in the Lemma; we fix it for the whole proof. The corresponding map ϕ is well-defined except perhaps in two values of t where $\beta(t) = 0$, and two others where $(\alpha(t), \beta(t))$ lies at infinity. It follows that the image of ϕ contains at least $q - 4$ points.

Let $\psi : S \to V$ be the map from Lemma 6; for two points $P = (u_0, v, y)$ and $P' = (u_0, v', y')$ on S, we want to find sufficient conditions for $\psi(P)$ and $\psi(P')$, or, equivalently, the sets $\{v, -A - u_0 - v, u_0 + y^2\}$ and $\{v', -A - u_0 - v', u_0 + y'^2\}$, to be disjoint.

Note that $v \mapsto -A - u - v$ and $y \mapsto -y$ are automorphisms of S; these automorphisms generate a Klein 4-group G. If P and P' share an orbit under G, then $\psi(P)$ and $\psi(P')$ cannot be disjoint. Note there is at most one orbit under G for any given value of y^2, as $u = u_0$ is assumed to be fixed.

Assume now $\psi(P)$ and $\psi(P')$ are not disjoint. A case-by-case analysis shows that y'^2 is equal to one of y^2, $v - u_0$, $-A - 2u_0 - v$, or $-f(u_0)/h(u, u + y^2)$, where (12) is used to derive the last option.

Let us define a graph on the set of G-orbits on S with $u = u_0$ by putting an edge between two distinct orbits X and X' if there are non-disjoint points $P \in \psi(X)$ and $P' \in \psi(X')$. The above reasoning shows that in this graph,

every vertex has at most three neighbours. We want to find a maximal set Σ of pairwise nonadjacent vertices, meaning that if $X \neq X' \in \Sigma$, then all points in $\psi(X)$ are disjoint from all points in $\psi(X')$. Such a set Σ can be constructed greedily by selecting any vertex, adding it to Σ, deleting it and its neighbours with all the incident edges from the graph, and repeating this process until no vertices remain. As we include at least every fourth G-orbit, and as the orbits contain at most 4 points, we see that at least a fraction of $1/16$ of the points in the image of ϕ have pairwise disjoint images under ψ.   ♦

Proof of Theorem 1 (odd characteristic). Let $\mathbb{F}$ be a finite field of cardinality greater than 5, so that there exists some $u \in \mathbb{F}$ satisfying the conditions of Lemma 7; we fix such a u for the rest of the proof.

We first show how to compute rational points on the elliptic curve E, which we assume to be given by an equation $y^2 = f(x)$, for some cubic polynomial f with no double roots. By composing the maps ψ from Lemma 6 and ϕ from Lemma 7, we can compute rational points on the threefold V. Then, given a rational point $P = (x_1, x_2, x_3, y)$ on V, we apply the algorithm from Theorem 3 to $f(x_i)$ for $i = 1, 2, 3$ to compute a square root c of $f(x_i)$ for, say, $i = i_0$. Having done this, we see that (x_{i_0}, c) is a rational point on the elliptic curve E.

The next question is whether two different points on V can lead to the same point on E. This is rather subtle; it is even the case that one point on V can lead to several points on E, for example when $f(x_i)$ has odd order for $i = 1, 2, 3$. However, it is clear that if two points on V are *disjoint* in the sense defined above, then they can only give rise to different points on E. Indeed, the x-coordinate of the point on E computed from $P = (x_1, x_2, x_3, y)$ is either x_1, x_2, or x_3. We can therefore use Lemma 9 to show that, if we let the argument t of $\psi \circ \phi$ run through all of $\mathbb{F}$, then at least $(q-4)/16$ valid x-coordinates of points on E follow from the obtained rational points on V. This gives $(q - 4)/8$ rational points on E, as claimed.   ♦

Remark. It is an interesting question whether the surface S given in Lemma 6 is rational over the ground field $\mathbb{F}$. This question is addressed in [7], for any base field of characteristic different from 2. If we homogenise the equation for S given in (15), we obtain a diagonal ternary quadratic form over the function field $\mathbb{F}(u)$, whose coefficients have degrees 0, 2, and 3. Using the notation and definitions given in [7], we see that the equation has minimal *index* 6 if we use the weights $(3, 2, 1)$ for the variables, whereas a rational surface of this form must have index at most 3 for some weight vector. Therefore, unless some factors of the discriminant of the equation are removable, S is not rational over $\mathbb{F}$.

