FINDING POINTS ON CURVES OVER FINITE FIELDS

Joachim von zur Gathen
FB Mathematik-Informatik, Universität–GH Paderborn
33095 Paderborn, Germany
gathen@uni-paderborn.de

Igor Shparlinski
School of MPCE, Macquarie University
Sydney, NSW 2109, Australia
igor@mpce.mq.edu.au

Preliminary version
November 6, 1995

Abstract. We solve two computational problems concerning plane algebraic curves over finite fields: generating an (approximately) uniform random point, and finding all points deterministically in amortized polynomial time (over a prime field, for non-exceptional curves).

1. Introduction

Let q be a prime power, $\mathbb{F}_q$ a finite field with q elements, $f \in \mathbb{F}_q[x, y]$ of degree n, and $\mathcal{C} = \{(a, b) \in \mathbb{F}_q^2 : f(a, b) = 0\} = \{f = 0\}$ the plane curve defined by f. We consider two problems of finding points on this curve: probabilistically finding a uniformly distributed random point, and deterministically computing all its points.

Curves over finite fields play a role in several applications: factoring integers with elliptic curves, testing primality with elliptic curves (or more general algebraic varieties), algebro-geometric Goppa codes, and fast multiplication over finite fields. For these applications, special methods for finding points (if needed) are used. This paper presents the first general and systematic approach to the problem, to the authors' knowledge.

Throughout this paper, we will assume that f is squarefree, and denote by σ the number of absolutely irreducible components of $\mathcal{C}$ which are defined over $\mathbb{F}_q$. The case of an *exceptional curve*, corresponding to $\sigma = 0$, needs special treatment and is dealt with in Section 5. So now we assume that $\sigma \geq 1$.

In Section 2 we study the probabilistic variant of our question: generating uniform random points on $\mathcal{C}$. Let $\pi: \mathcal{C} \to \mathbb{F}_q$ be the projection onto the first coordinate, for $0 \leq i \leq n$ let

$$R_i = \{a \in \mathbb{F}_q : \#\pi^{-1}(\{a\}) = i\}$$

be the set of points with exactly i preimages, and $r_i = \#R_i$. A first attempt to generate a random point on $\mathcal{C}$ might be to pick a random point $a \in \mathbb{F}_q$, and then a random $b \in \mathbb{F}_q$ with $(a, b) \in \mathcal{C}$, if such a b exists. For $(a, b) \in \pi^{-1}(R_i)$, the probability of being chosen is $1/(qi)$, which in general is not the uniform probability $1/\#\mathcal{C} = 1/\sum_{1 \leq j \leq n} jr_j$.

In von zur Gathen & Shparlinski (1994), we have shown how to compute approximations to the r_i. Given these, we present in Section 2 a method to generate approximately uniform random points on $\mathcal{C}$. On the other hand, we also show that from any method that generates (exactly or approximately) uniform random points on $\mathcal{C}$, we obtain approximations to the r_i. The running time of our algorithm is exponential in n, and polynomial in $\log q$ and the quality of the approximation. For a certain fairly rough approximation quality, the running time becomes polynomial in all parameters.

The famous theorem of Weil says that

$$| \#\mathcal{C} - \sigma q | \leq n^2 q^{1/2}. \tag{1.1}$$

With deterministic methods, the "brute force" approach to computing all points on $\mathcal{C}$ via finding, for each $a \in \mathbb{F}_q$, all $b \in \mathbb{F}_q$ with $f(a, b) = 0$, takes $O\tilde{~}(n^2 q^{3/2})$ operations in $\mathbb{F}_q$, using the fastest known deterministic algorithms to factor the univariate polynomial $f(a, y)$, for all $a \in \mathbb{F}_q$ (Shoup 1990; Section 1.1 of Shparlinski 1992, von zur Gathen & Shoup 1992). We present in Section 3 a deterministic method that uses $O\tilde{~}(n^5 q)$ operations, i.e., polynomial time per point. The central tool for our estimates is a bound of Perel'muter's (1969) on a certain exponential sum. In order to use this, we have to study in Section 4 some geometric and arithmetic properties of the fibre square $\mathcal{C} \times_\pi \mathcal{C}$.

Our approach only works in the case of a prime field $\mathbb{F}_q$, with $q = p$ prime, and does not work for exceptional curves.

Shoup (1990) has exhibited a deterministic univariate factoring algorithm which for almost all polynomials runs in polynomial time.

Our deterministic result has two interpretations: the first is that the members of a "small" parametrized family $f(a, y)$ of univariate polynomials, for all $a \in \mathbb{F}_p$, can be factored deterministically in (amortized) polynomial time. The second is that all points on a plane algebraic curve over $\mathbb{F}_p$ can be found deterministically in (amortized) polynomial time.

Finally, Section 5 presents a discussion of the case of exceptional curves which has been excluded in the other sections.

2. Generating uniform random points

Recall from the introduction the distribution statistics r_i and R_i for the projection $\pi: \mathcal{C} \to \mathbb{F}_q$, so that $\sum_{1 \leq i \leq n} ir_i = \#\mathcal{C}$. We first assume that $r_0, \ldots, r_n$ are known, and then have the following algorithm.

Algorithm 2.1. *Random point.*
Input: *$f \in \mathbb{F}_q[x, y]$ of degree n, and $r_0 \ldots, r_n$.*
Output: *A uniform random point (a, b) on $\mathcal{C} = \{f = 0\} \subseteq \mathbb{F}_q^2$.*

1. *Choose $i \in \{1, \ldots, n\}$ at random with probability $ir_i/\#\mathcal{C}$.*
2. *Choose $a \in R_i$ at random, by trying random $a \in \mathbb{F}_q$ until an $a \in R_i$ is found.*
3. *Choose $b \in \pi^{-1}(\{a\})$ at random, and return (a, b).*

A basic subroutine, used in step 2, is to determine for given $a \in \mathbb{F}_q$ the i with $a \in R_i$. To this end, we compute with $O(\mathsf{M}(n) \log(nq))$ operations

$$f_a^* = \gcd(y^q - y, f(a, y)) \in \mathbb{F}_q[y], \tag{2.1}$$

so that $a \in R_i$, where $i = \deg f_a^*$.

Theorem 2.2. *Let $f \in \mathbb{F}_q[x, y]$ be squarefree and non-exceptional of degree n, and $q \geq 4n^4$. Then the algorithm correctly returns a uniformly random point of $\mathcal{C} = \{f = 0\}$. It uses an expected number of $O(\log(nq))$ random bit choices, at most n^2 random choices in $\mathbb{F}_q$, and $O(n^2\mathsf{M}(n) \log(nq))$ operations in $\mathbb{F}_p$.*

PROOF. For $(a,b) \in \mathcal{C}$ with $a \in R_i$, the probability of being chosen is

$$\frac{ir_i}{\#\mathcal{C}} \cdot \frac{1}{r_i} \cdot \frac{1}{i} = \frac{1}{\#\mathcal{C}}.$$

If i is chosen in step 1, then the expected number of $a \in \mathbb{F}_q$ that we try in step 2 is q/r_i. Thus the total expected number of trials is

$$\sum_{1 \le i \le n} \frac{ir_i}{\#\mathcal{C}} \cdot \frac{q}{r_i} \le \frac{n^2 q}{2\#\mathcal{C}} \le n^2. \tag{2.2}$$

The last inequality

$$\#\mathcal{C} \ge q - n^2 q^{1/2} \ge q/2 \tag{2.3}$$

follows from (1.1) and the assumption that $\sigma \ge 1$.

Once we have $a \in R_i$, we factor f_a^* completely into linear factors with $O(n \log q)$ operations and $O(\log n)$ random choices in $\mathbb{F}_q$ (Cantor & Zassenhaus 1981) and choose a random root as b. (In fact, it is sufficient to isolate a single random root b of f_a^*.) □

Next we only assume that we have approximations ρ_i to r_i and μ to $\#\mathcal{C}$, with

$$\mid \rho_i - r_i \mid \le \epsilon_i \text{ for } 0 \le i \le n, \mid \mu - \#\mathcal{C} \mid \le \gamma, \tag{2.4}$$

and execute Algorithm Random Point with these values, i.e., ρ_i for r_i and μ for $\#\mathcal{C}$. A natural choice would be $\mu = \sum_{1 \le i \le n} i\rho_i$, but we do not insist on this.

LEMMA 2.3. *Let $1 \le i \le n$, and $(a,b) \in \mathcal{C}$ with $a \in R_i$. Then for the probability α that (a,b) is chosen by the above modified algorithm we have*

$$\mid \alpha - \frac{1}{\#\mathcal{C}} \mid \le \frac{1}{\mu}\Big(\frac{\epsilon_i}{r_i} + \frac{\gamma}{\#\mathcal{C}}\Big). \tag{2.5}$$

PROOF. Denote by E_i and F_a the events that i and a are chosen in steps 1 and 2, respectively. Then we have

$$\alpha = \text{prob}(E_i) \cdot \text{prob}(F_a \mid E_i) \cdot \text{prob}(b \text{ is chosen} \mid F_a) = \frac{i\rho_i}{\mu} \cdot \frac{1}{r_i} \cdot \frac{1}{i} = \frac{\rho_i}{r_i \mu},$$

$$\left|\alpha - \frac{1}{\#\mathcal{C}}\right| = \left|\frac{\rho_i}{r_i\mu} - \frac{1}{\#\mathcal{C}}\right| \le \frac{1}{\mu}\Big(\mid \frac{\rho_i}{r_i} - 1 \mid + \mid 1 - \frac{\mu}{\#\mathcal{C}} \mid\Big) \le \frac{1}{\mu}\Big(\frac{\epsilon_i}{r_i} + \frac{\gamma}{\#\mathcal{C}}\Big). \quad \square$$

In order to implement these ideas, we still need to compute the r_i's or approximations to them. To calculate them exactly, it seems that one has to completely enumerate all points on $\mathcal{C}$, as in Section 3 (except that there is no need to factor any f_a^*), at a cost of $\Omega(nq)$ operations.

A faster calculation is given by Lemma 3.1 of von zur Gathen & Shparlinski (1994). It yields probabilistic approximations ρ_i and μ such that (2.4) holds with probability at least $1-\delta$, where

$$\epsilon_i = 2\Big(\frac{qr_i}{h}\log\frac{2}{\delta}\Big)^{1/2}, \quad \gamma = \Big(\frac{q\#\mathcal{C}}{h}\cdot 2n(n+1)\log(2n/\delta)\Big)^{1/2}. \tag{2.6}$$

Here h is a parameter of the method, and the cost is h gcd computations as in (2.1), for a total of $O\big(h\ \mathsf{M}(n)\log(nq)\big)$ operations in $\mathbb{F}_q$.

For a cost estimate, we assume that we have a prescribed tolerance of the error in (2.5). By Weil's theorem (1.1), $\#\mathcal{C}$ is about σq, and we write this requirement as

$$\mid \alpha - \frac{1}{\#\mathcal{C}} \mid \leq \frac{\beta}{q},$$

where we think of the *quality of approximation* β as being small, say $\beta \approx q^{-1/4}$.

Let $0 \leq i \leq n$. We say that R_i and r_i are *large* if

$$r_i \geq q/(2\cdot n!), \tag{2.7}$$

and that $P=(a,b)\in\mathcal{C}$ is *over a large fibre set* if $a \in R_i$ and R_i is large. The motivation for this definition will be discussed after the next theorem.

THEOREM 2.4. *Let $q \geq 16n^{4n}(n!)^2$, let $\mathcal{C} \subseteq \mathbb{F}_q^2$ be a non-exceptional curve of degree $n \geq 9$, and $\beta, \delta > 0$. Assume that*

$$\beta \leq \frac{3}{4}\Big(\frac{n!}{n^3}\Big)^{1/2}. \tag{2.8}$$

With

$$O\Big(n!\,\mathsf{M}(n)\log(nq)\log(n\delta^{-1})\beta^{-2}\Big)$$

operations in $\mathbb{F}_q$ we can compute approximations ρ_i to r_i, for $1 \leq i \leq n$, with the following property, for any $P \in \mathcal{C}$ over a large fibre set. If α is the probability that Algorithm Random Point, when run with ρ_i for r_i and μ for $\#\mathcal{C}$, outputs P, then

$$\mid \alpha - \frac{1}{\#\mathcal{C}} \mid \leq \frac{\beta}{q}$$

holds with probability at least $1-\delta$.

PROOF. We set

$$h = \left\lceil \frac{81 \cdot n!}{\beta^2} \cdot \log \frac{2n}{\delta} \right\rceil,$$

and compute with $O(h\,\mathsf{M}(n)\log(nq))$ operations in $\mathbb{F}_q$ values ρ_i and μ such that (2.4) and (2.6) hold. Let $P = (a, b) \in \mathcal{C}$ and $1 \le i \le n$ be such that $a \in R_i$. Then by Lemma 2.3 we have

$$\mid \alpha - \frac{1}{\#\mathcal{C}} \mid \le \frac{1}{\mu}\Big(\frac{\epsilon_i}{r_i} + \frac{\gamma}{\#\mathcal{C}}\Big), \tag{2.9}$$

and it is sufficient to show that the righthand side is at most β/q.

We first need some bounds on the parameters in (2.9). As in (2.2), we have

$$\#\mathcal{C} \ge q - n^2 q^{1/2} \ge q/2,$$

since $q \ge 16n^{4n}(n!)^2 \ge 4n^4$. Since

$$\beta \le \frac{3}{4}\Big(\frac{n!}{n^3}\Big)^{1/2} < \frac{3}{2\sqrt{2n}}\,\Big(\frac{n!}{n+1}\Big)^{1/2},$$

we find

$$\begin{aligned} \mid \#\mathcal{C} - \mu \mid \quad &\le \quad \gamma \;\le \Big(\frac{q \cdot nq \cdot \beta^2}{81\, n! \log(2n/\delta)} \cdot 2n(n+1)\log\frac{2n}{\delta}\Big)^{1/2} \\ &= \quad \frac{nq\beta}{9} \cdot \Big(\frac{2(n+1)}{n!}\Big)^{1/2} \;\le\; \frac{q}{6}, \end{aligned} \tag{2.10}$$

$$\mu \ge \#\mathcal{C} - \gamma \ge q - n^2 q^{1/2} - \gamma \ge q/2 - q/6 = q/3.$$