6   Elliptic Curves in Characteristic 2

In this section we complete the proof of Theorem 1 under the assumption that the characteristic of the base field $\mathbb{F}$ is 2 and that E is given by a nonsingular Weierstrass equation.

Recall that by [10, Appendix A] we know that E has a Weierstrass equation of one of two following forms:

$$Y^2 + a_3 Y = X^3 + a_4 X + a_6 \quad \text{if } j(E) = 0,$$
$$Y^2 + XY = X^3 + a_2 X^2 + a_6 \quad \text{if } j(E) \neq 0.$$

In the case when $\mathbb{F}$ is finite of order 2^r, let Tr stand for the trace map from $\mathbb{F}$ to $\mathbb{F}_2$, which is defined by

$$\mathrm{Tr}_{\mathbb{F}/\mathbb{F}_2}(x) := x + x^2 + x^{2^2} + \cdots + x^{2^{r-1}} .$$

For motivation, consider the problem of finding rational points on

$$Y^2 + Y = f(X).$$

Lemma 10. *If f is linear in X, then there exists a deterministic polynomial-time algorithm that returns a point of $Y^2 + Y = f(X)$ over a finite field $\mathbb{F}$.*

Proof. It is well known that the valid X-coordinates are exactly $x \in \mathbb{F}$ satisfying $\mathrm{Tr}(f(x)) = 0$ [2, Sect. 6.6]. First precompute $a \in \mathbb{F}$ such that $\mathrm{Tr}(f(a)) = 1$. Since $x \mapsto \mathrm{Tr}(f(x))$ is a linear map over $\mathbb{F}_2$, we can deterministically compute the required a using linear algebra. Now, one of x or $x + a$ must be a valid X-coordinate.

Given such an x, it remains to solve for Y. Here we have an advantage over the case of odd characteristic in that there exist deterministic polynomial-time algorithms for solving quadratics ([2, Chap. 6], [1, Sect. 7.4]).   ◆

For more general f, the new idea is to look for points on the threefold

$$f(x_1) + f(x_2) + f(x_3) = y^2 + y .$$

Elements of the form $y^2 + y$ are exactly those in Ker(Tr), and form an index two subgroup of $\mathbb{F}^+$. Thus one of the three terms must itself be of the form $y^2 + y$.

With this in mind, we define

$$g(x) = x^{-2} \cdot (x^3 + a_2 x^2 + a_6), \text{ and}$$
$$h(x) = x^3 + a_4 x + a_6 .$$

Now let V_1 and V_2 be threefolds given by the equations

$$V_1 : g(x) + g(y) + g(z) = w^2 + w$$
$$V_2 : h(x) + h(y) + h(z) = w^2 + a_3 w .$$

These have the same geometric definition as the threefold V given in the previous section.

As in the odd characteristic case, we will construct a computable rational map from a parametrisable surface to the appropriate threefold. Once we have

a point on the threefold it will be easy to get rational points on E. The surfaces we need are given by the equations

$$S_1 : x + y + xy(x + y)^{-1} + a_2 = w^2 + w$$
$$S_2 : x^2 y + y^2 x + a_6 = w^2 + a_3 w \ .$$

Lemma 11. *Let $\mathbb{F}$ be a field of characteristic 2. There exist rational maps $\phi_1 : S_1 \to V_1$ and $\phi_2 : S_2 \to V_2$ over $\mathbb{F}$ which are invertible on their images, given by*

$$\phi_1 : (x, y, w) \mapsto (x, \ y, \ xy(x + y)^{-1}, \ w)$$
$$\phi_2 : (x, y, w) \mapsto (x, \ y, \ x + y, \ w) \ .$$

Proof. First consider ϕ_1, the map that will be used in the case when $j(E) \neq 0$. Recall that $g(x) = x + a_2 + a_6 x^{-2}$. We have

$$g(x) + g(y) + g\left(\frac{xy}{x + y}\right) = x + y + \frac{xy}{x + y} + 3a_2 + a_6 \left(\frac{1}{x} + \frac{1}{y} + \frac{x + y}{xy}\right)^2$$
$$= x + y + \frac{xy}{x + y} + a_2$$
$$= w^2 + w$$

since (x, y, w) is a point on S_1. Hence $(x, \ y, \ xy(x + y)^{-1}, \ w)$ is a point on V_1.