Substituting the bound on r_i and all these estimates in (2.9), we find

$$\begin{aligned} &\frac{1}{\mu}\Big(\frac{\epsilon_i}{r_i} + \frac{\gamma}{\#\mathcal{C}}\Big) \le \frac{3}{q}\Big(\frac{q}{h}\Big)^{1/2}\Big(\log\frac{2n}{\delta}\Big)^{1/2}\Big(\frac{2}{r_i^{1/2}} + \Big(\frac{2n(n+1)}{\#\mathcal{C}}\Big)^{1/2}\Big) \\ &\le \frac{3}{q^{1/2}} \cdot \Big(\frac{\beta^2}{81 \cdot n! \cdot \log(2n/\delta)}\Big)^{1/2} \cdot (\log\frac{2n}{\delta})^{1/2} \cdot \frac{1}{q^{1/2}}\Big((8 \cdot n!)^{1/2} + 2(n^2+n)^{1/2}\Big) \\ &\le \frac{3\beta}{9q(n!)^{1/2}} \cdot (9 \cdot n!)^{1/2} = \frac{\beta}{q}. \quad \square \end{aligned}$$

The bound (2.8) on β was chosen so that the last inequality in (2.10) holds. If we choose β maximal under this condition, say $\beta = c(n!/n^3)^{1/2}$ for some constant c, then the running time of our algorithm is polynomial in $n \log q$; the quality of approximation is then rather poor.

What about those r_i that are not large? By Lemma 2.3 of von zur Gathen & Shparlinski (1994), we have

$$| r_i - q\lambda_i | \leq 2n^{2n}q^{1/2}$$

for some $\lambda_i \in \mathbb{Q}$ with $n!\lambda_i \in \mathbb{Z}$. It follows that either $n!\lambda_i \geq 1$ and

$$r_i \geq q/n! - 2n^{2n}q^{1/2} \geq q/2n!, \tag{2.11}$$

or $n!\lambda_i \leq 0$ and

$$r_i \leq 2n^{2n}q^{1/2}. \tag{2.12}$$

Thus each r_i is either large in the sense of (2.7), or very small. The total fraction of $P = (a, b) \in \mathcal{C}$ that are not over a large fibre set is at most

$$\frac{n \cdot 2n^{2n}q^{1/2}}{\#\mathcal{C}} \leq \frac{2n^{2n+1}q^{1/2}}{q/2} \leq \frac{n}{n!} = \frac{1}{(n-1)!},$$

under the assumptions of Theorem 2.4. For such a point, r_i might be only 1, say, and thus the estimates (2.4) and (2.6) do not yield a useful bound, close to 1, on the relative error ρ_i/r_i. In fact, each R_i corresponds to the points on a certain curve $\mathcal{D}_i$, the ith *fibre power* of $\mathcal{C}$, and for those R_i that are not large, this curve $\mathcal{D}_i$ is exceptional. It is not surprising that our methods are, again, not adequate to deal with this exceptional situation.

At our current state of knowledge, computing the r_i's exactly is feasible only for small values of q, and computing approximations only for not too large values of n and q. Thus it seems somewhat unfortunate that we used such approximations in our uniform random generation algorithm.

We next show conversely that this is so by necessity, namely that from an approximate uniform generation method we obtain approximations to the r_i's.

So let $\epsilon > 0$ and A be a probabilistic algorithm that outputs a point on $\mathcal{C}$ such that for each $P \in \mathcal{C}$ the probability α that P is output satisfies

$$| \alpha - \frac{1}{\#\mathcal{C}} | \leq \epsilon.$$

Now let $0 \leq i \leq n$ and $k \in \mathbb{N}$. We run A for k times, count the number t of times that the output $P = (a, b)$ satisfies $a \in R_i$, using (2.1), and return $\rho_i = t/k$. The general definition of an (ϵ, δ)-approximation requires, in our case, that

$$\text{prob } (| r_i - \rho_i | \leq \epsilon \, r_i) \geq 1 - \delta.$$

The estimator theorem of Karp *et al.* (1989) says that we have such a scheme if we choose

$$k = \lceil 4\beta_i \log_e(2/\delta)\epsilon^{-2}\rceil,$$

where $\beta_i \geq \#\mathcal{C}/r_i$. Recall that under the assumptions of Theorem 2.4 we have $\#\mathcal{C} \geq q/2$, and either $r_i \geq q/2n!$ or $r_i \leq 2n^{2n}q^{1/2}$. Thus if we use $\beta_i = n!$, then in the first case we obtain an (ϵ, δ)-approximation scheme for r_i, and in the second we expect to find no $a \in R_i$.

Since

$$\frac{\#\mathcal{C}}{n} \leq \sum_{1 \leq i \leq n} r_i \leq \sum_{1 \leq i \leq n} ir_i = \#\mathcal{C},$$

the r_i's are on average at least $\#\mathcal{C}/n^2$. To find approximations only to the "large" r_i's, we might use $\beta_i = \lambda n^2$, with some small number λ. The dependence on q of this algorithm is only $O(\log q)$, while stand-alone methods, such as in Theorem 3.4 of von zur Gathen *et al.* (1993), not using our assumed uniform random generator, use $\Omega(q^\alpha)$ operations for some $\alpha > 0$.

3. Deterministic construction of all points

In this section, we present a deterministic algorithm for finding all points on $\mathcal{C} = \{f = 0\}$ over a prime field $\mathbb{F}_p$. It employs a deterministic polynomial–time algorithm for finding all roots of the univariate polynomials $f(a, y)$, with $a \in \mathbb{F}_p$. This algorithm does not factor $f(a, y)$ completely for all a, but we show that there are only about $\sqrt{p}$ exceptional a, and for these we use an always successful deterministic algorithm with time about $\sqrt{p}$; thus the total time is proportional to p, which is about the size of $\mathcal{C}$. Everything is polynomial in the degree n.

As a first step, we factor f into irreducible factors in $\mathbb{F}_p[x, y]$. The bivariate factoring algorithms (Lenstra 1985, von zur Gathen 1984, von zur Gathen & Kaltofen 1985) can actually be made into deterministic reductions from bivariate to univariate factorization over finite fields. Thus f can be factored with $n^{O(1)}p^{1/2}$ operations in $\mathbb{F}_q$. From now on, we assume that f is irreducible.

The projection $\pi : \mathcal{C} = \{f = 0\} \to \mathbb{F}_p$ onto the first coordinate is called *separable* if and only if $h_y = \partial h/\partial y \neq 0$ for each irreducible factor $h \in \mathbb{F}[x, y]$ of f. A simple example of an inseparable projection is given by $f = x - y^p \in \mathbb{F}_p[x, y]$. The curve $\mathcal{C} = \{x = y^p\}$ is smooth, and all tangents to $\mathcal{C}$ are vertical.

Let $\varphi: \mathbb{F}_q \to \mathbb{F}_q$ denote the absolute Frobenius map, with $\varphi(a) = a^p$. (We will only use $q = p$.) For our algorithms, it is convenient to have π separable, and the next lemma describes a simple procedure for achieving this by factoring out φ.

LEMMA 3.1. *Let $f \in \mathbb{F}_q[x,y]$ be irreducible. We can compute in polynomial time $g \in \mathbb{F}_q[x,y]$ and an integer $k \le \log_q(\deg_y f)$ such that*

$$id \times \varphi^k: \mathbb{F}^2 \longrightarrow \mathbb{F}^2$$

gives a bijection between $\{f = 0\}$ and $\{g = 0\}$, $\deg_x g = \deg_x f, \deg_y g \le \deg_y f$, and $\pi: \{g = 0\} \to \mathbb{F}_q$ is separable.

PROOF. We write $f = \sum_{i,j} f_{ij} x^i y^j$, with each $f_{ij} \in \mathbb{F}_q$, and $p = \mathrm{char}\mathbb{F}_q$. Then

$$f_y = 0 \Longleftrightarrow \forall i, j \ (f_{ij} \ne 0 \Rightarrow p \mid j).$$

If $f_y = 0$ and

$$h = \sum_{\substack{i,j \\ p|j}} f_{ij} x^i y^{j/p},$$

then $f(a,b) = h(a, b^p)$ for all $(a,b) \in \mathbb{F}_q^2$, and thus $\mathrm{id} \times \varphi: \mathbb{F}_q^2 \to \mathbb{F}_q^2$ gives a bijection between $\{f = 0\}$ and $\{h = 0\}$. We repeat this process until we obtain a polynomial $g \in \mathbb{F}_q[x,y]$ and $k \in \mathbb{N}$ with $g_y \ne 0$ and $\mathrm{id} \times \varphi^k$ a bijection between $\{f = 0\}$ and $\{g = 0\}$. $\square$

ALGORITHM 3.2. *Finding all points.*
Input: *$f \in \mathbb{F}_p[x,y]$ of degree n, where p is a prime.*
Output: *A list of all points $(a,b) \in \mathbb{F}_p^2$ with $f(a,b) = 0$.*

1. *Set $h = 288 n^4 \lceil \log_2 p \rceil^2$.*
2. *For all $a \in \mathbb{F}_p$ do steps 3 through 7.*
3. *Compute $f_a = f(a,y) \in \mathbb{F}_p[y]$.*
4. *Compute $f_a^* = \gcd(y^p - y, f_a) \in \mathbb{F}_p[y]$.*
5. *For $0 \le t < h$ compute the two factors*

$$g_{a,t} = \gcd\Big((y-t)^{(p-1)/2} - 1, f_a^*\Big), \ g_{a,t}^* = \gcd(y - t, f_a^*) \in \mathbb{F}_p[y]$$

of f_a^*.

6. *Compute the common refinement of the partial factorizations from step 5.*

7. *If step 6 returns only linear factors $y - b$, then add all these (a, b) to the list. Otherwise completely factor f_a^* with the deterministic algorithm of von zur Gathen & Shoup (1992), and add all resulting (a, b) to the list.*

THEOREM 3.3. *Let p be a prime, $f \in \mathbb{F}_p[x, y]$ squarefree and non-exceptional, and $\pi: \mathcal{C} = \{f = 0\} \to \mathbb{F}_p$ separable. Then the algorithm correctly computes all points on $\mathcal{C}$. It uses*

$$O(n^5 p \log n \operatorname{loglog} n \log(np) \log^2 p)$$

or $O\tilde{}(n^5 p)$ operations in $\mathbb{F}_p$.

PROOF. For all $a, b \in \mathbb{F}_p$ we have

$$f(a, b) = 0 \Longleftrightarrow f_a^*(b) = 0 \Longleftrightarrow y - b \mid f_a^*.$$

Since step 7 returns all linear factors of f_a^*, the final list correctly contains all points of $\mathcal{C} = \{f = 0\}$.

It remains to analyze the running time. The crucial point is to understand when step 6 succeeds in completely factoring f_a^*. Denote by $S \subseteq \mathbb{F}_p$ the set of all a for which this is not the case, and $s = \#S$. Furthermore, $\mathcal{C}_a = \pi(\mathcal{C} \cap (\{a\} \times \mathbb{F}_p))$ consists of all $b \in \mathbb{F}_p$ with $(a, b) \in \mathcal{C}$. Thus

$$S = \{a \in \mathbb{F}_p : \exists b, c \in \mathcal{C}_a\ b \neq c; b, c \geq h, \text{ and } \forall t < h\ (y - b \mid g_{a,t} \Longleftrightarrow y - c \mid g_{a,t})\}.$$

Furthermore, let $\mathsf{M}(n)$ be a multiplication time, i.e., such that the product of two polynomials in $\mathbb{F}_p[x]$ of degree at most n can be computed with $O(\mathsf{M}(n))$ operations in $\mathbb{F}_p$. Then we can choose $\mathsf{M}(n) = n \log n \operatorname{loglog} n$, and a gcd can be computed with $O(\mathsf{M}(n) \log n)$ operations. The refinement cost in step 6, if done along a binary tree, is $O(\mathsf{M}(n) \log n)$ for each t, or $O(h\, \mathsf{M}(n) \log n)$ in total. For $a \in S$, an application of the algorithm from von zur Gathen & Shoup (1992) costs $O(\mathsf{M}(n)\, p^{1/2} \log(np))$ operations in $\mathbb{F}_p$. The gcds in steps 4 and 5 are computed by repeated squaring for the required power of y and $y - t$, reducing after each multiplication modulo f_a and f_a^*, respectively.

For each a in step 2, we find the following number of operations in $\mathbb{F}_p$:

- step 3: $O(n^2)$,
- step 4: $O(\mathsf{M}(n) \log(np))$,

- step 5: $O(h\,\mathsf{M}(n)\log(np))$,
- step 6: $O(h\,\mathsf{M}(n)\log n)$,
- step 7: 0 if $a \in \mathbb{F}_p \setminus S$, and $O(\mathsf{M}(n)\,p^{1/2}\log(np))$ if $a \in S$.

The total cost is

$$O\Big(p \cdot (n^2 + n^4 \mathsf{M}(n)\log(np)\log^2 p) + s\,\mathsf{M}(n)p^{1/2}\log(np)\Big) \tag{3.1}$$

operations, and we now show that s is $O(n^2(n^2+\log p)p^{1/2})$. This will imply the claim about the running time. We let

$$Q = \{u \in \mathbb{F}_p^\times : \exists v \in \mathbb{F}_p^\times \; u = v^2\} = \{u \in \mathbb{F}_p^\times : u^{(p-1)/2} = 1\}$$

be the set of nonzero squares in $\mathbb{F}_p$, and χ the quadratic character on $\mathbb{F}_p$, with

$$\chi(b) = \begin{cases} 1, & \text{if } b \in Q, \\ -1, & \text{if } b \notin Q,\ b \neq 0, \\ 0, & \text{if } b = 0. \end{cases}$$

For the time being, we work with an arbitrary integer parameter h; only at the end will we substitute the value from step 1. Set $H = \{0, \ldots, h-1\} \subseteq \mathbb{F}_p$, where we identify $\mathbb{F}_p$ with $\{0, \ldots, p-1\}$. Two distinct elements $b, c \in \mathbb{F}_p$ are *h-separated* if and only if $\chi(b-t) \neq \chi(c-t)$ for some $t \in H$. A set $B \subseteq \mathbb{F}_p$ is *h-separated* if any two distinct elements of B are.