Next consider ϕ_2, the map that will be used when $j(E) = 0$. We have

$$h(x) + h(y) + h(x + y) = x^3 + a_4 x + y^3 + a_4 y + (x + y)^3 + a_4(x + y) + 3a_6$$
$$= x^2 y + y^2 x + a_6$$
$$= w^2 + a_3 w$$

since (x, y, w) is a point on S_2.

Note that given a point in the image of one of these maps we can trivially find its preimage on the surface, so that both maps are invertible on their images. ♦

Remark. A useful geometric interpretation of these maps is that the image of ϕ_1 is contained in the intersection of V_1 with $x^{-1} + y^{-1} + z^{-1} = 0$, while the image of ϕ_2 is contained in the intersection of V_2 with $x + y + z = 0$.

These maps now play a critical role in the following main theorem.

Theorem 12. *There exists a deterministic polynomial-time algorithm that, given a finite field $\mathbb{F}$ of characteristic 2 with more than 4 elements and an elliptic curve E over $\mathbb{F}$, computes a nontrivial rational point on E.*

Proof. There are two cases to consider, since E can either have j-invariant zero or nonzero. In both cases our strategy is to deterministically find points on the appropriate surface, map them to the threefold, and from there get a point on E.

First assume that $j(E) \neq 0$. For arbitrary c the equation

$$x + y + \frac{xy}{x+y} = c$$

is equivalent to the genus 0 curve $C: \ x^2 + y^2 + xy + c(x+y) = 0$ except when $x = y$. However, if (x,y) is a point on C with $x = y$ then it must be the point $(0,0)$, so not much is lost. We have the generic solution $(0,c)$ and from this get all points of C through the rational parametrisation

$$y = tx + c$$
$$x = \frac{tc}{1+t+t^2} \ .$$

Thus we have a family of rational points on S_1 parametrised by t and w which can be mapped to points on V_1 via ϕ_1. It now remains to compute rational points of E.

For $a \in \mathbb{F}^*$ consider the set

$$\{u^2 + au \mid u \in F\}.$$

This set is an additive subgroup of $\mathbb{F}^+$ of index 2, so if $g(x) + g(y) + g(z) = w^2 + w$ then at least one of $g(x), g(y), g(z)$ is itself of the form $u^2 + u$. Discover which it is, call it x, and deterministically solve the quadratic to find u. From $u^2 + u = x^{-2}(x^3 + a_2 x^2 + a_6)$ we now have

$$(ux)^2 + x(ux) = x^3 + a_2 x^2 + a_6$$

and hence a point on E.

Suppose instead that $j(E) = 0$. We wish to compute points on S_2. Taking $y = u^2$, we transform the equation for S_2 as following:

$$xy(x+y) + a_6 = w^2 + a_3 w$$
$$x^2 u^2 + au^4 + a_6 = w^2 + a_3 w$$
$$a_3 xu + xu^4 + a_6 = (w + xu)^2 + a_3(w + xu) \ .$$

Now, choose y and compute its square root u (possible deterministically since squaring is an automorphism). There are at most four bad choices of y to avoid, corresponding to the roots of $u^4 + a_3 u$. If $u^4 + a_3 u \neq 0$, the equation $x(a_3 u + u^4) + a_6 = z^2 + a_3 z$ is linear in x and hence for any given z, we easily compute the unique value for x. Now the point $(x, y, z + xu)$ is a point on S_2, which we map to V_2 via ϕ_2.

It remains to find a point on E. Mirroring the argument in the previous case, one of $h(x), h(y)$, and $h(z)$ has the form $u^2 + a_3 u$. Discover which it is, call it x, and solve the quadratic $u^2 + a_3 u = h(x)$ for u. Output (x, u) as a rational point on E. ♦

Remark. This argument can be generalised to work over any perfect characteristic 2 field, but only gives an algorithm when the maps $u \mapsto u^2$ and $u \mapsto u^2 + au$ are algorithmically invertible.

An important question to analyze is how many of the $\mathbb{F}$-rational points of E are obtained by this algorithm. The next theorem will demonstrate that the number is quite large, in particular at least a constant proportion. We define disjointness for points on V_1 as in Definition 8.

Theorem 13. *Let $\mathbb{F}$ be a finite field of order $q = 2^r$ with $q > 4$. The number of disjoint points of V_1 that arise from Theorem 12 is at least $(q - 4)/6$.*

Proof. Throughout, assume that the parameter w from Theorem 12 is fixed. Allowing different values could improve the bound, but that analysis has not yet been done.