With this notation, we have for $a \in \mathbb{F}_p$

$$a \in S \Longrightarrow \mathcal{C}_a \text{ is not } h\text{-separated.}$$

The reverse implication is true if the non-h-separated $b, c \in \mathcal{C}_a$ are both at least h. If $a \in S$, then for at least one pair of distinct elements $b, c \in \mathcal{C}_a$,

$$h = \sum_{0 \le t < h} \chi\Big((t-b)(t-c)\Big).$$

Now we let $k \in \mathbb{N}$ and

$$\begin{aligned} w &= \sum_{a \in \mathbb{F}_p} \sum_{\substack{b,c \in \mathcal{C}_a \\ b \neq c}} \Big| \sum_{0 \le t < h} \chi\Big((t-b)(t-c)\Big)\Big|^{2k} \\ &= \sum_{0 \le t_1, \ldots, t_{2k} < h} \; \sum_{a \in \mathbb{F}_p} \sum_{\substack{b,c \in \mathcal{C}_a \\ b \neq c}} \chi\Big((t_1-b)(t_1-c)\cdots(t_{2k}-b)(t_{2k}-c)\Big). \end{aligned}$$

Then, by the above, $sh^{2k} \le w$. We consider the set

$$\mathcal{D}_0 = \{(a,b,c) \in \mathbb{F}_p^3 \colon f(a,b) = f(a,c) = 0, b \neq c\} \subseteq \mathbb{F}_p^3.$$

The fibre product $\mathcal{D} = \mathcal{C} \times_\pi \mathcal{C}$ is the closure of $\mathcal{D}_0$ in $\mathbb{F}_p^3$; it has degree at most $n(n-1) < n^2$ and is discussed in detail in Section 4. Then

$$w = \sum_{t \in H^{2k}} \sum_{P \in \mathcal{D}} \chi\Big(\psi_t(P)\Big),$$

where the inner sum is over all $\mathbb{F}_p$-rational points $P = (a,b,c) \in \mathcal{D}$ with $b \neq c$, ψ_t is the polynomial

$$\psi_t = (y - t_1) \cdots (y - t_{2k})(z - t_1) \cdots (z - t_{2k}) \in \mathbb{F}_p[y,z]$$

in indeterminates y and z, and $\psi_t\Big((a,b,c)\Big)$ is obtained by substituting b and c for y and z, respectively.

Theorem 4.4 below says that there are at most $(12kn^2h^{1/2})^{2k}$ values of $t \in H^{2k}$ for which $\rho(\psi_t)$ is a square in the global ring $\mathcal{O}_\mathcal{A}$ of some irreducible component $\mathcal{A} \subseteq \mathbb{F}^3$ of $\mathcal{D}$, where $\rho \colon \mathbb{F}[x,y,z] \longrightarrow \mathcal{O}_\mathcal{A}$ is the restriction map.

For other vectors $t \in \mathbb{F}^{2k}$, we may apply the bound on character sums along a curve from Perel'muter (1969) that gives

$$\sum_{P \in \mathcal{D}} \chi\Big(\psi_t(P)\Big) \le d \cdot \Big(n^2(n^2 + 2k)p^{1/2}\Big) \tag{3.2}$$

for some constant d. Perel'muter's bound holds for each irreducible component of $\mathcal{D}$; we also use the fact that no such component is vertical (see proof of Theorem 3.4 of von zur Gathen *et al.* 1993). Since their degrees sum to $\deg \mathcal{D} < n^2$, (3.2) follows. Therefore

$$w \le (12kn^2h^{1/2})^{2k}p + d \cdot n^2(n^2+2k)h^{2k}p^{1/2},$$

$$s \le (12kn^2h^{-1/2})^{2k}p + d \cdot n^2(n^2+2k)p^{1/2}.$$

Now, using $k = \lceil \log_2 p \rceil$ and h as in step 1 of Algorithm 3.2, we find

$$(12kn^2h^{-1/2})^{2k} \le 2^{-k} \le 2^{-\log_2 p} = p^{-1}.$$

Hence

$$s = O\Big(n^2(n^2 + \log p)p^{1/2}\Big).$$

Together with (3.1), this proves the estimate of the total cost. □

4. Squares on the fibre product

Let $\mathbb{F}$ be an algebraically closed field, $f \in \mathbb{F}[x,y]$ squarefree of degree $n \geq 1$, $\mathcal{C} = \{f = 0\} \subseteq \mathbb{F}^2$ the associated plane curve, and $\pi : \mathcal{C} \longrightarrow \mathbb{F}$ the first projection. We assume that π is separable. Then $\mathcal{D} = \mathcal{C} \times_\pi \mathcal{C} \subseteq \mathbb{F}^3$, the fibre square over π, can be defined as the closure in $\mathbb{F}^3$ of

$$\mathcal{D}_0 = \{(a,b,c) \in \mathbb{F}^3 : f(a,b) = f(a,c) = 0, b \neq c\}.$$

Furthermore, let $g = (f(x,y) - f(x,z))/(y-z) \in \mathbb{F}[x,y,z]$.

A smooth point $P = (a,b) \in \mathcal{C}$ is *critical* for π if and only if the tangent line $T_{P,\mathcal{C}}$ in $\mathbb{F}^2$ is vertical. If f is irreducible, this is equivalent to $f_y(a,b) = 0$, where $f_y = \partial f/\partial y \in \mathbb{F}[x,y]$; in general, we have to replace f by its (unique) irreducible factor on whose component P lies. Since π is separable, $\mathcal{C}$ has only finitely many critical points.

THEOREM 4.1. *Let $f \in \mathbb{F}[x,y]$ be squarefree and π separable.*

(i) $\mathcal{D} = \{f(x,y) = g = 0\}$.

(ii) $\mathcal{D} = \mathcal{D}_0 \cup \{(a,b,b) : (a,b) \in \mathcal{C}$ *is singular or critical*$\}$.

(iii) $(a,b,c) \in \mathcal{D}$ *with* $b \neq c$ *is singular on* $\mathcal{D}$ *if and only if either* (a,b) *or* (a,c) *is singular on* $\mathcal{C}$*, or both* (a,b) *and* (a,c) *are critical on* $\mathcal{C}$*. All points of* $\mathcal{D} \setminus \mathcal{D}_0$ *are singular on* $\mathcal{D}$.

(iv) $\deg \mathcal{D} \leq n(n-1) < n^2$.

PROOF. Let $\Delta = \{(a,b,b) \in \mathbb{F}^3 : a, b \in \mathbb{F}\}$ be the diagonal. Clearly $\mathcal{D} \setminus \Delta = \mathcal{D}_0$, and $\mathcal{D}_0 \subseteq \mathcal{D}_1 = \{f(x,y) = g = 0\}$. By definition, $\mathcal{D}$ is the closure of $\mathcal{D}_0$, and thus $\mathcal{D} \subseteq \mathcal{D}_1$. We prove that (ii) is valid with $\mathcal{D}_1$ instead of $\mathcal{D}$. Thus $\mathcal{D}_1 \cap \Delta$ is finite, and $\mathcal{D} = \mathcal{D}_1$ follows, hence (i), (ii), and (iv).

So let u, v be indeterminates over $\mathbb{F}[x,y]$. Then the Taylor expansion of f around (u,v) of order 1 is

$$f(x,y) = f(u,v) + f_x(u,v)(x-u) + f_y(u,v)(y-v) + h$$

in $\mathbb{F}[x,y,u,v]$, with some $h \in (x-u, y-v)^2$. Therefore

$$\begin{aligned} & g(x,y,z) \\ = \;& \frac{1}{y-z} \cdot (f_y(u,v)(y-v) - f_y(u,v)(z-w) + h(x,y,u,v) - h(x,z,u,v)) \\ = \;& f_y(u,v) + H, \end{aligned}$$

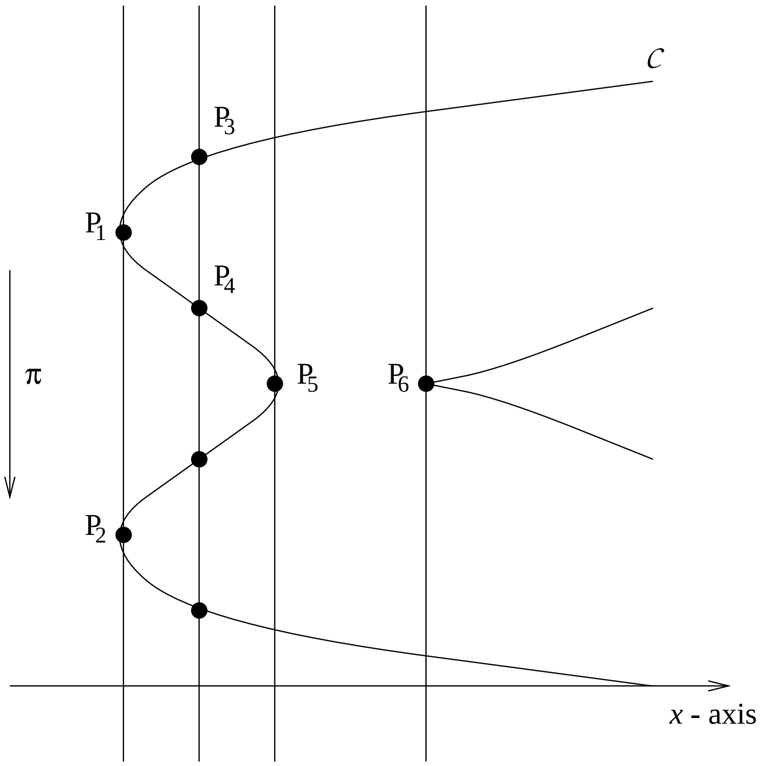

Figure 4.1: P_1, P_2, P_5 are critical for π, and P_6 is singular on $\mathcal{C}$. If $P_i = (a_i, b_i)$, then $(a_i, b_i, b_i) \in \mathcal{D} \cap \Delta$ for $i = 1, 2, 5, 6$. These four points are singular on $\mathcal{D}$. Furthermore, $(a_1, b_1, b_2) \in \mathcal{S} \subseteq \mathcal{D}$, and $(a_3, b_3, b_4) \in \mathcal{D} \setminus \mathcal{S}$.

with some $H \in (x-u, y-v, z-v)$. Thus for $(a, b) \in \mathcal{C}$

$$(a, b, b) \in \mathcal{D}_1 \iff f_y(a, b) = 0 \iff (a, b) \text{ is singular or critical on } \mathcal{C}.$$

For (iii), let $(a, b, c) \in \mathcal{D}$ with $b \neq c$. The Jacobian of $\mathcal{D}$ at (a, b, c) is

$$J(a, b, c) = \begin{pmatrix} f_x(a, b) & \dfrac{f_x(a, b) - f_x(a, c)}{b - c} \\ f_y(a, b) & \dfrac{f_y(a, b)}{b - c} \\ 0 & \dfrac{-f_y(a, c)}{b - c} \end{pmatrix}.$$

After multiplying the second column by $c - b$ and then adding the first column to the second, we obtain the matrix

$$A = \begin{pmatrix} f_x(a,b) & f_x(a,c) \\ f_y(a,b) & 0 \\ 0 & f_y(a,c) \end{pmatrix}.$$

Thus

$$\begin{aligned} &(a,b,c) \text{ is singular on } \mathcal{D} \iff \text{rank } (J(a,b,c)) \le 1 \\ &\iff \text{rank } (A) \le 1 \\ &\iff (a,b) \text{ or } (a,c) \text{ is singular on } \mathcal{C}, \text{ or both are critical on } \mathcal{C}. \quad \Box \end{aligned}$$

The condition that π be separable is necessary, since otherwise all points on $\mathcal{C}$ are critical. Recall the example $\mathcal{C} = \{x = y^p\}$, where $p = \operatorname{char} \mathbb{F}$, from Section 3. Then $f_y = 0$, $\mathcal{C}$ is smooth, and all tangent lines to $\mathcal{C}$ are vertical. Furthermore, $\mathcal{D}_0 = \emptyset$, $g = (y^p - z^p)/(y-z) = (y-z)^{p-1}$, and $\mathcal{C} \times_\pi \mathcal{C}$ equals $\{(a,b,b) \in \mathbb{F}^3 \colon a = b^p\}$, counted $p-1$ times. On the other hand, when $\mathcal{C} = \{y = g(x)\}$ is the graph of a polynomial $g \in \mathbb{F}_q[x]$, then π is separable, and $\mathcal{D} = \emptyset$.

We define

$$\mathcal{S} = \{(a,b,c) \in \mathcal{D} \colon (a,b) \text{ or } (a,c) \text{ is singular or critical on } \mathcal{C}\}.$$

We now let $\mathcal{A}$ be an irreducible component of $\mathcal{D}$, and want to estimate the number of t such that $\psi_t = \prod_{1 \le i \le 2k} (t_i - y)(t_i - z)$ is a square in $\mathcal{O}_\mathcal{A}$. We let $\rho \colon \mathbb{F}[x,y,z] \longrightarrow \mathcal{O}_\mathcal{A}$ be the restriction map.

Let $t \in \mathbb{F}^{2k}$, and $T = \{1, \ldots, 2k\}$. The overall goal of this section is to show in Theorem 4.4 that only few $\rho(\psi_t)$ are squares, when t is chosen from a finite subset H of $\mathbb{F}^{2k}$. For a simple example of a square, we take the parabola $f = x - y^2$, so that $\mathcal{C} = \{x = y^2\}$, and $\mathcal{D} = \{x - y^2 = y + z = 0\}$ is irreducible. If $k = 1$ and $t_2 = -t_1$, then

$$\rho(\psi_t) = \rho((t_1 - y)(t_1 - z)(t_2 - y)(t_2 - z)) = \rho((t_1 - y)^2 (t_1 + y)^2) \qquad (4.1)$$

is a square on $\mathcal{D}$.

The condition that $\rho(\psi_t)$ not be a square for (3.2) to hold is not an artifact of Perel'muter's proof, but without it (3.2) may actually fail to be true.