It was noted before that S_1 can be transformed into a genus 0 curve C : $x^2 + y^2 + xy + c(x + y) = 0$, with C having only gained the point $(0,0)$. Let $C'(\mathbb{F})$ be the points of C except for $(0,0)$, $(c,0)$, and $(0,c)$.

It can easily be confirmed that if (x,y) is a point on C', then $\sigma_1(x,y) = (x, \ xy(x+y)^{-1})$ and $\sigma_2(x,y) = (y,x)$ are points on C'. We conclude that the group $G = \langle \sigma_1, \sigma_2 \rangle$ acts on $C'(\mathbb{F})$, is isomorphic to Sym(3), and splits the points of C' into orbits of size 6. For the last statement, note that $x = y$ implies $(x,y) = (0,0)$ and $y = xy(x+y)^{-1}$ implies $y = 0$. Thus the stabiliser in Sym(3) of any point has index 6, giving an orbit of size 6.

Any coordinate only appears in its orbit, and each orbit yields the same set $(x, \ y, \ xy(x+y)^{-1})$. Thus each orbit when mapped via ϕ_1 yields a disjoint point on V_1.

It remains to count the number of orbits. If r is odd, $t^2 + t + 1$ is irreducible over $\mathbb{F}$ and hence all $t \in \mathbb{F}$ are valid. Thus C has $q+1$ points, but after discarding $(0,0)$, $(c,0)$, and $(0,c)$ we are left with $(q-2)/6$ orbits. If r is even, $t^2 + t + 1$ splits and hence there are $q - 2$ valid t, leaving us with $(q-4)/6$ orbits. ◆

Remark. We note that the case with $j(E) = 0$ yields a similar bound, since fixing w in S_2 yields a curve of genus 0 that also breaks up into orbits of size 6, each element of the orbit resulting in the same triples $(x, \ y, \ x + y)$.

Proof of Theorem 1 (even characteristic). Let $\mathbb{F}$ be a finite field of order $q = 2^r$ with $q > 4$, and let E be a nonsingular elliptic curve over $\mathbb{F}$. From Theorem 12 we obtain a deterministic polynomial-time algorithm that computes points on E. From Theorem 13 we see that this algorithm results in at least $(q-4)/6$ disjoint points on the threefold. This yields at least $(q-4)/6$ x-coordinates of E, and hence at least $(q-4)/3$ points of E.

This completes the proof of Theorem 1. ◆

Remark. If $\mathbb{F}$ is too small we simply check all pairs $(x,y) \in \mathbb{F}^2$ and obtain the set $E(\mathbb{F})$. This also holds for $\mathbb{F}$ of odd characteristic.

References

1. Bach, E. and Shallit, J.: Algorithmic Number Theory. The MIT Press, Cambridge (1996)
2. Berlekamp, E.R.: Algebraic coding theory. McGraw-Hill Book Co., New York (1968)
3. Bourbaki, N.: Algebra II. Chapters 4–7. Elements of Mathematics. Springer-Verlag, Berlin (2003) Translated from the 1981 French edition by P. M. Cohn and J. Howie, Reprint of the 1990 English edition
4. Bumby, R.T.: Sums of four squares. In: Number theory (New York, 1991–1995). Springer, New York (1996) 1–8
5. Cohen, H.: A course in computational algebraic number theory. Volume 138 of Graduate Texts in Mathematics. Springer-Verlag, Berlin (1993)
6. Reid, M.: Undergraduate algebraic geometry. Volume 12 of London Mathematical Society Student Texts. Cambridge University Press, Cambridge (1988)
7. Schicho, J.: Proper parametrization of surfaces with a rational pencil. In: Proceedings of the 2000 International Symposium on Symbolic and Algebraic Computation (St. Andrews), New York, ACM (2000) 292–300 (electronic)
8. Schoof, R.: Elliptic curves over finite fields and the computation of square roots mod p. Math. Comp. **44**(170) (1985) 483–494
9. Shafarevich, I.R.: Basic algebraic geometry. 1. Second edn. Springer-Verlag, Berlin (1994) Varieties in projective space, Translated from the 1988 Russian edition and with notes by Miles Reid.
10. Silverman, J.H.: The arithmetic of elliptic curves. Volume 106 of Graduate Texts in Mathematics. Springer-Verlag, New York (1992). Corrected reprint of the 1986 original
11. Skałba, M.: Points on elliptic curves over finite fields. Acta Arith. **117**(3) (2005) 293–301
12. van de Woestijne, C.: Deterministic equation solving over finite fields. PhD thesis, Universiteit Leiden (2006)