In the sequel, we define several combinatorial objects on the index set T. We first collect pairs of equal values of t_i in a systematic way. Namely, we take the lexicographically first maximal matching on the directed graph with vertex set T, and where (i,j) are connected if and only if $i < j$ and $t_i = t_j$. Then $T_1 \subseteq T$ is defined as the set of these first coordinates i, and $\tau_1 \colon T_1 \to T$

is defined by $\tau_1(i) = j$ if (i, j) occurs in that matching. As an example, if $t_3 = t_5 = t_8 = t_{11} = t_{13}$ and no other t_i equals these, then $T_1 = \{3, 5\}$, $\tau_1(3) = 8$, and $\tau_1(5) = 11$.

Next, we set

$$T_2 = \{i \in T \setminus (T_1 \cup \tau_1(T_1)) \colon \mathcal{A} \cap \{y = t_i\} \subseteq \mathcal{S} \text{ or } \mathcal{A} \cap \{z = t_i\} \subseteq \mathcal{S}\}.$$

Then the t_i for

$$i \in T_3 = T \setminus (T_1 \cup \tau_1(T_1) \cup T_2)$$

are pairwise distinct, and $(T_1, \tau_1(T_1), T_2, T_3)$ is a partition of T. Next, we let

$$S_0 = T_3 \times \{0\}, \quad S_1 = T_3 \times \{1\}$$

be two disjoint copies of T_3, and define a bipartite undirected graph $G = (S_0 \cup S_1, E)$ as follows. For $i, j \in T_3$, $(i, 0)$ and $(j, 1)$ are connected in G if and only if there is some $(a, b, c) \in \mathcal{A} \setminus \mathcal{S}$ such that $b = t_i$ and $c = t_j$.

In the example (4.1) of a parabola, we have $T_1 = T_2 = \emptyset$, and

$$G = \begin{tikzcd} (1,0) \arrow[dr, dash] & (1,1) \\ (2,0) \arrow[ur, dash] & (2,1) \end{tikzcd} .$$

Lemma 4.2. *If $t \in \mathbb{F}^{2k}$ is such that $\rho(\psi_t) \in \mathcal{O}_{\mathcal{A}}$ is a square, then each vertex in G has degree at least one.*

Proof. By symmetry, it is sufficient to show the claim for a vertex $(i, 0) \in S_0$.

Since $i \notin T_2$, we can choose some $P = (a, t_i, c) \in \mathcal{A} \setminus \mathcal{S}$; then $c \neq t_i$. Let

$$U_0 = \{j \in T \colon t_j = t_i\}, \quad U_1 = \{j \in T \colon t_j = c\},$$

$\rho \colon \mathbb{F}[x, y, z] \to \mathcal{O}_{\mathcal{A}}$ the restriction to $\mathcal{A}$, $\mathcal{R} = \mathcal{O}_{P,\mathcal{A}}$ the local ring at P, which is a Unique Factorization Domain (see e.g. Shafarevich 1974, Theorem II.3.2), and $\lambda = (\mathcal{O}_{\mathcal{A}} \to \mathcal{O}_{P,\mathcal{A}}) \circ \rho$ the composition of ρ with the localization at P. Then $i \in U_0$, and $U_0, U_1 \subseteq T \setminus T_2$.

For every $j \in T \setminus (U_0 \cup \tau_1(U_0) \cup \{i\})$, we have $t_j \neq t_i$, and thus $\lambda(y - t_j)$ is a unit in $\mathcal{R}$. Similarly, each $\lambda(z - t_j)$ with $t_j \neq c$ is a unit in $\mathcal{R}$.

Since $(a, t_i) \in \mathcal{C}$ is not critical for π, we have $f_y(a, t_i) \neq 0$, and therefore $\lambda(y - t_i) \in \mathcal{R}$ is a local parameter in $\mathcal{R}$. Similarly, each $\lambda(z - t_j)$ with $t_j = c$ is a local parameter in $\mathcal{R}$.

By the above, there is a unit $u \in \mathcal{R}$ such that

$$\begin{aligned}\lambda(\psi_t) &= \prod_{j\in T} \lambda(y-t_j)\cdot \prod_{j\in T}\lambda(z-t_j)\\ &= u\cdot \prod_{j\in U_0\cup\tau_1(U_0)\cup\{i\}} \lambda(y-t_j)\cdot \prod_{j\in U_1}\lambda(z-t_j)\end{aligned}$$

is a square in $\mathcal{R}$. Thus the total number of local parameters in the product is even. We have $\#U_0 = \#\tau_1(U_0)$ and $i \notin U_0 \cup \tau_1(U_0)$. It follows that in the left hand product, the number of local parameters is odd, and therefore also in the right hand product. Thus there exists some $j \in T_3$ with $t_j = c$; then $\{(i,0),(j,1)\} \in E$. □

We now take a maximal "disjoint" matching (V_0, V_1) in G of the following type. The sets $V_0, V_1 \subseteq T_3$ are disjoint, G induces a perfect matching on $(V_0 \times \{0\}) \cup (V_1 \times \{1\})$, and this matching is maximal. Furthermore, let $\mu : V_0 \longrightarrow V_1$ be the corresponding bijection, with $\mu(i) = j$ if and only if $\{(i,0),(j,1)\}$ occurs in the matching.

For every $i \in V_2 = T_3 \setminus (V_0 \cup V_1)$, $(i,0)$ is connected to some $(j,1) \in T_3 \times \{1\}$, and by the maximality of the matching, we have $j \in V_0 \cup V_1$. We take $\mu: V_2 \longrightarrow V_0 \cup V_1$ such that $\mu(i) = j$ for some such j, and note that (V_0, V_1, V_2) is a partition of T_3.

Finally, we indicate how to describe t_i for $i \in V_0$ succinctly if $\{(i,0),(j,1)\} \in E$ and t_j is known. For this, we take an arbitrary total order $\prec$ on $\mathbb{F}$. For each $t \in \mathbb{F}$, $\mathcal{C} \cap \{y = t\}$ has at most n points, say $(a_1,t),\ldots,(a_l,t)$ with $l \le n$ and $a_1 \prec \cdots \prec a_l$. If $j = \mu(i)$ and $t = t_j$, then $(a_r, t_i, t_j) \in \mathcal{D} \setminus \mathcal{S}$ for one of those points, with $1 \le r \le l$. We choose the smallest such r; then $\mathcal{C} \cap \{x = a_r\}$ consists again of at most n points. We let v be the position of (a_r, t_i) in this list, ordered according to $\prec$, and set $\tau_3(i) = (r, v)$. Then t_i is determined by $j = \mu(i)$, t_j, and $\tau_3(i)$.

Similarly, we define $\tau_3: V_2 \longrightarrow \{1,\ldots,n\}^2$ so that for $i \in V_2$, t_i is determined by $j = \mu(i)$, t_j, and $\tau_3(i)$.

We have thus associated to any $t \in \mathbb{F}^{2k}$ with $\rho(\psi_t)$ a square the following data:

$$T_1, \tau_1, T_2, V_0, \mu, \tau_3, \text{ and } t_i \text{ for } i \in T_1 \cup T_2 \cup V_1. \tag{4.2}$$

Lemma 4.3. *If $\rho(\psi_t)$ is a square in $\mathcal{O}_{\mathcal{A}}$, then t is determined by the data in (4.2).*

Proof. $(T_1, \tau_1(T_1), T_2, V_0, V_1, V_2)$ is a partition of T, and $t_i = t_{\tau_2(i)}$ for each $i \in T_1$. Thus it remains to show that each t_i with $i \in V_0 \cup V_2$ is determined by (4.2). But that is precisely what the construction of μ and τ_3 achieves. □

We are now ready for the main result of this section, an upper bound on the number of ψ_t which are squares. The bound is rather coarse, but sufficient for our purposes.

THEOREM 4.4. *Let $\mathbb{F}$ be an algebraically closed field, $f \in \mathbb{F}[x, y]$ squarefree, $\mathcal{C} = \{f = 0\}$ with $\pi: \mathcal{C} \to \mathbb{F}$ separable and $H \subseteq \mathbb{F}$ be a finite set with h elements. The number of $t \in H^{2k}$ such that $\rho(\psi_t)$ is a square in $\mathcal{O}_{\mathcal{A}}$ for some irreducible component $\mathcal{A}$ of $\mathcal{C} \times_\pi \mathcal{C}$ is at most $(12kn^2h^{1/2})^{2k}$.*

PROOF. We first fix a component $\mathcal{A}$ of $\mathcal{D}$, and show the corresponding bound. By Lemma 4.3, it is sufficient to give an upper bound on the number of choices for the data in (4.2).

The six sets $T_1, \tau_1(T_1), T_2, V_0, V_1, V_2$ form a partition of T, and there are at most 6^{2k} choices for this partition.

Suppose that these sets are chosen, with cardinalities $c_1, c_2, c_3, c_4, c_5, c_6$, respectively. Then $c_1 = c_2$, $c_3 < n^2$, and $c_4 = c_5$. The number of choices for τ_1 is at most $(2k)^{c_1}$, for μ at most $(2k)^{c_4+c_6}$, for τ_3 at most $(n^2)^{c_4+c_6}$, and for all t_i's required in (4.1) at most $h^{c_1+c_5} \cdot (n^2)^{c_3}$. Since $c_1 + c_5 \leq 2k/2 = k$, the total comes to

$$m = 6^{2k} \cdot (2k)^{c_1+c_4+c_6} \cdot (n^2)^{c_3+c_4+c_6} \cdot h^{c_1+c_5}. \tag{4.3}$$

Since $\deg \mathcal{D} \leq n(n-1)$ by Theorem 4.1 (i), $\mathcal{D}$ has at most $n(n-1) < n^2$ irreducible components. So the total number of t considered is at most n^2m, and

$$n^2m \leq 6^{2k} \cdot (2k)^{2k} \cdot (n^2h^{1/2})^{2k}.$$

Here we use that either $c_1 + c_2 + c_5 > 0$ and then $n^2 \cdot (n^2)^{c_3+c_4+c_6} \leq (n^2)^{2k}$, or $c_2 + c_3 + c_4 + c_6 > 0$ and then $n^2(h)^{c_1+c_5} \leq h^k$. $\square$

5. Exceptional polynomials

In this section, we deal with the somewhat troublesome case excluded so far: exceptional polynomials, for which $\sigma = 0$. No analogue of the deterministic result of Theorem 3.3 is known for them, while the probabilistic results of Section 2 carry over easily.

We first note that it is not surprising that they are difficult to deal with, since any subset of $\mathbb{F}_q^2$ is an exceptional curve. If $c \in \mathbb{F}_q$ is a nonsquare and $f = x^2 + cy^2$, then f is exceptional and

$$\{f = 0\} = \{(0,0)\}, \tag{5.1}$$

and by translation and finite unions the claim follows. If char $\mathbb{F}_q \geq 3$, then (5.1) also holds for $f = x^{q-1} + y^{q-1}$. If $b \in \mathbb{F}_q^2 \setminus \mathbb{F}_q$ with $b^2 \in \mathbb{F}_q$, then $b^{q-1} = (b^2)^{(q-1)/2} = -1$. Thus f is the product of all $x - by$ with these b, and thus f is exceptional, too.

Now given an arbitrary $f \in \mathbb{F}_q[x,y]$ of degree n, there are well-known probabilistic algorithms with time polynomial in $n \log q$ that factor f into its irreducible factors over $\mathbb{F}_q$ (von zur Gathen & Kaltofen 1985) and test each such factor for absolute irreducibility (Kaltofen 1985). For simplicity, assume now that f is irreducible over $\mathbb{F}_q$, and not absolutely irreducible. Then Kaltofen's algorithm can be used to find a field extension K of $\mathbb{F}_q$ with $[K\colon \mathbb{F}_q] \leq n$ and a proper factorization of f over K. If g and h are two distinct factors, then the first coordinate of any common root is a root of

$$\mathrm{res}_y(g,h) \in K[x].$$

Thus it is easy to calculate all common roots of g and h, to check which ones are in $\mathbb{F}_q^2$, and to determine whether they are indeed roots of f. All roots of f are found in this way; there are at most $n^2/4$ of them (von zur Gathen *et al.* 1993).

Theorem 5.1. *Let $f \in \mathbb{F}_q[x,y]$ have degree n. There is a probabilistic algorithm using $(n \log q)^{O(1)}$ operations in $\mathbb{F}_q$ that determines whether f is exceptional and, if it is, finds all points of $\{f = 0\}$.*

Acknowledgements

Parts of this work were done on visits of the first author to Macquarie University, and of the second author to Universität Paderborn. The hospitality of both institutions is gratefully acknowledged.

References

D. G. Cantor and H. Zassenhaus, A new algorithm for factoring polynomials over finite fields. *Math. Comp.* **36** (1981), 587–592.

J. von zur Gathen, Hensel and Newton methods in valuation rings. *Math. Comp.* **42** (1984), 637–661.

J. von zur Gathen and E. Kaltofen, Factorization of multivariate polynomials over finite fields. *Math. Comp.* **45** (1985), 251–261.

J. von zur Gathen and V. Shoup, Computing Frobenius maps and factoring polynomials. *Comput complexity* **2** (1992), 187–224.

J. von zur Gathen and I. E. Shparlinski, Components and projections of curves over finite fields. In *Proc. 5th Int. Symp. on Algorithms and Computation ISAAC '94*, vol. 834 of *Springer Lecture Notes in Computer Science*, 1994, 297–305.

J. von zur Gathen, M. Karpinski, and I. E. Shparlinski, Counting curves and their projections. In *Proc. 25th ACM Symp. Theory of Computing*, 1993, 805–812.

E. Kaltofen, Fast parallel absolute irreducibility testing. *J. Symb. Computation* **1** (1985), 57–67.

R. M. Karp, M. Luby, and N. Madras, Monte-Carlo approximation algorithms for enumeration problems. *J. Algorithms* **10**(3) (1989), 429–448.

A. K. Lenstra, Factoring multivariate polynomials over finite fields. *J. Comput. System Sci.* **30** (1985), 235–248.

G. I. Perel'muter, Оценка суммы вдоль алгебраическои кривои (Bounds on sums along algebraic curves). *Mat. Zametki* **5** (1969), 373–380.

I. R. Shafarevich, *Basic algebraic geometry*. Grundlehren Band 213. Springer Verlag, 1974.

V. Shoup, On the deterministic complexity of factoring polynomials over finite fields. *Inform. Process. Lett.* **33** (1990), 261–267.

I. E. Shparlinski, *Computational and algorithmic problems in finite fields*, vol. 88 of *Mathematics and its applications*. Kluwer Academic Publishers, 1992.